

DEPARTMENT  
DEPARTMENT  
DEPARTMENT  
SCIENCE  
SCIENCE  
SCIENCE  
COMPUTER  
COMPUTER  
COMPUTER



*Updating Derived Relations:  
Detecting Irrelevant  
and Autonomously  
Computable Updates*

UNIVERSITY OF WATERLOO  
UNIVERSITY OF WATERLOO  
UNIVERSITY OF WATERLOO

*José A. Blakeley  
Neil Coburn  
Per-Åke Larson*

*CS-86-17*

*May, 1986*

# Updating Derived Relations: Detecting Irrelevant and Autonomously Computable Updates\*

José A. Blakeley, Neil Coburn, and Per-Åke Larson†

Data Structuring Group  
Department of Computer Science  
University of Waterloo  
Waterloo, Ontario, Canada

## ABSTRACT

Consider a database containing not only base relations but also stored derived relations (also called materialized or concrete views). When a base relation is updated, it may also be necessary to update some of the derived relations. This paper gives sufficient and necessary conditions for detecting when an update of a base relation cannot affect a derived relation (an irrelevant update), and for detecting when a derived relation can be correctly updated using no data other than the derived relation itself and the given update operation (an autonomously computable update). The class of derived relations considered is restricted to those defined by *PSJ*-expressions, that is, any relational algebra expression constructed from an arbitrary number of project, select and join operations. The class of update operations consists of insertions, deletions, and modifications, where the set of tuples to be deleted or modified is specified by a *PSJ*-expression.

## 1. Introduction

In a relational database system, the database may contain *derived relations* in addition to base relations. A derived relation is defined by a relational expression (query) over the base relations. A derived relation may be *virtual*, which corresponds to the traditional concept of a view, or *materialized*, meaning that the relation resulting from evaluating the expression over the current database instance is actually stored. In the sequel all derived relations are assumed to be materialized, unless stated otherwise. As base relations are modified by update operations, the derived relations may also have to be changed. A derived relation can always be brought up-to-date by re-evaluating the relational expression defining it, provided that the necessary base relations are available. However, doing so after every update operation appears extremely wasteful and would probably be unacceptable, both from a performance and a cost point of view.

---

\* This research was supported by Cognos, Inc., Ottawa under contract WRI 502-12, by NSERC under grant No. A-2460, and by CONACYT scholarship No. 35957.

† Electronic mail: {jablakeley,ncoburn,palarson}@waterloo.csnet.

Consider a database  $\mathbf{D} = \{D, S\}$  consisting of a set of base relations  $D = \{R_1, R_2, \dots, R_m\}$  and a set of derived relations  $S = \{E_1, E_2, \dots, E_n\}$ , where each  $E_i \in S$  is a relational algebra expression over some subset of  $D$ . Suppose that an update operation  $U$  is posed against the database  $D$  specifying an update of base relation  $R_u \in D$ . To keep the derived relations consistent with the base relations, those derived relations whose definition involves  $R_u$  may have to be updated as well. The general *update problem for derived relations* consists of: (1) determining which derived relations may be affected by the update  $U$ , and (2) performing the necessary updates to the affected derived relations efficiently.

As a first step towards the solution of this problem, we consider the following two important subproblems. Given an update operation  $U$  and a potentially affected derived relation  $E_i$ ,

- determine the conditions under which the update  $U$  has no effect on the derived relation  $E_i$ , regardless of the database instance. In this case, the update  $U$  is said to be *irrelevant* to  $E_i$
- if the update  $U$  is not irrelevant to  $E_i$ , then determine the conditions under which  $E_i$  can be correctly updated using only  $U$  and the current instance of  $E_i$ , for every instance of the database. That is, no additional data from the base relations  $D$  is required. In this case,  $U$  is said to be *autonomously computable* over  $E_i$ .

The update problem for derived relations is part of an ongoing project at the University of Waterloo on the use of derived relations. The project is investigating a new approach to structuring the database in a relational system at the internal level [A 75]. In current systems there is a one-to-one correspondence between conceptual relations and stored relations, that is, each conceptual relation exists as a separate stored relation (file). This is a simple and straightforward solution, but its drawback is that the processing of a query often requires data to be collected from several stored relations. Instead of directly storing each conceptual relation, we propose structuring the stored database as a set of derived relations. The choice of relations should be guided by the actual or anticipated query load so that frequently occurring queries can be processed rapidly. To speed up query processing, some data may be redundantly stored in several derived relations.

The structure of the stored database should be completely transparent at the user level. This requires a system capable of automatically transforming any user update against a conceptual relation, into equivalent updates against all stored relations affected. The same type of transformation is necessary to process user queries. That is, any query posed against the conceptual relations must be transformed into an equivalent query against the stored relations. The query transformation problem has been addressed in a paper by Larson and Yang [LY 85].

Although our main motivation for studying the problem stems from the above project, its solution also has applications in other areas of relational databases. Buneman and Clemons [BC 79] proposed using views (that is, virtual derived relations) for the support of alerters. An alerter monitors the database and reports when a certain state (defined by the view associated with the alerter)

has been reached. Hammer and Sarin [HS 78] proposed a method for detecting violations of integrity constraints. Certain types of integrity constraints can be seen as defining a view. If we can show that an update operation has no effect on the view associated with an alterer or integrity constraint, then the update cannot possibly trigger the alterer or result in a database instance violating the integrity constraint. The use of derived relations (called concrete views) for the support of real-time queries was suggested by Gardarin et. al. [GSV 84], but it was rejected because of the lack of an efficient update mechanism. Our results have direct application in this area.

The detection of irrelevant or autonomously computable updates also has applications in distributed databases. Suppose that a derived relation is stored at some site and that an update request, possibly affecting the derived relation, is submitted at the same site. If the update is autonomously computable, then the derived relation can be correctly updated locally without requiring data from remote sites. On the other hand, if the request is submitted at a remote site, then we need to send only the update request itself to the site of the derived relation. As well, the results presented here provide a starting point for devising a general mechanism for database snapshot refresh [AL 86, BLT 86, L 86].

## 2. Notation and Basic Assumptions

We assume that the reader is familiar with the basic ideas of relational databases as in Maier [M 83]. A *derived relation* is a relation instance resulting from the evaluation of a relational algebra expression over a database instance. We consider a restricted but important class of derived relations, namely those defined by a relational algebra expression constructed from any combination of project, select and join operations, called a *PSJ-expression*. We often identify a derived relation with its defining expression even though, strictly speaking, the derived relation is the result of evaluating that expression.

We state the following without proof: every valid *PSJ-expression* can be transformed into an equivalent expression in a standard form consisting of a Cartesian product, followed by a selection, followed by a projection. It is easy to see this by considering the query tree corresponding to a *PSJ-expression*. The standard form is obtained by first pushing all projections to the root of the tree and thereafter all selection and join conditions. From this it follows that any *PSJ-expression* can be written in the form  $E = \pi_{\mathbf{A}} \sigma_C (r_{i_1} \times r_{i_2} \times \cdots \times r_{i_k})$ , where  $R_{i_1}, R_{i_2}, \dots, R_{i_k}$  are relation schemes,  $C$  is a selection condition, and  $\mathbf{A} = \{A_1, A_2, \dots, A_l\}$  are the attributes of the projection. We can therefore represent any *PSJ-expression* by a triple  $E = (\mathbf{A}, \mathbf{R}, C)$ , where  $\mathbf{A} = \{A_1, A_2, \dots, A_l\}$  is called the *attribute set*,  $\mathbf{R} = \{R_{i_1}, R_{i_2}, \dots, R_{i_k}\}$  is the *relation set* or *base*, and  $C$  is a *selection condition* composed from the conditions of all the select and join operations of the relational algebra expression defining  $E$ . The attributes in  $\mathbf{A}$  will often be referred to as the *visible* attributes of the derived relation. For simplicity, we assume that each relation of  $\mathbf{R}$  occurs only once in the relational algebra form of the *PSJ-expression*, that is, we do not allow self-joins. We also use the notation:

$\alpha(C)$	The set of all attributes appearing in condition $C$
$\alpha(R)$	The set of all attributes of relation $R$
$V(E, d)$	The relation resulting from evaluating the relational expression $E$ over the instance $d$ of $D$

The update operations considered are insertions, deletions, and modifications. Each update operation affects only one (conceptual) relation. The following notation will be used for update operations:

INSERT ( $R_u, T$ )	Insert into relation $R_u$ the set of tuples $T$
DELETE ( $R_u, \mathbf{R}_D, C_D$ )	Delete from relation $R_u$ all tuples satisfying condition $C_D$ , where $C_D$ is a selection condition over the relations $\mathbf{R}_D, \mathbf{R}_D \subseteq D$
MODIFY ( $R_u, \mathbf{R}_M, C_M, \mathbf{F}_M$ )	Modify all tuples in $R_u$ that satisfy the condition $C_M$ , where $C_M$ is a selection condition over the relations $\mathbf{R}_M, \mathbf{R}_M \subseteq D$ . $\mathbf{F}_M$ is a set of expressions, each expression specifying how an attribute of $R_u$ is to be modified

A DELETE or MODIFY operation must specify the set of tuples from  $R_u$  to be updated. Selecting the set of tuples to be deleted from or modified in  $R_u$  can be seen as a query to the database. In the same way as derived relations, these “selection queries” are restricted to those defined by *PSJ*-expressions. For the operation DELETE ( $R_u, \mathbf{R}_D, C_D$ ), the set of tuples to be deleted from  $R_u$  is selected by the *PSJ*-expression  $E_D = (\alpha(R_u), \mathbf{R}_D, C_D)$ . Similarly, for the operation MODIFY ( $R_u, \mathbf{R}_M, C_M, \mathbf{F}_M$ ), the set of tuples to be modified in  $R_u$  is selected by the *PSJ*-expression  $E_M = (\alpha(R_u), \mathbf{R}_M, C_M)$ .

$\mathbf{F}_M$  is assumed to contain an update expression for each attribute in  $R_u$ . We restrict the update expressions in  $\mathbf{F}_M$  to unconditional functions that can be computed “tuple-wise”. Unconditional means that the expression does not include any further conditions (all conditions are in  $C_M$ ). Tuple-wise means that, for any tuple in  $R_u$  selected for modification, the value of the expression can be computed from the values of the attributes of that tuple alone. The type of expressions we have in mind are simple, for example,  $H := H + 5, I := 5$ . Further details are given in section 4.3. We make the assumption that all the attributes involved in the update expressions are from relation  $R_u$ . That is, both the attributes modified and the attributes from which the new values are computed, are from relation  $R_u$ . If the attributes from which the new values are computed, are from a relation  $R_v, R_v \neq R_u$ , then it is unclear which tuple in  $R_v$  should be used to compute the new values.

All attribute names in the base relations are taken to be unique. We also assume that all attributes have discrete and finite domains. Any such domain can be mapped onto an interval of integers, and therefore we will in the sequel treat all attributes as being defined over some interval of integers. For Boolean expressions, the logical connectives will be denoted by “ $\vee$ ” for OR, juxtaposition or “ $\wedge$ ” for AND, “ $\neg$ ” for NOT, “ $\Rightarrow$ ” for implication, and “ $\Leftrightarrow$ ” for equivalence. To indicate that all variables of a condition  $C$ , are universally quantified, we write  $\forall C$ ; similarly for existential quantification. If we need to explicitly identify which variables are quantified, we write  $\forall_X(C)$  where  $X$  is a set of

variables.

An *evaluation* of a condition is obtained by replacing all the variable names (attribute names) by values from the appropriate domains. The result is either *true* or *false*. A *partial evaluation* (or *substitution*) of a condition is obtained by replacing some of its variables by values from the appropriate domains. Let  $C$  be a condition and  $t$  a tuple over some set of attributes. The partial evaluation of  $C$  with respect to  $t$  is denoted by  $C[t]$ . The result is a new condition with fewer variables.

### 3. Basic Concepts

Detecting whether an update operation is irrelevant or autonomously computable involves testing whether or not certain Boolean expressions are valid, or equivalently, whether or not certain Boolean expressions are unsatisfiable.

**Definition:** Let  $C(x_1, x_2, \dots, x_n)$  be a Boolean expression over variables  $x_1, x_2, \dots, x_n$ .  $C$  is *valid* if  $\forall x_1, x_2, \dots, x_n C(x_1, x_2, \dots, x_n)$  is *true*, and  $C$  is *unsatisfiable* if  $\nexists x_1, x_2, \dots, x_n C(x_1, x_2, \dots, x_n)$  is *true*, where each variable  $x_i$  ranges over its associated domain.  $\square$

A Boolean expression is valid if it always evaluates to *true*, unsatisfiable if it never evaluates to *true*, and satisfiable if it evaluates to *true* for some values of its variables. Proving the validity of a Boolean expression is equivalent to disproving the satisfiability of its complement. Proving the satisfiability of Boolean expressions is, in general, *NP*-complete. However, for a restricted class of Boolean expressions, polynomial algorithms exist. Rosenkrantz and Hunt [RH 80] developed such an algorithm for conjunctive Boolean expressions. Each expression  $B$  must be of the form:  $B = B_1 \wedge B_2 \wedge \dots \wedge B_m$  where each  $B_i$  is an atomic condition. An atomic condition must be of the form  $x \text{ op } y + c$  or  $x \text{ op } c$ , where  $\text{op} \in \{=, <, \leq, >, \geq\}$ ,  $x$  and  $y$  are variables, and  $c$  is a (positive or negative) constant. Each variable is assumed to range over the integers. The algorithm runs in  $O(n^3)$  time where  $n$  is the number of distinct variables in  $B$ .

In this paper, we are interested in the case when each variable ranges over a finite *interval* of integers. For this case, Larson and Yang [LY 85] developed an algorithm whose running time is  $O(n^2)$ . However, it does not handle expressions of the form  $x \text{ op } y + c$  where  $c \neq 0$ . We have developed a modified version of the algorithm by Rosenkrantz and Hunt for the case when each variable ranges over a finite interval of integers. The full details of the modified algorithm are given in the Appendix.

An expression not in conjunctive form can be handled by first converting it into disjunctive normal form and then testing each conjunct separately. Several of the theorems in sections 4 and 5 will require testing the validity of expressions of the form  $C_1 \Rightarrow C_2$ . The implication can be eliminated by converting to the form  $(\neg C_1) \vee C_2$ . Similarly, expressions of the form  $C_1 \Leftrightarrow C_2$  can be converted to  $C_1 C_2 \vee (\neg C_1)(\neg C_2)$ .

The concepts covered by the three definitions below were introduced in Larson and Yang [LY 85]. As they will be needed in sections 4 and 5 of this paper, we include them here for completeness.

**Definition:** Let  $C$  be a Boolean expression over the variables  $x_1, x_2, \dots, x_n$ . The variable  $x_i$  is said to be *nonessential* in  $C$  if  $\forall x_1, \dots, x_i, \dots, x_n, x_i' (C(x_1, \dots, x_i, \dots, x_n) = C(x_1, \dots, x_i', \dots, x_n))$ . Otherwise,  $x_i$  is *essential* in  $C$ .  $\square$

A nonessential variable can be eliminated from the condition simply by replacing it with any value from its domain. This will in no way change the value of the condition. For example, the variable  $H$  is nonessential in the following two conditions: (1)  $(I > 5)(J = I)((H > 5) \vee (H < 10))$ , and (2)  $(I > 5)(H > 5)(H \leq 5)$ .

**Definition:** Let  $C_0$  and  $C_1$  be Boolean expressions over the variables  $x_1, x_2, \dots, x_n$ . The variable  $x_i$  is said to be *computationally nonessential* in  $C_0$  with respect to  $C_1$  if

$$\forall x_1, \dots, x_i, \dots, x_n, x_i' [C_1(x_1, \dots, x_i, \dots, x_n) C_1(x_1, \dots, x_i', \dots, x_n) \Rightarrow (C_0(x_1, \dots, x_i, \dots, x_n) = C_0(x_1, \dots, x_i', \dots, x_n))] .$$

Otherwise,  $x_i$  is *computationally essential* in  $C_0$ .  $\square$

If a variable  $x_i$  (or a subset of the variables  $x_1, x_2, \dots, x_n$ ) is computationally nonessential in  $C_0$  with respect to  $C_1$ , we can correctly evaluate the condition  $C_0$  without knowing the exact value of  $x_i$ . That is, given any tuple  $t = (x_1, x_2, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$  where the full tuple (including  $x_i$ ) is known to satisfy  $C_1$ , we can correctly determine whether or not  $t$  satisfies  $C_0$ . This can be done by determining a surrogate value for  $x_i$  as explained in Larson and Yang [LY 85].

*Example:* Let  $C_1 \equiv (H > 5)$  and  $C_0 \equiv (H > 0)(I = 5)(J > 10)$ . It is easy to see that if we are given a tuple  $(i, j)$  for which it is known that the full tuple  $(h, i, j)$  satisfies  $C_1$ , then we can correctly evaluate  $C_0$ . If  $(h, i, j)$  satisfies  $C_1$  then the value of  $h$  must be greater than 5, and consequently it also satisfies  $(H > 0)$ . Hence, we can correctly evaluate  $C_0$  for the tuple  $(i, j)$  by assigning to  $H$  any surrogate value greater than 5.  $\square$

**Definition:** Let  $C$  be a Boolean expression over the variables  $x_1, x_2, \dots, x_n, y_1, y_2, \dots, y_m$ . The variable  $y_i$  is said to be *uniquely determined* by  $C$  and  $x_1, \dots, x_n$  if

$$\forall x_1, \dots, x_n, y_1, \dots, y_m, y_1', \dots, y_m' [C(x_1, \dots, x_n, y_1, \dots, y_m) C(x_1, \dots, x_n, y_1', \dots, y_m') \Rightarrow y_i = y_i'] . \square$$

If a variable  $y_i$  (or a subset of the variables  $y_1, y_2, \dots, y_m$ ) is uniquely determined by a condition  $C$  and the variables  $x_1, \dots, x_n$ , then given any tuple  $t = (x_1, \dots, x_n)$ , such that the full tuple  $(x_1, \dots, x_n, y_1, \dots, y_m)$  is known to satisfy  $C$ , the missing value of the variable  $y_i$  can be correctly reconstructed. How to reconstruct the values of uniquely determined variables was also shown in Larson and Yang [LY 85]. If the variable  $y_i$  is not uniquely determined, then we cannot guarantee that its value is reconstructible for *every* tuple. However, it may still be reconstructible for *some* tuples.

*Example:* Let  $C \equiv (I = H)(H > 7)(K = 5)$ . It is easy to prove that  $I$  and  $K$  are uniquely determined by  $H$  and the condition  $C$ . Suppose that we are given a tuple that satisfies  $C$  but only the value of  $H$  is known. Assume that  $H = 10$ . Then we can immediately determine that the values of  $I$  and  $K$  must be 10 and 5, respectively.  $\square$

**Definition:** Let  $E = (A, R, C)$  be a derived relation and let  $B$  be the set of all attributes uniquely determined by the attributes in  $A$  and the condition  $C$ . Then  $A^+ = A \cup B$  is called the *extended attribute set* of  $E$ .

Note that  $A^+$  is the maximal set of attributes for which values can be reconstructed for every tuple of  $E$ .

#### 4. Detecting Irrelevant Updates

This section considers irrelevant updates. We deal with insertions, then deletions, and finally the most difficult case, modifications. First we define what it means for an update to be irrelevant.

**Definition:** Let  $d$  denote an instance of the base relations  $D$  and  $d'$  the resulting instance after applying the update operation  $U$  to  $d$ . Let  $E$  be a derived relation. The update operation  $U$  is *irrelevant* to  $E$  if  $V(E, d') = V(E, d)$  for all instances  $d$  and  $d'$ .  $\square$

If the update operation  $U$  does not modify any of the relations over which  $E$  is defined then, obviously,  $U$  cannot have any effect on the derived relation. In this case  $U$  is said to be *trivially irrelevant* to  $E$ .

##### 4.1. Irrelevant Insertions

An insert operation  $\text{INSERT}(R_u, T)$  is irrelevant to a derived relation if none of the new tuples will be visible in the derived relation. Note that this should hold regardless of the state of the database. The following theorem was proven in [BLT 86] and is included here for completeness only.

**Theorem 1:** The operation  $\text{INSERT}(R_u, T)$  is irrelevant to the derived relation  $E = (A, R, C)$ ,  $R_u \in R$ , if and only if  $C[t]$  is unsatisfiable for every tuple  $t \in T$ .

**Proof:** (Sufficiency) Consider an arbitrary tuple  $t \in T$ . If  $C[t]$  is unsatisfiable, then for every assignment of values to the free variables in  $C[t]$ , it will evaluate to *false*. Therefore, there cannot exist any tuple defined over the Cartesian product of the relations in  $R - \{R_u\}$  that would combine with  $t$  to satisfy  $C$  and hence cause an insertion into  $E$ .

(Necessity) Consider an arbitrary tuple  $t \in T$  and assume that  $C[t]$  is satisfiable but that the  $\text{INSERT}$  operation is irrelevant to  $E$ . Without loss of generality we may assume that  $\alpha(C) - \alpha(C[t])$  contains a single variable. Since  $C[t]$  is satisfiable there exists some value,  $x$ , for the variable in  $\alpha(C) - \alpha(C[t])$  for which  $C[t]$  evaluates to *true*. We can construct a database instance  $d$ , using the value  $x$ , such that the insertion of  $t$  into  $R_u$  will cause a tuple to be inserted into the derived relation  $E$ .



To construct  $d$  we form, for each relation  $R_i$  in  $\mathbf{R}-\{R_u\}$  an instance,  $r_i$ , containing a single tuple. The values for  $r_i$  are determined as follows: if an attribute of  $r_i$  is the attribute in  $\alpha(C)-\alpha(C[t])$  then it is assigned the value  $x$ ; otherwise it is assigned an arbitrary value, say, the smallest in its domain.

The database instance  $d$  consists of the empty relation  $R_u$  and the set of instances of the relations  $\mathbf{R}-\{R_u\}$  each of which contains a single tuple. Hence,  $V(E, d) = \emptyset$ . However, if we obtain  $d'$  from  $d$  by inserting tuple  $t$  into relation  $R_u$ , then  $V(E, d')$  will contain one tuple. Therefore, the INSERT operation is not irrelevant to the derived relation  $E$ .  $\square$

#### 4.2. Irrelevant Deletions

A delete operation is irrelevant to a derived relation if none of the tuples in the derived relation will be deleted. We have the following theorem.

**Theorem 2:** The operation  $\text{DELETE}(R_u, \mathbf{R}_D, C_D)$  is irrelevant to the derived relation  $E = (A, \mathbf{R}, C)$ ,  $R_u \in \mathbf{R}$ , if and only if the condition  $C_D \wedge C$  is unsatisfiable.

**Proof:** Let  $\mathbf{B} = \mathbf{R} \cup \mathbf{R}_D = \{R_{i_1}, R_{i_2}, \dots, R_{i_k}\}$ .  $\mathbf{B}$  is called the *combined base* of the derived relation and the delete operation. We first show that we can extend the base of both  $E$  and the delete query  $E_D = (\alpha(R_u), \mathbf{R}_D, C_D)$  to  $\mathbf{B}$  without affecting their result in any way. Without loss of generality, we can assume that  $\mathbf{R}_D - \mathbf{R} = \{R_{i_1}\}$  so that  $\mathbf{R} = \{R_{i_2}, R_{i_3}, \dots, R_{i_k}\}$ . Let  $t$  be a tuple in the Cartesian product  $r_{i_2} \times r_{i_3} \times \dots \times r_{i_k}$  (the base before adding  $R_{i_1}$ ). If  $t$  satisfies  $C$ , then  $t[A]$  (the projection of  $t$  onto  $A$ ) will be visible in the derived relation, otherwise it will not. Extending the base to  $r_{i_1} \times r_{i_2} \times \dots \times r_{i_k}$  may give rise to a number of "copies" of  $t$  in the extended base. The copies differ only in the attributes of  $R_{i_1}$ . Since

$$\alpha(C) \subseteq \bigcup_{j=2}^k \alpha(R_{i_j})$$

then  $\alpha(R_{i_1}) \cap \alpha(C) = \emptyset$ . Hence, if  $t$  satisfies  $C$ , then all its copies will satisfy  $C$ . Similarly, if  $t$  does not satisfy  $C$ , then none of its copies will satisfy  $C$  either. The projection onto  $A$  will finally reduce all copies of  $t$  to a single tuple, exactly  $t[A]$ . This proves that extending the base of  $E$  does not change the resulting derived relation. In the same way, we can show that extending the base of the DELETE operation has no effect. We now complete the proof of the theorem.

(Sufficiency) Let  $t$  be a tuple over the combined base  $\mathbf{B}$  and assume that  $t$  satisfies  $C$ . Then  $t[A]$  is visible in the derived relation. If  $C_D \wedge C$  is unsatisfiable, then  $t$  cannot at the same time satisfy  $C_D$ . Hence  $t[A]$  will not be deleted from the derived relation.

(Necessity) Assume that  $C_D \wedge C$  is satisfiable. We can then construct an instance of each relation in  $\mathbf{B}$  such that deleting a tuple from  $r_u$ , ( $R_u \in \mathbf{B}$ ), will indeed change the derived relation. Let  $\alpha(C) \cup \alpha(C_D) = \{x_1, x_2, \dots, x_l\}$ . Because  $C_D \wedge C$  is satisfiable, there exists a value combination  $X^0 = \langle x_1^0, x_2^0, \dots, x_l^0 \rangle$  such that  $C[X^0]C_D[X^0]$  is true. We now construct one

tuple  $t_{i_j}$  for each relation  $R_{i_j} \in B$ . The attribute values of  $t_{i_j}$  are assigned as follows: if the attribute occurs in  $\alpha(C) \cup \alpha(C_D)$ , assign it the corresponding value from  $X^0$ , otherwise assign it an arbitrary value in its domain, the minimum value, for example. We now have a database instance where each relation, as well as the Cartesian product  $r_{i_1} \times r_{i_2} \times \dots \times r_{i_k}$ , contains one tuple. The tuple in the Cartesian product obviously satisfies  $C$  and hence the derived relation also contains one tuple. It also satisfies  $C_D$  and hence the relation  $r_u$ , will be empty after the deletion operation has been performed. Therefore, evaluating  $E$  over the new instance of the database will result in the empty set. This proves that the stated condition is necessary.  $\square$

*Example:* Consider two relations  $R_1(H, I, J), R_2(K, L)$  and the following derived relation and delete operation:

$E = (\{H, L\}, \{R_1, R_2\}, (J = K)(K > 10)(I = 5))$ , and  
DELETE  $(R_1, \{R_1\}, (J < 5)(I < 10))$ .

To show that the deletion is irrelevant to the derived relation we must prove that the following condition holds:

$\forall I, J, K \neg [(J = K)(K > 10)(I = 5)(J < 5)(I < 10)]$ . This is equivalent to proving that  $\exists I, J, K [(J = K)(K > 10)(I = 5)(J < 5)(I < 10)]$ , which can be simplified to  $\exists I, J, K [(J = K)(I = 5)(K > 10)(K < 5)]$ . The condition  $(K > 10)(K < 5)$  can never be satisfied and therefore the delete operation is irrelevant to the derived relation.  $\square$

**Corollary 2.1:** If  $C$  and  $C_D$  are independently satisfiable and  $\alpha(C_D) \cap \alpha(C) = \emptyset$ , then  $C_D \wedge C$  is satisfiable and, therefore, the DELETE operation is not irrelevant to the derived relation  $E$ .  $\square$

### 4.3. Irrelevant Modifications

Modifications are somewhat more complicated than insertions or deletions. Consider a tuple that is to be modified. It will not affect the derived relation if one of the following conditions applies:

- it does not qualify for the derived relation, neither before nor after the modification
- it does qualify for the derived relation both before and after the modification and, furthermore, all the attributes visible in the derived relation remain unchanged

Some additional notation is needed at this point. Consider a modify operation MODIFY  $(R_u, \mathbf{R}_M, C_M, \mathbf{F}_M)$  and a derived relation  $E = (A, \mathbf{R}, C)$ . Let  $\alpha(R_u) = \{B_1, B_2, \dots, B_l\}$ . For simplicity we will associate an update expression with every attribute in  $R_u$ , that is,  $\mathbf{F}_M = \{f_{B_1}, f_{B_2}, \dots, f_{B_l}\}$  where each update expression is of the form  $f_{B_i} \equiv (B_i := \langle \text{arithmetic expression} \rangle)$ . If an attribute  $B_i$  is not to be modified, we associate with it a *trivial update expression* of the form  $f_{B_i} \equiv (B_i := B_i)$ . If the attribute is assigned a fixed value  $c$ , then the corresponding update expression is  $f_{B_i} \equiv (B_i := c)$ . The notation  $\rho(f_{B_i})$  will be used to denote the right hand side of the update expression  $f_{B_i}$ , that is, the expression after the assignment sign. The notation  $\alpha(\rho(f_{B_i}))$  denotes the vari-

ables mentioned in  $\rho(f_{B_i})$ . For example, if  $f_{B_i} \equiv (B_i := B_j + c)$  then  $\rho(f_{B_i}) = B_j + c$  and  $\alpha(\rho(f_{B_i})) = \{B_j\}$ .

By substituting every occurrence of an attribute  $B_i$  in  $C$  by  $\rho(f_{B_i})$  a new condition is obtained. We will use the notation  $C(\mathbf{F}_M)$  to denote the condition obtained by performing this substitution for every variable  $B_i \in \alpha(R_u) \cap \alpha(C)$ .

A modification may result in a value outside the domain of the modified attribute. We make the assumption that such an update will not be performed, that is, the entire tuple will remain unchanged. Each attribute  $B_i$  of  $R_u$  must satisfy a condition of the form  $(B_i \leq U_{B_i})(B_i \geq L_{B_i})$  where  $L_{B_i}$  and  $U_{B_i}$  are the lower and upper bound, respectively, of its domain. Hence, the updated value of  $B_i$  must satisfy the condition  $(\rho(f_{B_i}) \leq U_{B_i})(\rho(f_{B_i}) \geq L_{B_i})$  and this must hold for every  $B_i \in \alpha(R_u)$ . The conjunction of all these conditions will be denoted by  $C_B(\mathbf{F}_M)$ , that is,

$$C_B(\mathbf{F}_M) \equiv \bigwedge_{B_i \in \alpha(R_u)} (\rho(f_{B_i}) \leq U_{B_i})(\rho(f_{B_i}) \geq L_{B_i})$$

**Theorem 3:** The modify operation  $\text{MODIFY}(R_u, \mathbf{R}_M, C_M, \mathbf{F}_M)$  is irrelevant to the derived relation  $E = (\mathbf{A}, \mathbf{R}, C)$ ,  $R_u \in \mathbf{R}$ , if and only if

$$\begin{aligned} & \forall [(C_M \wedge C_B(\mathbf{F}_M)) \\ & \Rightarrow ((\neg C) \wedge (\neg C(\mathbf{F}_M))) \vee (C \wedge C(\mathbf{F}_M) \bigwedge_{B_i \in I} (B_i = \rho(f_{B_i}))) ] \end{aligned}$$

where  $I = \mathbf{A} \cap \alpha(R_u)$ .

**Proof:** (Sufficiency) In the same way as in the proof of Theorem 2, let  $\mathbf{B} = \mathbf{R} \cup \mathbf{R}_M$  be the combined base. Consider a tuple  $t$  from the combined base  $\mathbf{B}$  such that  $t$  satisfies  $C_M$  and the corresponding modified tuple, denoted by  $t'$ , satisfies  $C_B(\mathbf{F}_M)$ . Because the above condition holds for every tuple, it must also hold for  $t$ . Hence, either the first or the second conjunct of the consequent must evaluate to *true*. They cannot both be *true* simultaneously.

If the first conjunct is *true*, both  $C[t]$  and  $C[t']$  must be *false*. This means that neither the original tuple  $t$ , nor the modified tuple  $t'$ , will contribute to the derived relation. Hence changing  $t$  to  $t'$  will not affect the derived relation.

If the second conjunct is *true*, both  $C[t]$  and  $C[t']$  must be *true*. In other words, the tuple  $t$  contributed to the derived relation and after being modified to  $t'$ , it still remains in the derived relation. The last part of the conjunct must also be satisfied, which ensures that all attributes of  $R_u$  visible in the derived relation have the same values in  $t$  and  $t'$ . Hence the derived relation will not be affected.

(Necessity) Assume that the above condition does not hold. That means that there exists at least one assignment of values to the attributes, that is, a tuple  $t$  such that the antecedent is *true* but the consequent is *false*. Denote the corresponding modified tuple by  $t'$ . There are three cases to consider.

Case 1:  $C[t] = \text{true}$  and  $C[t'] = \text{false}$ . In the same way as in the proof of Theorem 2, we can from  $t$  construct an instance  $d$  where each relation in  $\mathbf{B}$  contains a single tuple and where the derived relation

contains exactly one tuple as well. For this instance, the modification operation will produce a new instance  $d'$  where the only change is to the (single) tuple of relation  $R_u$ . The Cartesian product of the relations in  $\mathbf{B}$  then contains exactly one tuple, which agrees with  $t'$  on the corresponding attributes. Hence, the derived relation  $V(E, d')$  will be empty since  $C[t'] = false$ . This proves that the update operation is not irrelevant to the derived relation.

- Case 2:  $C[t] = false$  and  $C[t'] = true$ . Can be proven in the same way as Case 1, with the difference that the derived relation is initially empty and the modification results in a tuple being inserted into the derived relation.
- Case 3:  $C[t] = true$ ,  $C[t'] = true$  but  $\bigwedge_{B_i \in I} (B_i = \rho(f_{B_i}))$  is *false*, that is,  $t[B_i] \neq t'[B_i]$  for some  $B_i \in \mathbf{A} \cap \alpha(R_u)$ . In the same way as above, we can construct an instance where each relation in  $\mathbf{B}$  contains only a single tuple, and where the derived relation also contains a single tuple, both before and after the modification. However, in this case the value of attribute  $B_i$  will change as a result of performing the MODIFY operation. Since  $B_i \in \mathbf{A}$ , this change will be visible in the derived relation. This proves that the update is not irrelevant to the derived relation.  $\square$

The following example illustrates the theorem.

*Example:* Suppose the database consists of the two relations  $R_1(H, I)$  and  $R_2(J, K)$  where  $H, I, J$  and  $K$  each have the domain  $[0, 30]$ . Let the derived relation and modify operation be defined as:

$$E = (\{I, J\}, \{R_1, R_2\}, (H > 10)(I = K))$$

$$\text{MODIFY } (R_1, \{R_1\}, (H > 20), \{(H := H + 5), (I := I)\}) .$$

Thus the condition given in Theorem 3 becomes

$$\begin{aligned} \forall H, I, K [(H > 20)(H + 5 \geq 0)(H + 5 \leq 30) \\ \Rightarrow (\neg((H > 10)(I = K))) \wedge (\neg((H + 5 > 10)(I = K))) \\ \vee (H > 10)(I = K)(H + 5 > 10)(I = K)(I = I)] \end{aligned}$$

which can be simplified to

$$\begin{aligned} \forall H, I, K [(H > 20)(H \leq 25) \\ \Rightarrow (\neg((H > 10)(I = K))) \wedge (\neg((H > 5)(I = K))) \\ \vee (H > 10)(I = K)] . \end{aligned}$$

By inspection we see that if  $I = K$ , then the second term of the consequent will be satisfied whenever the antecedent is satisfied. If  $I \neq K$ , the first term of the consequent is always satisfied. Hence, the implication is valid and we conclude that the update is irrelevant to the derived relation.  $\square$

**Corollary 3.1:** If  $B_i = \rho(f_{B_i})$  for all  $B_i \in I$  and for all  $B_i \in (\alpha(R_u) \cap \alpha(C))$  then the MODIFY operation is irrelevant to the derived relation  $E$ .

**Proof:** Since  $B_i = \rho(f_{B_i})$  for all  $B_i \in I$  then  $\bigwedge_{B_i \in I} (B_i = \rho(f_{B_i}))$  is *true*. Also, since  $B_i = \rho(f_{B_i})$  for all  $B_i \in (\alpha(R_u) \cap \alpha(C))$  then  $C(\mathbf{F}_M) = C$ . Therefore, the condition in the theorem becomes

$$\forall [C_M C_B(\mathbf{F}_M) \Rightarrow ((\neg C) \vee C)]$$

which is equivalent to

$$\forall [C_M C_B(\mathbf{F}_M) \Rightarrow (\text{true})]$$

Hence, we conclude that the implication is valid and that the MODIFY operation is irrelevant to the derived relation.  $\square$

### 5. Autonomously Computable Updates

If an update operation is not irrelevant to a derived relation, then some data from the database is needed to correctly update the derived relation. The simplest case is when all the data needed is contained in the derived relation itself. In other words, the new state of the derived relation can be computed solely from the current state of the derived relation and the information contained in the update expression.

**Definition.** Consider a derived relation  $E$  and an update operation  $U$ , both defined over base relations  $D$ . Let  $d$  denote an instance of  $D$  before applying  $U$  and  $d'$  the corresponding instance after applying  $U$ . The effect of the operation  $U$  on  $E$  is said to be *autonomously computable* if there exists a function  $F_U$  such that

$$V(E, d') = F_U(V(E, d))$$

for all database instances  $d$  and  $d'$ . Apart from the information in  $U$  itself, the only other data required by  $F_U$  must be contained in the current instance of  $E$ .  $\square$

#### 5.1. Insertions

Consider an operation INSERT  $(R_u, T)$  where  $T$  is a set of tuples to be inserted into  $R_u$ . Let the derived relation be  $E = (A, R, C)$ ,  $R_u \in R$ . The effect of the INSERT operation on the derived relation is autonomously computable if

A. given a tuple  $t \in T$  we can correctly decide whether  $t$  will always (regardless of the database instance) satisfy the selection condition  $C$  and hence should be inserted into the derived relation

and,

B. the values for all attributes visible in the derived relation can be obtained from  $t$ .

Note that if  $t$  could cause the insertion of more than one tuple into the derived relation, then the update is not autonomously computable. Suppose that  $t$  generates two different tuples to be inserted:  $t_1$  and  $t_2$ . Then  $t_1$  and  $t_2$  must differ in at least one attribute visible in the derived relation; otherwise only one

tuple would be inserted. Suppose that they differ on  $A_i \in A$ .  $A_i$  cannot be an attribute of  $R_u$  because the exact value of every attribute in  $R_u$  is given by  $t$ . Hence, the values of  $A_i$  in  $t_1$  and  $t_2$  would have to be obtained from tuples elsewhere in the database.

**Theorem 4A:** Let  $E = (A, R, C)$  be a derived relation and  $t$  a tuple to be inserted into relation  $R_u$ , where  $R_u \in R$ . Whether or not  $t$  will create an insertion into the derived relation is guaranteed to be autonomously computable if and only if one of the following holds:

I.  $R = \{R_u\}$

or

II.  $R \neq \{R_u\}$  and all the variables of  $C[t]$  are nonessential and the current instance of  $E$  is non-empty.

**Proof:** (Sufficiency)

Case I: Since  $R = \{R_u\}$  then  $\alpha(C) \subseteq \alpha(R_u)$ . Hence,  $C[t]$  can be completely evaluated, i.e. will yield *true* or *false*.

Case II: The fact that all variables in  $\alpha(C[t])$  are nonessential guarantees that  $C[t]$  will evaluate to the same value regardless of the values assigned to those variables. Since the current instance of  $E$  is non-empty, the Cartesian product of all relations in  $R - \{R_u\}$  will contain at least one tuple. Combining  $t$  with a tuple from this Cartesian product gives a tuple with fixed values for all variables in  $\alpha(C)$  and the condition can be evaluated. Whatever the values of the attributes in  $\alpha(C[t])$  are, the condition will always evaluate to the same truth value. Hence, whatever the current instance of the database the decision will always be the same.

(Necessity) Assume that whether or not  $t$  will create on insertion into the derived relation is autonomously computable but that neither of the two cases holds. Since the second case contains three conditions, three possibilities arise:

- $(R \neq \{R_u\})$  and  $(R = \{R_u\})$ . This is obviously a contradiction.
- $(R \neq \{R_u\})$  and there exists some variable,  $x \in \alpha(C[t])$ , which is essential in  $C[t]$ . Without loss of generality we can assume that  $x$  is the only variable in  $\alpha(C[t])$ . This means that there exists two different values  $x'$  and  $x''$  such that  $C[t, x']$  is *true* and  $C[t, x'']$  is *false*. In the same way as in the proof of Theorem 2, we can construct two different instances  $d'$  and  $d''$  of  $D$ . Instance  $d'$  is constructed from  $x'$  and instance  $d''$  from  $x''$ , such that, except for the given values of  $x$ , all the corresponding attribute values agree. In both instances relation  $R_u$  is empty and every other relation in  $D$  consists of a single tuple. Hence,

$$V(E, d') = V(E, d'')$$

Now insert tuple  $t$  into relation  $R_u$ . Since  $C[t, x']$  is *true*,  $V(E, d')$  must have a new tuple inserted, whereas  $V(E, d'')$  will not, as  $C[t, x'']$  is *false*. Consequently, whether or not insertion of  $t$  will affect the derived relation depends on the existence of tuples not seen in the derived relation.

- ( $\mathbf{R} \neq \{R_u\}$ ) and the current instance of  $E$  is empty. There are two situations which would cause  $E$  to be empty; either no tuple in the Cartesian product of the base relations satisfies  $C$  or one of the base relations is empty. If  $R_v \in \mathbf{R}$ ,  $R_v \neq R_u$ , is empty then even if  $C[t]$  is true,  $t$  will not cause an insertion into  $E$ . Consequently, whether or not the insertion of  $t$  will affect the derived relation depends on the existence of tuples in the other relations in the base of  $E$ . That is, on the existence of tuples not seen in the derived relation.  $\square$

**Corollary 4A.1:** If  $\alpha(C) \subseteq \alpha(R_u)$  and the current instance of  $E$  is non-empty then whether or not  $t$  causes an insertion into the derived relation is autonomously computable.  $\square$

**Theorem 4B:** Assume that a tuple in  $T$  has been shown to cause the insertion of a new tuple into the derived relation. The values of all visible attributes in the new tuple are guaranteed to be autonomously computable if and only if  $A \subseteq \alpha(R_u)$ .

**Proof:** (Sufficiency) Obvious.

(Necessity) Without loss of generality we can assume that  $A - \alpha(R_u)$  contains only one attribute  $x \in \alpha(R_i)$ ,  $R_i \neq R_u$ . Assume that  $t \in T$  causes the insertion of a new tuple. To insert the new tuple into the derived relation we must determine the value of  $x$ . Even if the value of  $x$  is uniquely determined by the attribute values of  $t$ , this is not sufficient. The value of  $x$  must correspond to the  $x$  value in some tuple in  $R_i$ , and the existence of such a tuple cannot be guaranteed without checking the current instance of the relation  $R_i$ .  $\square$

## 5.2. Deletions

To handle deletions autonomously, we must be able to determine, for every tuple in the derived relation, whether or not it satisfies the delete condition. This is covered by the following theorem.

**Theorem 5:** The effect on the derived relation  $E = (A, \mathbf{R}, C)$  of the operation DELETE  $(R_u, \mathbf{R}_D, C_D)$ ,  $R_u \in \mathbf{R}$ , is guaranteed to be autonomously computable if and only if every attribute in  $\alpha(C_D) - A$  is computationally nonessential in  $C_D$  with respect to  $C$ .

**Proof:** (Sufficiency) If the variables in  $\alpha(C_D) - A$  are all computationally nonessential, we can correctly evaluate the condition by assigning surrogate values.

(Necessity) Without loss of generality we can assume that  $\alpha(C_D) - A$  consists of a single attribute  $x$ . Assume that  $x$  is computationally essential in  $C_D$  with respect to  $C$ . We can then construct two tuples  $t_1$  and  $t_2$  over the attributes in  $A \cup \alpha(C) \cup \alpha(C_D)$  such that they both satisfy  $C$ ,  $t_1$  satisfies  $C_D$  but  $t_2$  does not, and  $t_1$  and  $t_2$  agree on all attributes except attribute  $x$ . Each of  $t_1$  and  $t_2$  can now be extended into an instance of  $D$ . Both instances will give the same instance of the derived relation, consisting of a single tuple  $t_1[A]$  (or  $t_2[A]$ ). In one instance, the tuple should be deleted from the derived relation, in the other one it should not. The decision depends on the value of attribute  $x$  which is not visible in the derived relation. Hence the decision cannot be made without additional data.  $\square$

*Example:* Consider two relations  $R_1(H, I), R_2(J, K)$ . Let the derived relation and delete operation be defined as:

$$E = (\{J, K\}, \{R_1, R_2\}, (I = J)(H < 20))$$

$$\text{DELETE } (R_1, \{R_1\}, (I = 20)(H < 30))$$

The attributes in  $\alpha(C_D) - A = \{H, I\} - \{J, K\} = \{H, I\}$  must be computationally nonessential in  $C_D$  with respect to  $C$  in order for the deletion to be autonomously computable. That is, the following condition must hold:

$$\forall H, I, H', I', J, K [(I = J)(H < 20)(I' = J)(H' < 20)$$

$$\Rightarrow ((I = 20)(H < 30)) = ((I' = 20)(H' < 30))].$$

The conditions  $(H < 30)$  and  $(H' < 30)$  will both be *true* whenever  $(H < 20)$  and  $(H' < 20)$  are *true*. For any choice of values that make the antecedent *true*, we must have  $J = I = I'$ . Because  $I = I'$ , the conditions  $I = 20$  and  $I' = 20$  are either both *true* or both *false*, and hence the consequent will always be satisfied. Therefore, the variables  $H$  and  $I$  are computationally nonessential in  $C_D$  with respect to  $C$ . This guarantees that for any tuple in the derived relation we can always correctly evaluate the delete condition by assigning surrogate values to the variables  $H$  and  $I$ .

To further clarify the concept of computationally nonessential, consider the following instance of the derived relation  $E$ .

$$E: \begin{array}{cc} J & K \\ \hline 10 & 15 \\ 20 & 25 \end{array}$$

We now have to determine on a tuple by tuple basis which tuples in the derived relation should be deleted. Consider tuple  $t_1 = (10, 15)$  and the condition  $C \equiv (I = J)(H < 20)$ . We substitute for the variables  $J$  and  $K$  in  $C$  the values 10 and 15, respectively, to obtain  $C[t_1] \equiv (I = 10)(H < 20)$ . Any values for  $H, I$  that make  $C[t_1] = \text{true}$ , are valid surrogate values, say  $I = 10, H = 19$ . We can then evaluate  $C_D$  using these surrogate values, and find that  $(10 = 20)(19 < 30) = \text{false}$ . Therefore, tuple  $t_1 = (10, 15)$  should not be deleted from  $E$ . Similarly, for  $t_2 = (20, 25)$  we obtain  $C[t_2] \equiv (I = 20)(H < 20)$ . Surrogate values for  $H$  and  $I$  that make  $C[t_2] = \text{true}$  are  $I = 20, H = 19$ . We then evaluate  $C_D$  using these surrogate values and find that  $(20 = 20)(19 < 30) = \text{true}$ . Therefore, tuple  $t_2 = (20, 25)$  should be deleted from  $E$ .  $\square$

**Corollary 5.1:** If  $\alpha(C_D) \subseteq A^+$  then the DELETE operation is autonomously computable.  $\square$

### 5.3. Modifications

Deciding whether modifications can be performed autonomously is much more complicated than for either insertions or deletions. In general, a modify operation may generate insertions into, deletions from, and modifications of existing tuples of the derived relation as a result of updating a base relation. We summarize the conditions imposed by these possibilities in the following four steps and give one theorem for each step.



- A. Prove that every tuple selected for modification which does not satisfy  $C$  before modification, will not satisfy  $C$  after modification. This means that no new tuples will be inserted into the derived relation.
- B. Prove that we can autonomously compute which tuples in the derived relation should be modified. Call this the modify set.
- C. Prove that we can autonomously compute which tuples in the modify set will not satisfy  $C$  after modification and hence can be deleted from the derived relation.
- D. Prove that, for every tuple in the modify set which will not be deleted, we can autonomously compute the new values for all attributes in  $A$ .

**Theorem 6A:** The operation  $\text{MODIFY}(R_u, \mathbf{R}_M, C_M, \mathbf{F}_M)$  is guaranteed not to create any new tuples which need to be inserted into the derived relation  $E = (A, \mathbf{R}, C)$ ,  $R_u \in \mathbf{R}$ , if and only if

$$\forall [(\neg C)_{C_M} C_B(\mathbf{F}_M) \Rightarrow \neg C(\mathbf{F}_M)]$$

**Proof:** (Sufficiency) Assume that the condition holds. Consider a tuple  $t$  in the Cartesian product of the relations in the combined base  $\mathbf{B} = \mathbf{R} \cup \mathbf{R}_M$ , and assume that  $t$  is selected for modification. Let  $t'$  denote the corresponding tuple after modifications. Assume that  $t$  does not satisfy  $C$  and hence will not have created any tuple in the derived relation. Because the above condition holds for every tuple, it must also hold for  $t$  and hence  $t'$  cannot satisfy  $C$ . Consequently, modifying  $t$  to  $t'$  does not cause any new tuple to appear in the derived relation.

(Necessity) If the condition does not hold we can, in the same way as in the proof of Theorem 2, construct an instance such that each relation contains only one tuple and the derived relation is empty before modification but contains one tuple after modification. This then shows that the condition is necessary.  $\square$

**Corollary 6A.1:** If  $B_i = \rho(f_{B_i})$  for all  $B_i \in (\alpha(R_u) \cap \alpha(C))$  or if  $(\alpha(R_u) \cap \alpha(C)) = \emptyset$  then the  $\text{MODIFY}$  operation will not cause any insertions into the derived relation.

**Proof:** In either case the values of the variables which appear in the selection condition for  $E$ , that is in  $C$ , are unchanged. Hence, no new tuples will need to be inserted into  $E$ .  $\square$

**Corollary 6A.2:** If  $C_M \wedge C(\mathbf{F}_M)$  is unsatisfiable then the  $\text{MODIFY}$  operation will not cause any insertions into the derived relation.

**Proof:** If  $C_M \wedge C(\mathbf{F}_M)$  is unsatisfiable then  $\neg(C_M \wedge C(\mathbf{F}_M))$  must always be true. Hence  $\neg C_M \vee \neg C(\mathbf{F}_M)$  is always true. Therefore,  $C_M \Rightarrow \neg C(\mathbf{F}_M)$  is valid which means the implication given in the theorem must be valid.  $\square$

**Theorem 6B:** The condition  $C_M$  is guaranteed to be autonomously computable if and only if every attribute in  $\alpha(C_M) - A$  is computationally nonessential in  $C_M$  with respect to  $C$ .  $\square$

The following proof is similar to the proof of Theorem 5, it is included here for completeness.

**Proof:** (Sufficiency) If the variables in  $\alpha(C_M) - A$  are all computationally nonessential, we can correctly evaluate the condition by assigning surrogate values.

(Necessity) Without loss of generality we can assume that  $\alpha(C_M) - A$  consists of a single attribute  $x$ . Assume that  $x$  is computationally essential in  $C_M$  with respect to  $C$ . We can then construct two tuples  $t_1$  and  $t_2$  over the attributes in  $A \cup \alpha(C) \cup \alpha(C_M)$  such that they both satisfy  $C$ ,  $t_1$  satisfies  $C_M$  but  $t_2$  does not, and  $t_1$  and  $t_2$  agree on all attributes except attribute  $x$ . Each of  $t_1$  and  $t_2$  can now be extended into an instance of  $D$ . Both instances will give the same instance of the derived relation, consisting of a single tuple  $t_1[A]$  (or  $t_2[A]$ ). In one instance, the tuple in the derived relation should be modified, in the other one it should not. The decision depends on the value of attribute  $x$  which is not visible in the derived relation. Hence the decision cannot be made without additional data.  $\square$

**Corollary 6B.1:** If  $\alpha(C_M) \subseteq A^+$  then the condition  $C_M$  is autonomously computable.  $\square$

**Theorem 6C:** The decision whether or not a tuple in  $E$  chosen for modification by  $C_M$ , will still satisfy  $C$  after modification is guaranteed to be autonomously computable if and only if every attribute in  $\alpha(C(F_M)) - A$  is computationally nonessential in  $C(F_M)$  with respect to the condition  $C \wedge C_M \wedge C_B(F_M)$ .

**Proof:** (Sufficiency) If every attribute  $x \in (\alpha(C(F_M)) - A)$  is computationally nonessential with respect to  $C \wedge C_M \wedge C_B(F_M)$ , we can correctly evaluate the condition by assigning surrogate values.

(Necessity) Without loss of generality assume that  $\alpha(C(F_M)) - A$  contains only a single attribute  $x$  and that  $x$  is computationally essential. In the same way as in the proof of Theorem 5, we can then construct two tuples  $t_1$  and  $t_2$  over the attributes in  $A \cup \alpha(C) \cup \alpha(C_M) \cup \alpha(C(F_M)) \cup \alpha(C_B(F_M))$ , such that they only differ in the value of  $x$ , and such that  $t_1$  and  $t_2$  satisfy  $C$ ,  $C_M$  and  $C_B(F_M)$ . Let  $t_1'$  and  $t_2'$  denote the corresponding tuples after modification. Because  $x \in (\alpha(C(F_M)) - A)$  then the tuples  $t_1'$  and  $t_2'$  must have different values for at least one attribute, and  $x$  must occur in the update expression for at least one attribute in  $\alpha(C)$ . One of them,  $t_1'$  say, will satisfy  $C$  while the other one will not. We can now extend  $t_1$  and  $t_2$  to obtain two different database instances where each relation contains only one tuple and where in both cases the derived relation contains the same tuple. In one case (for the instance obtained from  $t_2$ ) the single tuple in the derived relation should be deleted after the modification, while in the other case it should not. The decision depends on the value of  $x$ , which is not visible in the derived relation.  $\square$

**Corollary 6C.1:** If  $B_i = \rho(f_{B_i})$  for all  $B_i \in (\alpha(R_u) \cap \alpha(C))$  or if  $(\alpha(R_u) \cap \alpha(C)) = \emptyset$  then the MODIFY operation will not cause any deletions from the derived relation.

**Proof:** Similar to the proof of Corollary 6A.1.  $\square$

**Corollary 6C.2:** If  $\alpha(C(F_M)) \subseteq A^+$  then whether or not a tuple in  $E$ , chosen for modification by  $C_M$ , will satisfy  $C$  after modification is autonomously computable.  $\square$

Recall that  $\alpha(\rho(f_{B_i}))$  denotes the set of attributes occurring in the right hand side expression of  $f_{B_i}$ . Define the set  $Z$  as

$$Z = \bigcup_{B_i \in A} \alpha(\rho(f_{B_i}))$$

that is,  $Z$  is the set of attributes from which the new values for the attributes in  $A$  are computed.

**Theorem 6D:** For all tuples in the derived relation which are not to be deleted after modification, the new values for the attributes in  $A$  are guaranteed to be autonomously computable if and only if every attribute in  $Z - A$  is uniquely determined by the condition  $C \wedge C_M \wedge C(F_M) \wedge C_B(F_M)$  and the attributes in  $A$ .

**Proof:** (Sufficiency) Assume that every attribute in  $Z - A$  is uniquely determined by the condition  $C \wedge C_M \wedge C(F_M) \wedge C_B(F_M)$  and the attributes in  $A$ . Therefore, the value of every attribute in  $\alpha(\rho(f_{B_i}))$  can be uniquely determined and, hence, the modified values of attributes in  $A$  are autonomously computable.

(Necessity) Let the modifications be autonomously computable. Without loss of generality we can assume that  $Z - A$  consists of a single attribute  $x$ . Suppose that  $x$  is not uniquely determined by the condition  $C \wedge C_M \wedge C(F_M) \wedge C_B(F_M)$  or by the attributes in  $A$ . We can then construct two tuples  $t_1$  and  $t_2$  over the attributes in  $A \cup \alpha(C) \cup \alpha(C_M) \cup \alpha(C(F_M)) \cup \alpha(C_B(F_M))$  such that  $t_1$  and  $t_2$  both satisfy  $C, C_M, C(F_M)$ , and  $C_B(F_M)$ , and both tuples agree on the values of all attributes except  $x$ . Each of  $t_1$  and  $t_2$  can now be extended into an instance of  $D$ . Both instances will give the same instance of the derived relation, consisting of a single tuple  $t_1[A]$  (or  $t_2[A]$ ). In both instances the tuple in the derived relation should be modified. However, the values of those modified attributes which are calculated using  $x$  will be different depending on whether we use  $t_1$  or  $t_2$ . Hence, the values are not autonomously computable.  $\square$

**Corollary 6D.1:** If  $Z \subseteq A^+$  then the modified values are autonomously computable.  $\square$

**Corollary 6D.2:** If  $B_i = \rho(f_{B_i})$  for all  $B_i \in (\alpha(R_u) \cap A)$  then the tuples that remain in  $E$  will not have any attribute values changed.  $\square$

We give an example which proceeds through the four steps associated with Theorems 6A - 6D, at each step testing the appropriate condition.

*Example:* Suppose the database consists of the two relations  $R_1(H, I)$  and  $R_2(J, K)$  where  $H, I, J$  and  $K$  each have the domain  $[0, 30]$ . Let the derived relation and modify operation be defined as:

$$E = (\{I, J\}, \{R_1, R_2\}, (H > 10)(I = K)(J < 25))$$

MODIFY( $R_2, \{R_2\}, ((J < 15) \vee (J \geq 25))(K \geq 20)$ ),  $\{(J := K + 5), (K := K)\}$ ).

Step A:

$$\forall H, I, J, K [( \neg ((H > 10)(I = K)(J < 25)) )]$$

$$\begin{aligned}
& \wedge ((J < 15) \vee (J \geq 25)(K \geq 20)) \\
& \wedge ((K + 5 \leq 30)(K + 5 \geq 0)) \\
& \Rightarrow \neg((H > 10)(I = K)(K + 5 < 25))
\end{aligned}$$

Clearly, as far as testing validity is concerned, we need not consider the values of  $H$  and  $I$ . Therefore, assume  $(H > 10)$  and  $(I = K)$ . For the  $\neg C$  condition of the antecedent to evaluate to *true*, we must have  $(J < 25)$ . Hence for the  $C_M$  condition to evaluate to *true* we must also have  $(K \geq 20)$ . Therefore in the consequent  $(K + 5 < 25)$ , and hence the consequent is *true*. This shows that the given implication is valid. Therefore, the given modify operation will not introduce new tuples into  $E$ .

Step B:  $\alpha(C_M) - A = \{K\}$

$$\begin{aligned}
& \forall H, I, J, K, K' [((H > 10)(I = K)(J < 25))((H > 10)(I = K')(J < 25)) \\
& \Rightarrow (((J < 15) \vee (J \geq 25)(K \geq 20)) = ((J < 15) \vee (J \geq 25)(K' \geq 20)))]
\end{aligned}$$

For the antecedent to be *true*, we must have  $I = K = K'$ . Hence the two conditions in the consequent are equal regardless of the value of  $J$ . Therefore,  $K$  is computationally nonessential in  $C_M$  with respect to  $C$ , and  $C_M$  is autonomously computable.

Step C:  $\alpha(C(F_M)) - A = \{H, K\}$

$$\begin{aligned}
& \forall H, H', I, J, K, K' \\
& [((H > 10)(I = K)(J < 25))((J < 15) \vee (J \geq 25)(K \geq 20)) \\
& \wedge ((K + 5 \leq 30)(K + 5 \geq 0)(K \leq 30)(K \geq 0)) \\
& \wedge ((H' > 10)(I = K')(J < 25))(J < 15) \vee (J \geq 25)(K' \geq 20)) \\
& \wedge ((K' + 5 \leq 30)(K' + 5 \geq 0)(K' \leq 30)(K' \geq 0)) \\
& \Rightarrow (((H > 10)(I = K)(K + 5 < 25)) = ((H' > 10)(I = K')(K' + 5 < 25)))]
\end{aligned}$$

For the antecedent to be *true* we must have  $I = K = K'$  and  $H > 10$  and  $H' > 10$ . If so, the truth values of the two conditions of the consequent depend solely on  $K$  and  $K'$ . Because  $K = K'$  they will always have the same truth value. Therefore,  $H$  and  $K$  are computationally nonessential in  $C(F_M)$  with respect to  $C \wedge C_M \wedge C_B(F_M)$ , and  $C(F_M)$  is autonomously computable.

Step D:  $Z - A = \{K\}$

Since  $C$  guarantees that  $I = K$  in  $E$ , the value of  $K$  is uniquely determined by the value of  $I$  which is visible in  $E$ . Therefore, the new values of modified tuples in  $E$  are autonomously computable.

In summary, consider a numeric example for the given schema.

*Before*

$R_1$ :	$H$	$I$	$R_2$ :	$J$	$K$	$E$ :	$I$	$J$
	11	5		20	5		5	20
	11	15		10	15		15	10
	11	22		10	22		22	10
	11	20		30	20			

*After*

$R_1$ :	$H$	$I$	$R_2$ :	$J$	$K$	$E$ :	$I$	$J$
	11	5		20	5		5	20
	11	15		20	15		15	20
	11	22		27	22			
	11	20		25	20			

Step A provides a warrant that the last tuple, which does not satisfy  $C$  before modification, will not satisfy  $C$  after. Step B guarantees that we can determine which tuples of  $E$  to modify; the second and the third. Step C allows us to determine which modified tuples of  $E$  will be deleted since they will no longer satisfy condition  $C$ ; the third one. Step D ensures that we can compute the new values for the remaining modified tuple.  $\square$

## 6. Conclusion

Necessary and sufficient conditions for detecting when an update operation is irrelevant to a derived relation (or view, or integrity constraint) have not previously been available for any nontrivial class of updates and derived relations. The concept of autonomously computable updates is completely new. Limiting the class of derived relations to those defined by  $PSJ$ -expressions does not seem to be a severe restriction, at least not as it applies to structuring the stored database in a relational system. The class of update operations considered is fairly general. In particular, this seems to be one of a few papers on update processing where modify operations are considered explicitly and separately from insert and delete operations. Previously, modifications have commonly been treated as a sequence of deletions followed by insertion of the modified tuples.

Testing the conditions given in the theorems above is efficient in the sense that it does not require retrieval of any data from the database. According to our definitions, if an update is irrelevant or autonomously computable, then it is so for *every* instance of the base relations. The fact that an update is not irrelevant does not mean that it will always affect the derived relation. Determining whether or not it will, requires checking the current instance. The same applies for autonomously computable updates.

It should be emphasized that the theorems hold for any class of Boolean expressions. However, actual testing of the conditions requires an algorithm for proving the satisfiability of Boolean expressions. Currently, efficient algorithms exist only for a restricted class of expressions, the main restriction being on the atomic conditions allowed. An important open problem is to find efficient

algorithms for more general types of atomic conditions. The core of such an algorithm is a procedure for testing whether a set of inequalities/equalities can all be simultaneously satisfied. The complexity of such a procedure depends on the type of expressions (functions) allowed and the domains of the variables. If linear functions with variables ranging over the real numbers (integers) are allowed, the problem is equivalent to finding a feasible solution to a linear programming (integer programming) problem.

We have not imposed any restrictions on valid instances of base relations, for example, functional dependencies or inclusion dependencies. Any combination of attribute values drawn from their respective domains represents a valid tuple. Any set of valid tuples is a valid instance of a base relation. If relation instances are further restricted, then the given conditions are still sufficient, but they may not be necessary.

If an update is not autonomously computable some additional data may be required. An open problem is to determine the minimal amount of additional data required from the database, and how to retrieve it efficiently.

## Appendix

The theorems presented in this paper require that statements be proven at *run-time*, that is as updates are being performed on a particular database instance. What is required is that certain types of Boolean expressions be tested for unsatisfiability (or equivalently, tested for satisfiability) or that implications involving Boolean expressions be proven valid. The latter problem can be translated into one of showing that a Boolean expression is unsatisfiable. Hence, in either case we can proceed by testing satisfiability.

Rosenkrantz and Hunt [RH 80] gave an algorithm for testing the satisfiability of conjunctive Boolean expressions where the atomic conditions come from a restricted class. Their algorithm is based on Floyd's all-pairs-shortest-path algorithm and therefore has an  $O(n^3)$  worst case complexity, where  $n$  is the number of variables in the expression. The algorithm presented here is a modification of that given by Rosenkrantz and Hunt; there are three main differences. First, we assume that each variable has a finite domain whereas Rosenkrantz and Hunt allow infinite domains. Second, if the expression is satisfiable our algorithm not only verifies the satisfiability but also produces an assignment of values to the variables which satisfies the expression. Third, although the worst case complexity of our algorithm remains  $O(n^3)$  under certain circumstances (which are not unreasonable to expect) the complexity is reduced to  $O(n^2)$ .

The algorithm given here tests the satisfiability of a restricted class of Boolean expressions. Each variable is assumed to take its values from a finite, ordered set. Since there is an obvious mapping from such sets to the set of integers, we always assume that the domain consists of a finite interval of the integers. It is assumed that each Boolean expression,  $B$ , over the variables  $x_1, x_2, \dots, x_n$ , is in conjunctive form, i.e.  $B = B_1 \wedge B_2 \wedge \dots \wedge B_m$ , and that each atomic condition,  $B_i$ , is of the form  $(x_i \text{ op } x_j + c)$  or  $(x_i \text{ op } c)$  where  $\text{op} \in \{=, <, \leq, >, \geq\}$  and  $c$  is an integer constant.

The first step is to *normalize*  $B$  so that the resulting expression,  $N = N_1 \wedge N_2 \wedge \dots \wedge N_{m_N}$ , only has atomic conditions of the form  $(x_i \leq x_j + c)$ . Conditions of the form  $(x_i \text{ op } c)$  are handled by modifying the given domain bounds for the variable  $x_i$ .

We build a weighted, directed graph  $G = (V, E)$  representing  $N$ . Each variable in  $N$  is represented by a node in  $G$ . For the atomic condition  $(x_i \leq x_j + c)$  we construct an arc from node " $x_j$ " to node " $x_i$ " having weight  $c$ . Hence,  $|V| = (\text{the number of variables in } N)$  and  $|E| = m_N$ . The graph,  $G$ , can be *reduced* in size by removing nodes of in-degree zero. The justification for doing this is that if  $x_j$  is such a node then  $N$  does not have any conditions of the form  $(x_j \leq x_i + c)$ . Hence, the upper limit value of  $x_j$  is not constrained by the value assigned to any other variable. Therefore, we allow  $x_j$  to be assigned its (modified) upper bound. Also, for each node  $x_i$  in

$G$ , such that there is an arc of weight  $c$  from  $x_j$  to  $x_i$ , we replace the upper bound on  $x_i$  with  $\min\{(\text{upper bound on } x_i), (\text{upper bound on } x_j) + c\}$ . We can then remove node  $x_j$  and its incident arcs from  $G$ .

The graph is represented by an  $n \times n$  array  $A$  where, initially,  $A(i, j) = c$  if and only if  $N$  contains an expression of the form  $(x_i \leq x_j + c)$ . If two nodes do not have an arc between them the corresponding array entry is labelled with  $\infty$  (i.e. an arbitrarily large positive value). Performing on  $A$ , the operation which corresponds to graph reduction on  $G$ , may produce an array with some rows and columns which will be unused. We therefore include in our algorithm a *compaction* routine which moves all the relevant information remaining in  $A$  to the upper left-hand corner.

After compaction, any variable not represented by a row and column in  $A$  has been assigned a value which does not contravene any condition in  $N$ .  $A$  is used as the input to a modified version of Floyd's algorithm to determine either that  $N$  is unsatisfiable or to produce an assignment which satisfies  $N$ . The idea is that we give each remaining variable an initial *trial* value equal to its (modified) upper bound. At each iteration we adjust the values (downward) to reflect the current values in  $A$  and the previous set of trial values. The iterations continue until we find an assignment to the variables which satisfies  $N$  or until we determine that  $N$  is unsatisfiable. This takes at most  $|V| \leq n$  iterations.

To be more specific, given a graph with nodes  $x_1, \dots, x_n$ , the  $k$ th step of Floyd's algorithm produces the least weight path between each pair of nodes, with intermediate nodes from the set  $\{x_1, x_2, \dots, x_k\}$ . In terms of the Boolean expression this corresponds to forming, from the conditions in  $N$ , the most restrictive condition between each pair of variables. The only conditions of  $N$  which may be used at the  $k$ th step are those involving the variables  $x_1, \dots, x_k$ . The new trial value for  $x_i$  is found by taking the minimum of its previous trial value and  $\min\{(\text{previous trial value of } x_j) + A[i, j]\}$  for  $1 \leq j \leq k$ .

There are three possible situations that indicate that the algorithm should terminate. We test each of these conditions after each iteration:

1. Is there a negative weight cycle? In this case  $N$  is unsatisfiable
2. Does the current trial assignment violate any variable's lower bound? Again,  $N$  is unsatisfiable
3. Does the current trial assignment satisfy the lower bound for each variable and satisfy  $N$ ? In this case  $N$  is satisfiable.

Since the longest cycle can contain at most  $|V|$  arcs we conclude that this is the maximum number of iterations required. If after  $|V|$  iterations we



have not found a negative weight cycle or violated any lower bound then the current trial assignment must satisfy  $N$ .

The following are pseudo-code versions of the algorithms required to carry out this method. The first is Satisfiability which acts as a mainline for the entire procedure. It in turn uses a number of procedures which are listed in the order they are called in Satisfiability. They are: Normalize, Reduce, CheckBound, CheckExpr, Compact, TestA, and CalcTrial.

The following is a description of the variables used in the algorithms:

- $n$  the number of variables in  $B$
- $B$  a Boolean expression of the form  $B_1 \wedge \dots \wedge B_m$  over the variables  $x_1, \dots, x_n$ . Each atomic condition  $B_i$  is of the form  $(x_j \text{ op } x_k + c)$  or  $(x_j \text{ op } c)$  where  $\text{op} \in \{=, <, \leq, >, \geq\}$  and  $c$  is an integer
- $U$  and  $L$  are  $n$ -dimensional, integer vectors. Each  $x_i$  in  $B$  is assumed to have a finite interval of the integers as its domain. Initially,  $U[i]$  and  $L[i]$  give, respectively, the upper and lower domain bounds on  $x_i$ , with  $L[i] \leq U[i]$  for  $1 \leq i \leq n$ . As the algorithm proceeds these bounds may be adjusted to reflect the constraints imposed on the variables by the conditions in  $B$ . The modified entries in  $U$  are taken as the “trial” assignment of values to the  $x_i$ ’s. These entries may be further modified in the search for an assignment which satisfies  $B$ . At the termination of the procedure if  $B$  is found to be satisfiable then the assignment  $x_i := U[i]$  will satisfy  $B$
- $SAT$  is a boolean variable which is *true* if  $B$  is satisfiable, *false* otherwise
- $N$ , a Boolean expression of the form  $N_1 \wedge \dots \wedge N_p, p \leq m$ , over the variables  $x_1, \dots, x_n$ , is the “normalized” version of  $B$ . Each atomic condition  $N_i$  is of the form  $(x_j \leq x_k + c)$
- $A$  is an  $n \times n$  array of integers ( including  $\infty$  ) used initially to record a directed graph representing  $B$ . Subsequently,  $A$  is reduced and compacted by removing nodes of in-degree 0 from the graph. The compacted  $A$  is then used as input to Floyd’s algorithm to give the shortest paths in this graph
- $indeg$  is a  $1 \times n$  array of integers, where  $indeg[i]$  gives the number of non-infinity entries in row  $i$  of  $A$  (i.e. the in-degree of node  $x_i$  in the directed graph representing  $N$ ). A row of  $A$  which, after reduction, is no longer being considered has its  $indeg$  set to  $-1$
- $size$  gives the order of  $A$  after it has been “compacted”

- *row* is a  $1 \times n$  array of integers, where for  $1 \leq i \leq size$ , *row*[*i*] gives the row number of the corresponding row of *A* before compaction

Note:  $\infty$  represents a large positive number whose value, which is larger than that of any integer, remains unchanged by the addition or subtraction of arbitrary integers

Procedure

**Satisfiability**( $B, L, SAT, U, n$ )

**Input:**  $B, L, U, n$

**Output:**  $SAT, U$  where the assignment  $x_i := U[i]$  satisfies  $B$  if  $SAT$  is true

**Local:**  $A, N, indeg, row, size$

```

begin
  SAT := true
  Normalize( B, L, N, SAT, U )
  if SAT = false then return fi
  for i := 1 to n do /* initialize indeg and A */
    indeg[i] := 0
    for j := 1 to n do A[i,j] = ∞ od
  od
  for each (xj ≤ xk + c) in N do
  /* build the matrix for the directed graph representing N */
    if A[j, k] = ∞ then
      indeg[j] := indeg[j] + 1
      A[j, k] = c
    else if c < A[j, k] then
      A[j, k] = c
    fi
  od
  /* "remove" rows representing nodes of in-degree zero */
  Reduce(A, L, SAT, U, indeg, n)
  if SAT = false then return fi
  size := n /* initialize size and row */
  for i := 1 to size do row[i] := i od
  /* test the "trial" values against the lower bounds */
  CheckBound( L, SAT, U, row, size )
  if SAT = false then return fi
  CheckExpr( N, SAT, U ) /* test the "trial" values in N */
  if SAT = true then return fi
  /* move the remaining rows of A to the upper left-hand corner */
  Compact(A, indeg, n, row, size)
  /* run Floyd's Algorithm (with some added tests) on the compacted A */
  for k := 1 to size do
    for i := 1 to size do
      for j := 1 to size do
        if A[i, j] > A[i, k] + A[k, j] then

```

```

        A[i, j] := A[i, k] + A[k, j]
      fi
    od
  od
  TestA(A, SAT, size) /* check for negative cycles */
  if SAT = false then return fi
  CalcTrial(A, U, row, size) /* calculate the new "trial" values */
  CheckBound( L, SAT, U, row, size )
  /* test the new "trial" values against the lower bounds */
  if SAT = false then return fi
  CheckExpr( N, SAT, U ) /* test the new "trial" values in N */
  if SAT = true then return fi
od
return
end

```

---

Procedure

**Normalize**(*B, L, N, SAT, U*)

**Input:** *B, L, U*

**Output:** *L, N, SAT, U*

**begin**

*SAT* := *true*

*N* := *true* /\* a trivial atomic condition \*/

**for each** *B<sub>i</sub>* **in** *B* **do**

**case** *B<sub>i</sub>* **of form**

    (*x<sub>j</sub> op c*) : **case** *op* **of**

      ≥ : **if** (*c* > *U*[*j*]) **then** *SAT* := *false* **return**

**else if** (*c* > *L*[*j*]) **then** *L*[*j*] := *c*

**fi**

      > : **if** (*c* ≥ *U*[*j*]) **then** *SAT* := *false* **return**

**else if** (*c* + 1 > *L*[*j*]) **then** *L*[*j*] := *c* + 1

**fi**

      ≤ : **if** (*c* < *L*[*j*]) **then** *SAT* := *false* **return**

**else if** (*c* < *U*[*j*]) **then** *U*[*j*] := *c*

**fi**

      < : **if** (*c* ≤ *L*[*j*]) **then** *SAT* := *false* **return**

**else if** (*c* - 1 < *U*[*j*]) **then** *U*[*j*] := *c* - 1

**fi**

```

= : if (c < L[j]) or (c > U[j]) then SAT := false return
    else U[j] := L[j] := c fi
endcase
(xj op xk + c) : case op of
    ≥ : N := (xk ≤ xj - c) ∧ N
    > : N := (xk ≤ xj - (c + 1)) ∧ N
    < : N := (xj ≤ xk + (c - 1)) ∧ N
    ≤ : N := (xj ≤ xk + c) ∧ N
    = : N := (xj ≤ xk + c) ∧ (xk ≤ xj - c) ∧ N
endcase
endcase
od
return
end

```

---

Procedure

**Reduce**(*A, L, SAT, U, indeg, n*)

**Input:** *A, L, U, indeg, n*

**Output:** *A, SAT, U, indeg*

**Local:**

- *currchanged* is an integer variable used to indicate the number of the last row to have its *indeg* decremented during the current pass of *A*
- *prevchanged* is an integer variable used to indicate the number of the last row to have its *indeg* decremented during the previous pass of *A*

**begin**

*prevchanged* := *n*

*SAT* := *true*

**while** *prevchanged* > 0 **do**

*currchanged* := 0

**for** *i* := 1 to *prevchanged* **do**

**if** *indeg*[*i*] = 0 **then**

**for** *j* := 1 to *n* **do**

**if**  $U[j] > U[i] + A[j, i]$  **then**

$U[j] := U[i] + A[j, i]$

**if**  $U[j] < L[j]$  **then**

*SAT* := *false*

```

        return
      fi
    fi
     $indeg[j] := indeg[j] - 1$ 
     $currchanged := j$ 
  od
   $indeg[i] := -1$ 
fi
od
 $prevchanged := currchanged$ 
od
return
end

```

---

Procedure  
**CheckBound**( $L, SAT, U, row, size$ )

**Input:**  $L, U, row, size$

**Output:**  $SAT$

```

begin
   $SAT := true$ 
  for  $i := 1$  to  $size$  do
    if  $U[row[i]] < L[row[i]]$  then  $SAT := false$  return fi
  od
  return
end

```

---

Procedure  
**CheckExpr**( $N, SAT, U$ )

**Input:**  $N, U$

**Output:**  $SAT$

```

begin
   $SAT := true$ 
  for each  $(x_j \leq x_k + c)$  in  $N$  do
    if  $U[j] > U[k] + c$  then  $SAT := false$  return fi
  od
end

```

```

    od
  return
end

```

---

Procedure  
**Compact**( $A, indeg, n, row, size$ )

**Input:**  $A, indeg, n$

**Output:**  $A, row, size$

```

begin
  size := 0
  for i := 1 to n do
    if indeg[i] > 0 then
      size := size + 1
      row[size] := i
    fi
  od
  for i := 1 to size do
    for j := 1 to size do
      A[i, j] := A[row[i], row[j]]
    od
  od
  return
end

```

---

Procedure  
**TestA**( $A, SAT, size$ )

**Input:**  $A, size$

**Output:**  $SAT$

```

begin
  SAT := true
  for i := 1 to size do
    if A[i, i] < 0 then SAT := false return fi
  od
  return
end

```

---

Procedure

**CalcTrial**( $A, U, row, size$ )

**Input:**  $A, U, row, size$

**Output:**  $U$

**Local:**  $T$  is an  $1 \times size$  array of integers used for temporary storage

**begin**

**for**  $i := 1$  **to**  $size$  **do**

$T[i] = \min\{U[row[i]], \min_{1 \leq k \leq size} \{U[row[k]] + A[i, k]\}\}$

**od**

**for**  $i := 1$  **to**  $size$  **do**

$U[row[i]] := T[i]$

**od**

**return**

**end**



**References**

- [A 75] ANSI/X3/SPARC Study Group on Database Management Systems, Interim Report, FDT (ACM SIGMOD bulletin), 7, 2, (1975).
- [AL 80] Adiba, M., and Lindsay, B.G., "Database Snapshots," Proc. 6th International Conf. on Very Large Databases, (Montreal, 1980), 86-91.
- [BLT 86] Blakeley, J.A., Larson, P.-Å., and Tompa, F.W., "Efficiently Updating Materialized Views," SIGMOD 1986, (to appear).
- [BC 79] Buneman, O.P., and Clemons, E.K., "Efficiently Monitoring Relational Databases," ACM Trans. on Database Systems, 4, 3 (1979), 368-382.
- [GSV 84] Gardarin, G., Simon, E., and Verlaine, E., "Querying Real Time Relational Data Bases," IEEE-ICC International Conference (Amsterdam, 1984), 757-761.
- [HS 78] Hammer, M. and Sarin, S.K., "Efficiently Monitoring of Database Assertions," Supplement, Proc. ACM SIGMOD International Conf. on Management of Data, (Austin, TX, 1978), 38-48.
- [LY 85] Larson, P.-Å. and Yang, H.Z., "Computing Queries from Derived Relations," Proc. 11th International Conf. on Very Large Databases, (Stockholm, 1985), 259-269.
- [L 86] Lindsay, B., et.al., "A Snapshot Differential Refresh Algorithm," Research Report RJ 4992, IBM Almaden Research Center (1986).
- [M 83] Maier, D., *The Theory of Relational Databases*, Computer Science Press, Rockville, MD, 1983.
- [RH 80] Rosenkrantz, D.J. and Hunt, H.B. III, "Processing Conjunctive Predicates and Queries," Proc. 6th International Conf. on Very Large Data Bases, (Montreal, 1980), 64-72.