

14

UPDATING LOGICAL DATABASES

Ronald Fagin, Gabriel M. Kuper,
Jeffrey D. Ullman, and Moshe Y. Vardi

ABSTRACT

We suggest a new approach to database updates, in which a database is treated as a collection of theories. We investigate two issues: simultaneous multiple update operations and equivalence of databases under update operations.

1. INTRODUCTION

One of the main problems in database theory is the problem of view updating, i.e., how to translate an update on a user view into an update of the database [1-4, 6-8, 13]. The problem is that in general there is no unique database update corresponding to the view update. Another

Advances in Computing Research, Volume 3, pages 1-18
Copyright © 1986 by JAI Press Inc.
All rights of reproduction in any form reserved.
ISBN: 0-89232-611-5

problem is that of updating a database that must satisfy certain integrity constraints [12, 15]. The difficulty here is that the database after the update may no longer satisfy the constraints, in which case we may have to modify other things in the database to ensure that the integrity constraints still hold. As in the case of view updates, there is not necessarily a unique way to modify the database so that the constraints still hold.

Fagin et al. [5] suggest that the appropriate framework for studying the semantics of updates is to treat the database as a consistent set of sentences in first-order logic, i.e., a *theory*. A theory is a description of the world, but is not necessarily a complete description; every model of the theory is a possible state of the world. Thus the database can be viewed as an exact description of our knowledge about the world. This framework was propounded in other papers (e.g., [9, 11, 14]).

When one tries to update a theory by inserting or deleting some sentence, several new theories can accomplish the update. Fagin et al. [5] argue that we should try to minimize the change that is needed to accomplish the update. Unfortunately, even under this minimality constraint, there may be several theories that accomplish the update, with no reasonable way to choose among them. One approach to this, suggested in [5], is to define the result of the update to be the disjunction of all the possible theories that accomplish the update with minimal change. Two difficulties with this approach are that it requires us to have sentences of a rather complicated syntax (e.g., disjunctions of tuples in a relational database) and that the number of sentences in the database may grow doubly exponentially with each update.

The fact that several theories can accomplish a given update motivates an alternative approach: viewing the database as a collection of theories rather than a single theory. We call a collection of theories a *flock*. The advantage of this approach is that it is easier to deal with the multiplicity of flocks than with the multiplicity of theories. With the new approach, the sentences we get are of no greater complexity than those that were in the database or those that were inserted, and the number of sentences does not grow as fast as before.

In this paper, after presenting the two approaches to updates, databases as theories vs. databases as flocks, we investigate two basic issues. First, we study *batch operations*, in which many sentences, rather than a single sentence, are inserted or deleted simultaneously. We then observe that two theories or flocks that are logically equivalent may not be equivalent after an update is performed. We give necessary and sufficient conditions for *equivalence forever*, i.e., equivalence that is preserved under updates.

2. UPDATES OF THEORIES

Our basic units of information are *sentences*, i.e., formulas without free variables in some first-order logic. We do not allow inconsistent sentences, and we do not allow the deletion of valid sentences. A *theory* is a consistent set of sentences. We shall use the letters S and T to denote theories, and the letters σ and τ to denote sentences.

We start by describing the framework developed in [5]. (We use \subseteq to denote inclusion and \subset to denote proper inclusion.)

Definition 1

1. A theory T *accomplishes the deletion* of σ from S if $T \not\vdash \sigma$.
2. A theory T *accomplishes the insertion* of σ into S if $\sigma \in T$.

Definition 2. Let T_1 , T_2 , and T be theories.

1. T_1 has *fewer insertions* than T_2 with respect to T if $T_1 - T \subset T_2 - T$.
2. T_1 has *fewer deletions* than T_2 with respect to T if $T - T_1 \subset T - T_2$.
3. T_1 has *fewer changes* than T_2 , with respect to T , if T_1 has fewer deletions than T_2 , or T_1 and T_2 have the same deletions ($T - T_1 = T - T_2$) and T_1 has fewer insertions than T_2 .

Definition 3. A theory T *accomplishes an update* (i.e., a deletion or an insertion) u of S *minimally* if T accomplishes u and there is no theory T' that accomplishes u and has fewer changes than T with respect to S .

THEOREM 1. [5] *Let S and T be theories and let σ be a sentence. Then*

1. T *accomplishes the deletion of σ from S minimally iff T is a maximal subset of S that is consistent with $\neg\sigma$.*
2. $T \cup \{\sigma\}$ *accomplishes the insertion of σ into S minimally iff T is a maximal subset of S that is consistent with σ .*

There could be many theories that accomplish an update minimally. Suppose that T_1, \dots, T_n are the theories that accomplish an update u of S minimally. It is argued in [5] that the result of u should be a theory T such that

$$\text{Mod}(T) = \bigcup_{1 \leq i \leq n} \text{Mod}(T_i)$$

where $\text{Mod}(S)$ is the set of models of the theory S .

Definition 4. Let T_1, \dots, T_n be theories. The *disjunction* of these theories is defined to be the theory

$$\bigvee_{1 \leq i \leq n} T_i = \{\tau_1 \vee \dots \vee \tau_n \mid \tau_i \in T_i, 1 \leq i \leq n\}$$

It is shown in [5] that

$$\text{Mod}\left(\bigvee_{1 \leq i \leq n} T_i\right) = \bigcup_{1 \leq i \leq n} \text{Mod}(T_i)$$

Thus they suggest that if T_1, \dots, T_n are the theories that accomplish an update u minimally, then the result of u should be $\bigvee_{1 \leq i \leq n} T_i$.

3. FLOCKS

In this section, we shall describe another approach to updates, namely using collections of theories. We call these collections *flocks*. The intuitive idea is that since we have many possible theories that accomplish an update minimally, we reflect this ambiguity by keeping all these theories.

Definition 5. A flock S is a set of theories. The models of S are

$$\text{Mod}(S) = \bigcup_{S \in S} \text{Mod}(S)$$

An update still consists of the insertion or deletion of a single sentence. To update a flock we have to update each theory in the flock. Formally:

Definition 6. Let $S = \{S_1, \dots, S_n\}$ be a flock. A flock $T = \{T_1, \dots, T_n\}$ accomplishes an update u of S minimally if T_i accomplishes the update of S_i minimally for $1 \leq i \leq n$.

Again, there could be many flocks that accomplish an update minimally. Suppose that T_1, \dots, T_n are the flocks that accomplish an update u of S minimally. As in [5] we contend that the result of u should be a flock T such that

$$\text{Mod}(T) = \bigcup_{1 \leq i \leq n} \text{Mod}(T_i)$$

It is easy to show that the flock $\bigcup_{1 \leq i \leq n} T_i$ has this property. This motivates the following definition:

Definition 7. Let S be a flock, and let S_1, \dots, S_n be the flocks that accomplish an update u of S minimally. Then the result of u is the flock $\bigcup_{1 \leq i \leq n} S_i$.

LEMMA 2. Let $S = \{S_1, \dots, S_n\}$ be a flock. For each theory S_i , let $S_i^1, \dots, S_i^{j_i}$ be the theories that accomplish the update u of S_i minimally. Then the result of applying u to S is the flock

$$S' = \{S_i^k \mid 1 \leq i \leq n, 1 \leq k \leq j_i\}$$

PROOF. Let S' be the result of the update. If $S \in S'$ then, by Definition 7, $S \in S_j$ for some S_j that accomplishes the update minimally. But then, by Definition 6, S accomplishes the update of some $S_i \in S$ minimally. i.e., S is one of the theories $S_i^1, \dots, S_i^{j_i}$.

Now let $S = S_i^k$ for some k , $1 \leq k \leq j_i$. Then S accomplishes the update u of S_i minimally. For each j , $1 \leq j \leq n$, $j \neq i$, let S^j be any theory that accomplishes the update u of S_j minimally. Then, by Definition 6, the flock $\{S^1, \dots, S^{j-1}, S, S^{j+1}, \dots, S^n\}$ accomplishes the update u of S minimally and so, by Definition 7, each theory in this flock is in S' . In particular, $S \in S'$. \square

In other words, to update a flock, consider each theory in the flock in turn. Take all theories that accomplish the update minimally and put them in the new flock.

Note that if a flock is a singleton, i.e., contains exactly one theory, its models as a theory and as a flock are the same. Also, the flock that we get after applying an update to such a flock has the same models as the theory we get by applying the update to the single member of that flock, as the following lemma shows.

LEMMA 3. Let $S = \{S\}$ be a singleton flock, and u an update. If S' is the result of applying the update u to the theory S , and S' is the result of applying u to the flock S , then S' and S' have the same models.

PROOF. By Lemma 2, the result of applying u to S is the flock $S' = \{S_i \mid 1 \leq i \leq j\}$, where S_1, \dots, S_j are the theories that accomplish the update u of S minimally. Similarly, the result of applying u to S is the theory $S' = \bigvee_{1 \leq i \leq j} S_i$. By Definition 5 and the comments at the end of Section 2, S' and S' have the same models. \square

Even though the result of an update has the same models under both approaches, under future updates their results may differ, as the following example shows.

Example 1. If we start with the flock $\{\{A, B\}\}$, and delete $A \wedge B$ from it using the flocks approach, we take all the maximal subtheories of $\{A, B\}$ that do not imply $A \wedge B$, namely $\{A\}$ and $\{B\}$. That is, the resulting flock is $\{\{A\}, \{B\}\}$. If we now delete A and then delete B , we end up with the flock containing only the empty theory, i.e., anything is a model of the result. On the other hand, if we start with the theory $\{A, B\}$, and delete $A \wedge B$, we get the theory $\{A \vee B\}$. This has the same models as the flock $\{\{A\}, \{B\}\}$. However, if we now delete A and then delete B , we still have the theory $\{A \vee B\}$, which does not have the same models as the empty theory. \square

In practice, singleton flocks are the most likely to be used as the starting state of the database (in fact, the starting state will probably be $\{\emptyset\}$). It would be interesting to characterize the flocks that are obtained from singleton flocks by a sequence of update operations. Another interesting question is the comparative merit of the two approaches: theories vs. flocks. We know that these approaches yield different results for the same updates. Which one of them is closer to correct?

4. BATCH OPERATIONS

Batch operations consist of deleting or inserting several sentences simultaneously.

Definition 8. Let S be a theory and let Σ be a set of sentences. We say that S' accomplishes the deletion of Σ from S if $S' \not\models \sigma$ for each $\sigma \in \Sigma$. We say that S' accomplishes the insertion of Σ into S if $\Sigma \subseteq S'$. We say that S' accomplishes an update u of S minimally if S' accomplishes u and there is no theory that accomplishes u with fewer changes.

The above definition is nonconstructive in the sense that it does not explicitly say how to find those theories that accomplish an update minimally. The following theorem gives a constructive equivalent condition, which generalizes Theorem 1.

THEOREM 4. *Let S and T be theories and Σ a set of sentences. Then*

1. *T accomplishes the deletion of Σ from S minimally iff T is a maximal subset of S such that $T \cup \{\neg\sigma\}$ is consistent for all σ in Σ .*
2. *$T \cup \Sigma$ accomplishes the insertion of Σ into S minimally iff T is a maximal subset of S that is consistent with Σ .*

PROOF

1. If T is a maximal subset of S that is consistent with $\neg\sigma$ for every $\sigma \in \Sigma$, then clearly T accomplishes the deletion of Σ from S . Assume that T does not accomplish the deletion minimally, i.e., there is a theory T' that accomplishes the deletion with fewer changes than T with respect to S . If T' has fewer deletions than T , then $S - T' \subset S - T$. But then $T' \cap S$ is also consistent with $\neg\sigma$, for all σ in Σ , contrary to the maximality of T . Therefore T' must have the same deletions as T with respect to S . Clearly, T' cannot have fewer insertions than T , since T has no insertions at all.

If T accomplishes the deletion minimally, it must be consistent with $\neg\sigma$ for every $\sigma \in \Sigma$. It is also clear that $T \subseteq S$ since if it contained sentences not in S we could remove them and get a theory that accomplished the update with the same deletions and with fewer insertions. If T is not a maximal subset of S that is consistent with all the $\neg\sigma$'s, then there is a theory that accomplishes the update with fewer deletions than T .

2. Let T be a maximal subset of S that is consistent with Σ . $T \cup \Sigma$ clearly accomplishes the insertion of Σ . Suppose that T' accomplishes the update with fewer deletions, and let $T'' = T' \cap S$. Then $S - T'' = S - T' \subset S - T$, and therefore $T \subset T'' \subseteq S$ and T'' is consistent with Σ , contradicting the maximality of T . Clearly, no theory can accomplish the insertion with the same deletions and with fewer insertions than $T \cup \Sigma$, since the only insertions are Σ .

If $T \cup \Sigma$ accomplishes the insertion of Σ minimally, we must have $T \subseteq S$ and T consistent with Σ . If T is not a maximal subset of S that was consistent with Σ , then we can find T' consistent with Σ that satisfies $T \subset T' \subseteq S$. But then $T' \cup \Sigma$ accomplishes the insertion with fewer deletions. This is a contradiction. \square

Using Definition 8, we can define the result of batch updates both for theories and for flocks. For theories, we define the result of the update to be the disjunction of all the theories that accomplish the update minimally, as in Definition 4. For flocks, we use Definitions 6 and 7. Namely, to update a flock consider each theory in the flock in turn, take all theories that accomplish the update of this theory minimally, and put them into the new flock. In the sequel, we reserve the term *update* (respectively, *deletion*, *insertion*) for the case where a single sentence is deleted or inserted, to distinguish it from *batch update* (respectively, *batch deletion*, *batch insertion*), where a set of sentences is deleted or inserted.

The following example shows that the batch deletion of Σ does not always give the same result as deleting the sentences in Σ one by one.

Example 2. Deleting $\{A, B\}$ from the theory $\{A, B, A \equiv B\}$ results in the theory $\{A \equiv B\}$. If, on the other hand, we delete first A , we get the theory $\{B \vee (A \equiv B)\}$, which remains unchanged after deleting B . Deleting first B and then A gives us the theory $\{A \vee (A \equiv B)\}$.

Deleting $\{A, B\}$ from the flock $\{\{A, B, A \equiv B\}\}$ results in the flock $\{\{A \equiv B\}\}$. If, on the other hand, we first delete A we get the flock $\{\{B\}, \{A \equiv B\}\}$, and if we then delete B we end up with the flock $\{\emptyset, \{A \equiv B\}\}$. This is different from the flock $\{\{A \equiv B\}\}$, since the union of the models of the first flock consists of all possible structures, whereas the models of $\{\{A \equiv B\}\}$ are only those models in which A and B are equivalent. \square

Similarly, the insertion of Σ does not give the same result as inserting the sentences in Σ one by one. The following theorem shows, however, that for flocks, batch insertions can be simulated by single updates.

THEOREM 5. *Let $\Sigma = \{\sigma_1, \dots, \sigma_n\}$ be a consistent set of sentences and let S be a flock. Then the result of inserting Σ into S is the same as first deleting $\neg(\sigma_1 \wedge \dots \wedge \sigma_n)$ and then inserting the σ_i 's one by one.*

PROOF. A theory S is consistent with Σ iff $S \not\models \neg(\sigma_1 \wedge \dots \wedge \sigma_n)$. Let \mathbf{T} be the result of deleting $\neg(\sigma_1 \wedge \dots \wedge \sigma_n)$ from S . We claim that $S' = \{T \cup \Sigma \mid T \in \mathbf{T}\}$ is the result of inserting Σ into S .

First, let T be a theory in \mathbf{T} . It accomplishes the deletion of $\neg(\sigma_1 \wedge \dots \wedge \sigma_n)$ from some $S \in \mathbf{S}$ minimally. We claim that $T \cup \Sigma$ accomplishes the insertion of Σ into S minimally. It is clear that $T \cup \Sigma$ accomplishes the insertion of Σ . If T' accomplishes the insertion with fewer deletions than $T \cup \Sigma$, then T' also accomplishes the deletion of $\neg(\sigma_1 \wedge \dots \wedge \sigma_n)$ from S with fewer deletions than T with respect to S , a contradiction. Clearly no theory can accomplish the insertion of Σ into S with fewer insertions than $T \cup \Sigma$ with respect to S , since the only insertions here are the sentences of Σ . This shows that each theory in S' is in the result of inserting Σ into S .

Now let T be a theory in the result of inserting Σ into S , i.e., T accomplishes the insertion of Σ into some $S \in \mathbf{S}$ minimally. Let $T' = T - (\Sigma - S)$. Then T' is consistent with Σ and so T' accomplishes the deletion of $\neg(\sigma_1 \wedge \dots \wedge \sigma_n)$ from S . If some theory S' accomplishes the deletion with fewer deletions than T' with respect to S , then $S' \cup \Sigma$ accomplishes the insertion of Σ with fewer deletions than T , a contradic-

tion. Therefore, T' accomplishes the deletion of $\neg(\sigma_1 \wedge \dots \wedge \sigma_n)$ from S minimally, i.e., $T' \in \mathbf{T}$, and therefore $T = T' \cup \Sigma$ is in S' . \square

Remarks

1. The theorem does not hold for theories. For example, let S be the theory $\{A, B\}$, and Σ the set $\{(A \neq B) \wedge C\}$. Then the result of inserting Σ into S is the theory

$$\{A \vee ((A \neq B) \wedge C), B \vee ((A \neq B) \wedge C), A \vee B, (A \neq B) \wedge C\}$$

On the other hand, the result of deleting $(A \equiv B) \vee \neg C$ from S is the theory $\{A \vee B\}$, and if we then just insert $(A \neq B) \wedge C$, we get the theory $\{A \vee B, (A \neq B) \wedge C\}$.

2. There are batch deletions from flocks that cannot be simulated by any sequence of single updates. For example, if we delete $\{A, B\}$ from the flock $\{\{A \vee B, A \vee \neg B, \neg A \vee B\}\}$, we get the flock $\{\{A \vee B, \{A \vee \neg B, \neg A \vee B\}\}$. It is shown in [10] that the latter flock cannot be obtained from any singleton flock by single updates (the proof is too long to present here).

5. EQUIVALENCE FOREVER

5.1. Definitions

Two theories or flocks are logically equivalent if they have the same models. Nevertheless, this does not guarantee that they will continue to have the same models after any sequence of updates, as the next example shows.

Example 3. The two theories $\{B\}$ and $\{B, A \vee B\}$ are logically equivalent. However, if we delete B from both of them we get the nonequivalent theories \emptyset and $\{A \vee B\}$.

The two flocks $\{\{B\}\}$ and $\{\{B, A \vee B\}\}$ are logically equivalent. After deleting B from both of them we get the nonequivalent flocks $\{\emptyset\}$ and $\{\{A \vee B\}\}$. \square

We say that two theories or flocks are *equivalent forever* if after applying any sequence of updates we always get two theories or flocks that have the same models. In the rest of this section we supply characterizations for equivalence forever.

We use the following definition.

Definition 9. We say that a theory S covers a theory T iff every sentence τ in T is logically equivalent to a conjunction $\sigma_1 \wedge \cdots \wedge \sigma_n$ of sentences in S . (An empty conjunction is by convention valid.)

5.2. Equivalence Forever for Theories

THEOREM 6. *Let S and T be finite theories. The following are equivalent.*

1. S and T are equivalent forever under updates.
2. S and T are equivalent forever under batch updates.
3. S and T are equivalent forever under deletions.
4. S and T are equivalent forever under batch deletions.
5. Each subset of S is logically equivalent to a subset of T , and vice versa.
6. S covers T , and vice versa.

PROOF. (2) \Rightarrow (1), (2) \Rightarrow (4), (4) \Rightarrow (3), (5) \Rightarrow (6), and (1) \Rightarrow (3) are obvious. We shall show (3) \Rightarrow (6), (6) \Rightarrow (5), and (5) \Rightarrow (2).

(3) \Rightarrow (6) We shall prove the following statement, which we call statement (*), inductively on k :

(*) Let T_1 and T_2 be finite theories that are equivalent forever under deletions. If there is a structure M that obeys $\tau \in T_1$ and that also obeys exactly k sentences in T_2 , then τ is equivalent to a conjunction of sentences in T_2 .

Statement (*) implies (3) \Rightarrow (6). For, let $S = T_1$ and $T = T_2$, and let τ be an arbitrary member of S . Let M be a structure which obeys τ . (There is such a structure since we deal only with consistent sentences.) Since T is finite, there is some k (possibly $k = 0$) such that M obeys exactly k sentences in T . Then statement (*) tells us that τ is equivalent to a conjunction of sentences in T , as desired.

If T_1 and T_2 are finite theories and if M is a structure, it is convenient for us to define $\sigma(M, T_1, T_2)$ to be the sentence

$$\bigvee \{ \sigma \mid \sigma \in T_1 \cup T_2 \text{ and } M \text{ violates } \sigma \}$$

It is easy to see that there is a single maximal theory which results from deleting this sentence from T_1 , namely, the set of all sentences in T_1 which are true in M . Of course, the same is true about T_2 . We are now ready to prove statement (*), by induction on k .

$k = 0$: In this case, M obeys no sentence in T_2 . Let us denote by T'_1 (respectively, T'_2) the result of deleting $\sigma(M, T_1, T_2)$ from T_1 (respectively, T_2). Since M obeys no sentence in T_2 , it follows that T'_2 is the empty theory, which every structure obeys. Since T'_1 and T'_2 are equivalent (by equivalence forever of T_1 and T_2 under deletions), it follows that T'_1 consists of valid sentences. But τ belongs to T'_1 , since M obeys τ . It follows that τ is valid, and is therefore equivalent to a conjunction of sentences in T_2 .

Inductive step: Assume that the inductive hypothesis (*) holds, with k' substituted for k , for every $k' < k$, and for every choice of T_1 and T_2 . Let T_1 and T_2 be finite theories that are equivalent forever under deletions, and let M be a structure which obeys $\tau \in T_1$ and which also obeys exactly k sentences in T_2 . We must show that τ is equivalent to a conjunction of sentences in T_2 .

Let us denote by T'_1 (respectively, T'_2) the result of deleting $\sigma(M, T_1, T_2)$ from T_1 (respectively, T_2). Then T'_1 is a subset of T_1 which contains τ , and T'_2 is a subset of T_2 which contains exactly k sentences (namely, those sentences in T_2 which are true in M). By equivalence forever of T_1 and T_2 under deletions, we know that T'_1 and T'_2 are also equivalent forever under deletions. In particular, T'_1 and T'_2 are equivalent, and so T'_2 implies τ . If also τ were to imply T'_2 (that is, if τ were to imply every member of T'_2), then we would be done, since τ would be equivalent to the subset T'_2 of T_2 . So we can assume that τ does not imply T'_2 . Therefore, there is a structure M' which obeys τ but not T'_2 . Let k' be the number of members of T'_2 which M' obeys. Then $0 \leq k' < k$, since T'_2 contains k sentences, not all of which M' obeys. By inductive hypothesis (*), where T'_1 , T'_2 , and k' play the roles of T_1 , T_2 , and k respectively, it follows that τ is equivalent to a conjunction of members of T'_2 , and hence of T_2 .

(6) \Rightarrow (5) Let S' be a subset of S . For each $\sigma \in S'$, let T_σ be a subset of T such that σ is equivalent to the conjunction of members of T_σ . Let T' be the union of all sets T_σ where $\sigma \in S'$. We now show that S' is equivalent to T' . If $\tau \in T'$, find $\sigma \in S'$ such that $\tau \in T_\sigma$. Then σ implies τ , so S' implies τ . Hence, S' implies T' . Conversely, assume that $\sigma \in S'$. Then T_σ implies σ , and so T' implies σ . Hence T' implies S' .

(5) \Rightarrow (2) Assume that (5) holds. It suffices to show that if $S^{(1)}$ (respectively, $T^{(1)}$) is the result of applying a batch update u to S (respectively, T), then every subset of $S^{(1)}$ is equivalent to a subset of $T^{(1)}$ and vice versa.

Let the update u be the deletion of Σ (we shall remark at the end how to modify the proof to deal with the case where u is an insertion.) Let S' be a maximal subset of S that is consistent with

$\neg\sigma$, for every $\sigma \in \Sigma$. By assumption, there is a subset of T that is equivalent to S^\dagger . Let T^\dagger be a maximal such subset of T . We now show that T^\dagger is a maximal subset of T that is consistent with $\neg\sigma$, for every $\sigma \in \Sigma$. Clearly T^\dagger is consistent with every $\neg\sigma$, since T^\dagger is equivalent to S^\dagger , and S^\dagger is consistent with every $\neg\sigma$. If T^\dagger is not maximal, then find $T' \subsetneq T$ consistent with every $\neg\sigma$ such that $T^\dagger \subset T'$. By definition of T^\dagger , we know that T^\dagger is not equivalent to T' . By hypothesis, there is a subset of S that is equivalent to T' ; let S' be a maximal such subset. Since $S^\dagger \equiv T^\dagger \subset T' \equiv S'$, it follows that S' implies S^\dagger , and so $S' \cup S^\dagger$ is equivalent to S' . By maximality of S' , it follows that $S^\dagger \subseteq S'$. But $S^\dagger \neq S'$, since $S^\dagger \equiv T^\dagger \neq T' \equiv S'$. Hence, $S^\dagger \subset S'$. Since T' is consistent with $\neg\sigma$, for every $\sigma \in \Sigma$, and since $T' \equiv S'$, it follows that S' is consistent with every $\neg\sigma$. This contradicts maximality of S^\dagger .

Let us call each maximal subset of S (respectively, T) that is consistent with every $\neg\sigma$ an *S-candidate* (respectively, a *T-candidate*). We have shown that for each *S-candidate* S^\dagger there is a *T-candidate* T^\dagger such that $S^\dagger \equiv T^\dagger$. Similarly, for each *T-candidate* T^\dagger there is an *S-candidate* S^\dagger such that $S^\dagger \equiv T^\dagger$. Furthermore, this correspondence is bijective, since if there were two different *T-candidates* T^\dagger and T^\ddagger with $S^\dagger \equiv T^\dagger \equiv T^\ddagger$, then $T^\dagger \cup T^\ddagger$ would be a subset of S consistent with every $\neg\sigma$, contradicting the maximality of T^\dagger and T^\ddagger . That is, if S_1, \dots, S_n are all of the distinct *S-candidates*, then there is a listing T_1, \dots, T_n of all of the distinct *T-candidates* such that $S_i \equiv T_i$ for $1 \leq i \leq n$.

The result $S^{(1)}$ of the update on S is the theory $\bigvee \{S_i \mid 1 \leq i \leq n\}$, and analogously for $T^{(1)}$. We shall show $T^{(1)}$ covers $S^{(1)}$. As in the proof that (6) \Rightarrow (5), it then follows that every subset of $S^{(1)}$ is equivalent to a subset of $T^{(1)}$, as desired.

Let α be a member of $S^{(1)}$. We know that α is of the form $\alpha_1 \vee \dots \vee \alpha_n$, where $\alpha_i \in S_i$, for $1 \leq i \leq n$. By assumption, there is a subset T'_i of T which is equivalent to α_i . Since $\alpha_i \in S_i \equiv T_i$, it follows that T_i implies α_i , and hence T_i implies T'_i . So by the maximality of T_i , we know that $T'_i \subseteq T_i$. Let Q be the set $\{\tau_1 \vee \dots \vee \tau_n \mid \tau_i \in T'_i \text{ for } 1 \leq i \leq n\}$. Then $Q \subseteq T^{(1)}$. Let τ be the conjunction of the members of Q . The proof is complete if we show that α is equivalent to τ . Let τ'_i be the conjunction of members of T'_i , for $1 \leq i \leq n$, and let γ be the disjunction $\tau'_1 \vee \dots \vee \tau'_n$. Clearly τ'_i is equivalent to α_i , since both are equivalent to T'_i ($1 \leq i \leq n$). Hence, γ is equivalent to α . But τ is the conjunctive normal form of γ , and consequently τ is equivalent to α .

We close by remarking how the proof should be modified to deal with insertions rather than deletions. Assume that the update u is the

insertion of Σ . Let us call each maximal subset of S (respectively, T) that is consistent with Σ an *S-candidate* (respectively, a *T-candidate*). Just as before, it follows that if S_1, \dots, S_n are all of the distinct *S-candidates*, then there is a listing T_1, \dots, T_n of all of the distinct *T-candidates* such that $S_i \equiv T_i$, for $1 \leq i \leq n$. The result $S^{(1)}$ of the update on S is the theory $\bigvee \{S_i \cup \Sigma \mid 1 \leq i \leq n\}$, and analogously for $T^{(1)}$. Let α be a member of $S^{(1)}$. We know that α is of the form $\alpha_1 \vee \dots \vee \alpha_n$, where $\alpha_i \in S_i \cup \Sigma$, for $1 \leq i \leq n$. If no α_i is in Σ , then the proof proceeds as before. Assume now that some α_i is in Σ . For simplicity in description, assume that α_1 is in Σ , but $\alpha_i \in S_i$ for $2 \leq i \leq n$ (otherwise there is an obvious modification in the proof). As before, find a subset T'_i of T which is equivalent to α_i , for $i \geq 2$. Let Q be the set $\{\alpha_1 \vee \tau_2 \vee \dots \vee \tau_n \mid \tau_i \in T'_i \text{ for } 2 \leq i \leq n\}$. Then $Q \subseteq T^{(1)}$, and, as before, α is equivalent to the conjunction of members of Q . \square

5.3. Equivalence Forever of Flocks

We do not have, at present, a simple necessary and sufficient condition for equivalence forever of general flocks. However, for singleton flocks, i.e., flocks that contain only one theory, we can prove an analogue to Theorem 6.

THEOREM 7. *Let S and T be finite theories, and let $\mathbf{S} = \{S\}$ and $\mathbf{T} = \{T\}$ be singleton flocks. The following are equivalent.*

1. \mathbf{S} and \mathbf{T} are equivalent forever under updates.
2. \mathbf{S} and \mathbf{T} are equivalent forever under batch updates.
3. \mathbf{S} and \mathbf{T} are equivalent forever under deletions.
4. \mathbf{S} and \mathbf{T} are equivalent forever under batch deletions.
5. Each subset of S is logically equivalent to a subset of T , and vice versa.
6. S covers T and T covers S .

PROOF (2) \Rightarrow (1), (2) \Rightarrow (4), (4) \Rightarrow (3), (5) \Rightarrow (6) and (1) \Rightarrow (3) are obvious, and (6) \Rightarrow (5) was proven in Theorem 6. We now show (3) \Rightarrow (6) and (6) \Rightarrow (2).

(3) \Rightarrow (6) Assume that S does not cover T . Then there is a sentence τ in T that is not logically equivalent to any conjunction of sentences of S . Let Σ be the set of sentences in $S \cup T$ that are not implied by τ . Let R be the set of *maximal disjunctions* of sentences in Σ , i.e., the set R of all disjunctions of sentences in Σ such that if we add any

other sentence in Σ to the disjunction, the result is implied by τ . Formally, R consists of all sentences of the form $\sigma_1 \vee \cdots \vee \sigma_k$, where each σ_i is in Σ ,

$$\tau \not\models \sigma_1 \vee \cdots \vee \sigma_k$$

and if σ is any sentence in Σ distinct from all the σ_i 's, then

$$\tau \models \sigma_1 \vee \cdots \vee \sigma_k \vee \sigma$$

We now show that if we delete the sentences in R from the flock $S = \{S\}$, one by one, in any order, the resulting flock S' will be equivalent to $\{S - \Sigma\}$, and similarly deleting R from $T = \{T\}$ will result in $\{T - \Sigma\}$. We prove this for S , and an analogous proof holds for T .

Since no sentence in Σ is implied by τ , every σ in Σ can be extended to a maximal disjunction $\sigma \vee \sigma_1 \vee \cdots \vee \sigma_k$ that is in the set R . After deleting this disjunction, we get a flock of theories, none of which can contain any of the sentences $\sigma, \sigma_1, \dots, \sigma_k$. Therefore after deleting all of the sentences in R from S we get a flock S' of theories, each of which must be a subset of $S - \Sigma$.

We can show by induction on the number of deletions that the result is a singleton flock, consisting of one theory that is a superset of $S - \Sigma$. The basis for the induction is the initial flock S . We now show that if we have a flock consisting of one theory that is a superset of $S - \Sigma$ and a subset of S and we delete a sentence in the set R from it, we get a singleton flock that also consists of one theory that is a superset of $S - \Sigma$ and a subset of S .

Suppose that we have such a flock consisting of the theory S' and we delete from it a sentence $\sigma_1 \vee \cdots \vee \sigma_k$ of R . Let $\Sigma = \{\sigma_1, \dots, \sigma_k\} = \{\alpha_1, \dots, \alpha_m\}$. Since $\sigma_1 \vee \cdots \vee \sigma_k$ is maximal, τ implies $\sigma_1 \vee \cdots \vee \sigma_k \vee \alpha_i$ for $1 \leq i \leq m$, and consequently $\neg \alpha_i$ implies $\sigma_1 \vee \cdots \vee \sigma_k \vee \neg \tau$ for $1 \leq i \leq m$. Suppose that $\{\tau, \alpha_1, \dots, \alpha_m\}$ implies $\sigma_1 \vee \cdots \vee \sigma_k$, then τ implies $\neg \alpha_1 \vee \cdots \vee \neg \alpha_m \vee \sigma_1 \vee \cdots \vee \sigma_k$. But then it follows that τ implies $\sigma_1 \vee \cdots \vee \sigma_k$ —contradiction. But τ implies $S - \Sigma$ by definition and S' is a subset of S , so $(S' - \Sigma) \cup \{\alpha_1, \dots, \alpha_m\}$ does not imply $\sigma_1 \vee \cdots \vee \sigma_k$. Thus $S' - \{\sigma_1, \dots, \sigma_k\}$ is consistent with $\neg(\sigma_1 \vee \cdots \vee \sigma_k)$. It follows that the result of deleting $\sigma_1 \vee \cdots \vee \sigma_k$ from $\{S'\}$ is $\{S' - \{\sigma_1, \dots, \sigma_k\}\}$. This completes the induction and shows that the result of deleting R from $S = \{S\}$ is $S' = \{S - \Sigma\}$.

By the definition of Σ , we have $\tau \models S - \Sigma$, and therefore τ implies the conjunction of all the sentences in $S - \Sigma$. Since τ is not logically equivalent to any conjunction of a collection of sentences in S , it follows that $S - \Sigma \not\models \tau$. Therefore, there must be a model M of $S - \Sigma$ that is not a model of τ . Then M is a model of S' . However, since τ is

in $T - \Sigma$, M is not a model of T' . It follows that S and T are not equivalent forever under deletions.

(6) \Rightarrow (2) We show by induction on the number of updates that we always have

$$(\forall S' \in S')(\exists T' \in T')(S' \text{ covers } T' \wedge T' \text{ covers } S') \quad (1)$$

where S' and T' are the flocks we get from S and T by performing some updates. By our assumption, Condition 1 holds at the beginning, when both flocks are singletons.

Assume that Condition 1 holds after some insertions and deletions. We have to show that it continues to hold after deleting or inserting a set of sentences Σ . We first show this for deletion. We shall use S^1 and T^1 for the flocks before the deletion, S^2 and T^2 for the flocks afterwards.

Let S^2 be a theory in the flock S^2 . We first show that there is some theory in T^2 that covers S^2 . By the definition of deletion, S^2 must be a maximal subset of some theory S^1 in the flock S^1 that does not imply any sentence in Σ . By the inductive hypothesis, there is a theory T^1 in the flock T^1 such that S^1 covers T^1 , and T^1 covers S^1 . Let σ_i be any sentence in the theory S^2 . Since S^2 is a subset of S^1 and T^1 covers S^1 , there are sentences $\tau_{i1}, \dots, \tau_{im}$ in T^1 such that $\sigma_i \equiv \tau_{i1} \wedge \dots \wedge \tau_{im}$.

Let A be the set of all these τ_{ij} 's, for all σ_i 's in S^2 . We claim that A does not imply any sentence in Σ . Assume otherwise, i.e., $A \models \sigma$, for some σ in Σ . Since each σ_i in S^2 implies all the corresponding τ_{ij} 's in A , we have $S^2 \models A$, and therefore $S^2 \models \sigma$, a contradiction. Therefore A does not imply any sentence in Σ and can be extended to a maximal subset of T^1 with this property. Call the maximal subset T^2 . Since A covers S^2 , T^2 also covers S^2 . We shall now show that S^2 covers T^2 , thus completing the proof.

Let τ be any sentence in T^2 . We have to show that it is logically equivalent to a conjunction of sentences in S^2 . Since S^1 covers T^1 and T^2 is a subset of T^1 , there are $\sigma_1, \dots, \sigma_k$ in S^1 such that

$$\tau \equiv \sigma_1 \wedge \dots \wedge \sigma_k \quad (2)$$

We know that $T^2 \models S^2$, since T^2 covers S^2 . We also know that $T^2 \models \tau \models \sigma_i$, for each σ_i . If some σ_i were not in S^2 , the fact that S^2 is a maximal subset of S^1 not implying any sentence in Σ would entail that $S^2 \cup \{\sigma_i\}$ implies some sentence $\sigma \in \Sigma$. But then $T^2 \models \sigma$, a contradiction. This shows that each σ_i is in S^2 , and therefore S^2 covers T^2 .

Now let M be a model of some theory S' in the flock S' . By Condition 1, there is some theory T' in the flock T' , such that S'

covers T' . This implies that M is also a model of T' . Thus every model of S' is also a model of T' . Similarly, every model of T' is also a model of S' .

For insertion, note that by Theorem 5, inserting $\Sigma = \{\sigma_1, \dots, \sigma_k\}$ into a flock $S = \{S_1, \dots, S_n\}$ is the same as first deleting $\neg(\sigma_1 \wedge \dots \wedge \sigma_k)$ to get a flock $\{S'_1, \dots, S'_m\}$ and then inserting the sentences in Σ , all of which are consistent with the S'_i 's, and therefore the result is the flock $\{S'_1 \cup \Sigma, \dots, S'_m \cup \Sigma\}$. It is easy to see that Condition 1 is preserved by both of these steps. \square

Example 4. The flocks $\{A, B, A \wedge B\}$ and $\{A, B\}$ are equivalent forever. The flocks $\{A, B, A \vee B\}$ and $\{A, B\}$ are not equivalent forever. If we delete A and then B , we get $\{A \vee B\}$ from the first flock and $\{\emptyset\}$ from the second one.

For arbitrary flocks we only have a sufficient condition for equivalence forever.

THEOREM 8. *Let S and T be two flocks that satisfy the conditions*

$$(\forall S \in S)(\exists T \in T)(S \text{ covers } T \wedge T \text{ covers } S) \quad (3)$$

and

$$(\forall T \in T)(\exists S \in S)(T \text{ covers } S \wedge S \text{ covers } T) \quad (4)$$

Then S and T are equivalent forever.

PROOF. See proof of (6) \Rightarrow (2) in Theorem 7. \square

Remarks

1. By Theorem 6 we can replace " S covers T " in this theorem by the condition "for every subset of S , there is a logically equivalent subset of T ."
2. The above conditions are not necessary for equivalence forever. For example, it is shown in [10] that the two flocks

$$S = \{A, B, A \equiv B\}, \{A, A \equiv B\}, \{B, A \equiv B\}$$

and

$$T = \{A, A \equiv B\}, \{B, A \equiv B\}$$

are equivalent forever even though they do not satisfy Conditions 3 and 4.

6. CONCLUSIONS AND OPEN PROBLEMS

We have presented in this paper two different approaches to the problem of updating databases. Both approaches are based on looking at the database as a set of logical sentences, and investigating what happens when we insert or delete a fact. In one approach, the database is similar to a logical theory, but with the existence of a fact in the set having a greater significance than it being merely a logical consequence of them. In the second approach, a database is a set of theories, rather than a single theory.

Flocks have a smaller complexity—even though the number of sentences can still grow exponentially, the size of the individual sentences remains the same as in the original database. In particular, if the sentences represented individual tuples in a relational database, we would remain with tuples after the update, instead of getting disjunctions of tuples. Another advantage of flocks is that they seem to give better semantics in simple examples.

On the other hand, theories are mathematically more tractable. We do not have a characterization for equivalence forever for arbitrary flocks, nor can we tell if an arbitrary flock can come from a singleton flock, one that corresponds to a normal database. These problems do not arise if the database is regarded as a single theory.

ACKNOWLEDGMENTS

We would like to thank Paris Kanellakis for helpful comments on a previous draft of this paper. G.M.K.'s and J.D.U.'s work was supported by AFOSR grant 80-0212. M.Y.V.'s work was partly done while at Stanford University and was supported by a Weizmann Fellowship and AFOSR grant 80-0212. G.M.K.'s work was done while he was at Stanford University.

REFERENCES

1. Bancilhon F, Spyratos N: Update semantics of relational views. *ACM Trans Database Systems* 6:557-575, 1981.
2. Carlson CR, Arora AK: The updatability of relational views based on functional dependencies. *Proc IEEE COMPSAC*, pp 415-420, 1979.
3. Cosmadakis SS, Papadimitriou CH: Updates of relational views. *JACM* 31:742-760, 1984.
4. Dayal U, Bernstein PA: Translation of update operations on relational views. *ACM Trans Database Systems* 8:381-416, 1982.
5. Fagin R, Ullman JD, Vardi MY: On the semantics of updates in databases. *Proc 2nd ACM Symp on the Principles of Database Systems*, pp 352-365, 1983.

6. Jacobs BE: Application of database logic to the view update problem. Technical Report TR 960, Dept of Computer Science, University of Maryland, College Park, MD, 1980.
7. Keller AM: Updates to relational databases through views involving joins, in P Scheuermann (ed): *Improving Database Usability and Responsiveness*, Academic Press, New York, 1983, pp 363-384.
8. Klug AC: Theory of database mapping. Technical Report CSRG 98, Dept of Computer Science, University of Toronto, 1978.
9. Kowalski R: Logic as database language. Unpublished manuscript, Dept of Computing, Imperial College, London, 1981.
10. Kuper GM, Vardi MY: On updating flocks of theories. Technical Report, Stanford University, Stanford, CA, to appear.
11. Nicolas JM, Gallaire H: Database—theory vs. interpretation, in H Gallaire, J Minker (eds): *Logic and Databases*, Plenum Press, New York, pp 33-54, 1978.
12. Nicolas JM, Yazdanian K: Integrity checking in deductive databases, in H Gallaire, J Minker (eds): *Logic and Databases*, Plenum Press, New York, pp 325-344, 1978.
13. Osman IM: Updating defined relations. Natl Computer Conf, Vol. 48, AFIP Press, pp 733-740, 1979.
14. Reiter, R: Towards a logical reconstruction of relational database theory, in ML Brodie, J Mylopoulos, J Schmidt (eds): *On Conceptual Modelling: Perspectives from Artificial Intelligence, Databases, and Programming Languages*, Springer-Verlag, Berlin, pp 191-233, 1984.
15. Todd S: Automatic constraint maintenance and updating defined relations. *Proc IFIP 77*, North-Holland, Amsterdam, pp 145-148, 1977.