

# A Formal Approach to Recovery by Compensating Transactions \*

Henry F. Korth

Eliezer Levy

Abraham Silberschatz

Department of Computer Sciences  
University of Texas at Austin  
Austin, TX 78712 USA

## Abstract

Compensating transactions are intended to handle situations where it is required to undo either committed or uncommitted transactions that affect other transactions, without resorting to cascading aborts. This stands in sharp contrast to the standard approach to transaction recovery where cascading aborts are avoided by requiring transactions to read only committed data, and where committed transactions are treated as permanent and irreversible. We argue that this standard approach to recovery is not suitable for a wide range of advanced database applications, in particular those applications that incorporate long-duration or nested transactions. We show how compensating transactions can be effectively used to handle these types of applications. We present a model that allows the definition of a variety of types of correct compensation. These types of compensation range from traditional undo, at one extreme, to application-dependent, special-purpose compensating transactions, at the other extreme.

## 1 Introduction

The concept of transaction atomicity is the cornerstone of today's transaction management systems. Atomicity requires that an aborted transaction will have no effect on the state of the database. The most common method for achieving this is to maintain a recovery log

---

\*Work partially supported by a grant from the IBM Corporation, TARP grant 4355 and NSF grant IRI-8805215.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the VLDB copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Very Large Data Base Endowment. To copy otherwise, or to republish, requires a fee and/or special permission from the Endowment.

Proceedings of the 16th VLDB Conference  
Brisbane, Australia 1990

and provide the  $undo(T_i)$  operation which restores the data items updated by  $T_i$  to the value they had just prior to the execution of  $T_i$ . However, if some other transaction,  $T_j$ , has read data values written by  $T_i$ , undoing  $T_i$  is not sufficient. The (indirect) effects of  $T_i$  must be removed by aborting  $T_j$ . Aborting the affected transaction may trigger further aborts. This undesirable phenomenon, called *cascading aborts*, can result in uncontrollably many transactions being forced to abort because some other transaction happened to abort.

Since a committed transaction, by definition, cannot abort, it is required that if transaction  $T_j$  reads the values of data items written by transaction  $T_i$ , then  $T_j$  does not commit before  $T_i$  commits. A system that ensures this property is said to be *recoverable* [2]. One way of avoiding cascading aborts and ensuring recoverability is to prohibit transactions from reading *uncommitted* data values — those produced by transactions that have not committed yet. This principle has formed the basis for standard recovery in most contemporary database systems.

Unfortunately, there is a large range of database applications for which the standard recovery approach is excessively restrictive and even not appropriate. The common denominator of such applications is the need to allow transactions to read *uncommitted data* values.

In general, as indicated by Gray [6], early exposure of uncommitted data is essential in the realm of long-duration and nested transactions. When transactions are long-lived, it is unreasonable to prevent access to uncommitted data by forcing other transactions to wait until the updating transaction commits, since the wait will be of long duration. Also, long-duration and nested transactions are often used to model collaborative design activities [9]. In order to promote the cooperative nature of design environments, there is a need to expose incomplete (i.e., uncommitted) design objects. Such applications, and other that incorporate transactions of that nature, cannot be accommodated by the standard recovery approach since their executions entail cascading aborts and some of them are even non-recoverable.

An additional restriction imposed by standard recov-

ery is the inability to undo an already committed transaction. Suppose that a transaction was committed “erroneously.” By committed erroneously, we mean that from the system’s point of view there was nothing wrong with the committed transaction. However, external reasons, that were discovered later, rendered the decision to commit the transaction erroneous. Under the standard recovery approach there is no support for undoing such transactions.

This paper presents the method of a *compensating transactions* as a recovery mechanism in applications where exposure of uncommitted data and undoing of committed transactions must be facilitated. Our goals are to develop a better understanding of what compensation really is, when it is possible to employ it, and what the implications are on correctness of executions when compensation is used.

The remainder of this paper is organized as follows. We give an informal introduction to compensating transactions in Section 2. In Section 3, we present a transaction model suitable for the study of compensation. We then use this model in Section 4 to define criteria for “reasonable” compensation. After illustrating our definitions with examples in Section 5, we examine the theoretical consequences of our model in Section 6. Implementation issues are discussed in Section 7, and related work is described in Section 8.

## 2 Overview of Compensation

When the updates of a (committed or uncommitted) transaction  $T$  are read by some other transaction, we say that  $T$  has been *externalized*. The sole purpose of *compensation* is to handle situations where we want to undo an externalized transaction  $T$ , without resorting to cascading aborts. We refer to  $T$  as the *compensated-for transaction*. The transactions that are affected by (reading) the data values written by  $T$  are referred to as *dependent transactions* (of  $T$ ), and are referred to as a *set* using the notation  $dep(T)$ . The key point of our recovery paradigm is that we would like to *leave the effects of the dependent transactions intact* while preserving the consistency of the database, when undoing the compensated-for transaction. Compensation undoes  $T$ ’s effects in a *semantic* manner, rather than by physically restoring a prior state. All that is guaranteed by compensation is that a consistent state is established based on semantic information. This state may not be identical to the state that would have been reached, had the compensated-for transaction never taken place.

We propose the notion of *compensating transactions* as the vehicle for carrying out compensation. We use the notation  $CT$  to denote the compensating transaction for transaction  $T$ . A compensating transaction

has the fundamental properties of a transaction along with some special characteristics. It appears atomic to concurrently executing transactions (that is, transactions do not observe partially compensated states); it conforms to consistency constraints; and its effects are durable. However, a compensating transaction is a very special type of transaction. Under certain circumstances, it is required to *restore* consistency, rather than merely preserve it. It is durable in the strong sense that once a decision is made to initiate compensation, the compensating transaction must complete, since it does not make any sense to abort it. The choice of either to abort or to commit is present for the original transaction. A compensating transaction offers the ability to reverse this choice, but we do not go any further by providing the capability to abort the compensation. There are other special characteristics. Above all, a compensating transaction does not exist by its own right; it is always regarded within the context of the compensated-for transaction. It is always executed after the compensated-for transaction. Its actions are derivative of the actions of the compensated-for and the dependent transactions. In some situations, the actions of a compensating transaction can be extracted automatically from the program of the compensated-for transaction, the current state of the database, and the current state of the log. In other situations, it is the system programmer’s responsibility to pre-define a compensating transaction.

A mundane example taken from “real life” exemplifies some of the characteristics of compensation. Consider a database system that deals with transactions that represent purchasing of goods. Consider the act of a customer returning goods after they have been sold. The compensated-for transaction in that case is a particular purchase, and the compensating transaction encompasses the activity caused by the cancellation of the purchase. The compensating transaction is bound to the compensated-for transaction by the details of the particular sale (e.g., price, method of payment, date of purchase). The effects of purchasing transaction might have been externalized in different ways. For instance, it might have triggered a dependent transaction that issued an order to the supplier in an attempt to replenish the inventory of the sold goods. Furthermore, the customer might have been added to the store’s mailing list as a result of that particular sale. The actual compensation depends on the relevant policy. For example, the customer may be given store credit, or full refund. Whether to cancel the order from the supplier and whether to retain the customer in the mailing list are other application-dependent issues with which the compensating transaction must deal.

### 3 A Transaction Model

In the classical transaction model [14, 2] transactions are viewed as sequences of read and write operations that map consistent database states to consistent states when executed in isolation. The correctness criterion of this model is called *serializability*. A concurrent execution of a set of transactions is represented as an interleaved sequence of read and write operations, and is said to be serializable if it is equivalent to a serial (non-concurrent) execution.

This approach poses severe limitations on the use of compensation. First, sequences of uninterpreted reads and writes are of little use when the semantically-rich activity of compensation is considered. Second, the use of serializability as the correctness criterion for applications that demand interaction and cooperation among possibly long-duration transactions was questioned by the work on concurrency control in [11, 9, 3]. Since we target compensation as a recovery mechanism for these kind of applications, our model does not rely on serializability as the only correctness notion.

#### 3.1 Transactions and Programs

A transaction is a sequence of operations that are generated as a result of the execution of some program. The exact sequence that the program generates depends on the database state “seen” by the program. In the classical transaction model only the sequences are dealt with, whereas the programs are abstracted and are of little use. Given a concurrent execution of a set of transactions (i.e., an interleaved sequence of operations) compensation for one of the transactions,  $T$ , can be modeled as an attempt to cancel the operations of  $T$  while leaving the rest of the sequence intact. The validity of what remains from that execution is now in serious doubt, since originally transactions read data items updated by  $T$  and acted accordingly, whereas now  $T$ 's operations have vanished but its indirect impact on its dependent transactions is still apparent. The only formal way to examine a compensated execution is by comparing it to a hypothetical execution of only the dependent transactions, without the compensated-for transaction. We use the comparison of the compensated execution with the hypothetical execution that does not include the compensated-for transaction, as a key criterion in our exposition. Generating this hypothetical execution and studying it requires the introduction the *transactions' programs* which are, therefore, indispensable for our purposes.

A *transaction program* can be defined in any high-level programming language. Programs have local (i.e., private) variables. In order to support the private (i.e., non-database) state space of programs we define the

concept of an *augmented state*. The augmented state space is the database state space unioned with the private state spaces of the transactions' programs. The provision of an augmented state allows one to treat reading and updating the database state in a similar manner. Reading the database state is translated to an update of the augmented state, thereby modeling the storage of the value read in a local variable.

Thus, a *database*, denoted as  $db$ , is a set of data *entities*. The *augmented database*, denoted as  $adb$ , is a set of entities that is a superset of the database; that is,  $db \subset adb$ . An entity in the set  $(adb - db)$  is called a *private entity*. Entities have identifying *names* and corresponding *values*. A *state* is a mapping of entity names to entity values. We distinguish between the *database state* and the state of the augmented database, which is referred to as the *augmented state*. We use the notation  $S(e)$ , to denote the value of entity  $e$  in a state  $S$ . The symbols  $S$  and  $e$  (and their primed versions,  $S', e'$ , etc.) are used, hereafter, to denote a state and an entity, respectively.

Another deviation from the classical transaction model is the use of semantically-richer operations instead of the primitive read and write. Having such operations allows refining the notion of conflicting versus commutative operations [1, 16]. That is, it is possible to examine whether two operations commute and hence can be executed concurrently. By contrast, in the classical model, there is not much scope for such considerations since a write operation conflicts with any other operation on the same entity.

An *operation* is a function from augmented states to augmented states that is restricted as follows:

- An operation updates at most one entity (either a private or a database entity);
- an operation reads at most one *database* entity, but it may read an arbitrary number of private entities;
- an operation can both update and read only the same database entity.

We use the following shorthand notation for a single operation  $f$ :  $e_0 := f(e_1, \dots, e_k)$ . We say that  $f$  *updates* entity  $e_0$ , and *reads* entities  $e_1, \dots, e_k$ . The *arguments* of an operation are all the entities it reads. There are two special termination operations, *commit*, and *abort*, that have no effect on the augmented state. Operations are assumed to be executed *atomically*.

It is implicitly assumed that all the arguments of an operation are meaningful; that is, a change in their value cause a change in the value computed by the operation. The operations in our model reconcile two contradictory goals. On the one hand, operations are functions from augmented states to augmented states, thereby

giving the flexibility to define complex and semantically-rich operations. On the other hand, the mappings are restricted so that at most one database entity is accessed in the same operation, thereby making it feasible to allow atomic execution of an operation. Although only one database entity may be accessed by an operation, as many local variables (i.e., private entities) as needed may be used as arguments for the mapping associated with the operation. Having private entities as arguments to operations adds more semantics to operations. Having functions for operations allows us to conveniently compose operations by functional composition, thereby making sequences of operations functions too.

We are in a position now to introduce the notion of a transaction as a program. A *transaction program* is a sequence of *program statements*, each of which is either:

- An operation.
- A *conditional statement* of the form: **if  $b$  then  $SS1$  else  $SS2$** , where  $SS1$  and  $SS2$  are sequences of program statements, and  $b$  is a predicate that mentions only private entities and constants.

We impose the the following restrictions on the operations that are specified in the statements:

- The set of private entities is *partitioned* among the transaction programs. An operation in a program cannot read nor update a private entity that is not in its own partition;
- private entities are updated only once;
- An operation reads a private entity only after another operation has updated that entity.

**Example 1.** Consider the following sets of entities:  $db = \{a, b, c\}$ , and  $adb = db \cup \{u, v, w\}$ , and the following two transaction programs,  $T_1$  and  $T_2$ :

```

T1: begin
      u:=a;
      v:=b;
      if u > v then c:= f(c,v)
                else begin
                          w:=c;
                          b:= g(u,w)
                        end
      end

T2:  begin
      a:=0;
      b:=1
      end

```

Observe that operation  $T_1$  both updates and reads entity  $c$ .  $T_2$  demonstrates operations that read no entities.  $\diamond$

### 3.2 Histories and Correctness

We use the framework for alternative correctness criteria set forth in [11]. Explicit *input* and *output predicates* over the database state are associated with transactions. The input predicate is a pre-condition of transaction execution and must hold on the state that the transaction reads. The output condition is a post-condition which the transaction guarantees on the database state at the end of the transaction provided that there is no concurrency and the database state seen by the transaction satisfies the input condition. Thus, as in the standard model, transactions are assumed to be generated by correct programs, and responsibility for correct concurrent execution lies with the concurrency control protocol.

Observe that the input and output predicates are excellent means for capturing the semantics of a database system. We use the convention that predicates (and hence semantics) can be associated with a set of transactions, similarly to the way predicates are associated with nested transactions in [11]. That is, a set of transactions is supposed to collectively establish some desirable property, or complete a coherent task. This convention is most useful in domains where a set of sub-transactions are assigned a single complex task.

We do not elaborate on the generation of interleaved or concurrent executions of sets of transaction programs, since this is not central to understanding our results. However, the notion of a history, the result of this interleaving, is a central concept in our model. A *history* is a sequence of operations, defining both a total order among the operations, as well as a function from augmented states to augmented states that is the functional composition of the operations. We use the notation  $X = \langle f_1, \dots, f_n \rangle$  to denote a history  $X$  in which operation  $f_i$  precedes  $f_{i+1}$ ,  $1 \leq i < n$ . Alternatively, we use the functional composition symbol 'o' to compose operations as functions. That is,  $X = f_1 \circ \dots \circ f_n$  denotes the function from augmented states to augmented states defined by the same history  $X$ . We use the upper case letters at the end of the alphabet, e.g.,  $X, Y, Z$ , to denote both the sequence and the function a history defines.

The equivalence symbol ' $\equiv$ ' is used to denote equality of histories as functions. That is, if  $X$  and  $Y$  are histories, then  $X \equiv Y$  means that for all augmented states  $S$ ,  $X(S) = Y(S)$ . Observe that since histories and operations alike are functions, the function composition symbol 'o' is used to compose histories as well as operations.

When a (concurrent) execution of a set of transaction programs  $A$  is initiated on a state  $S$  and generates a history  $X$ , we say that  $X$  is a *history of  $A$*  whose *initial state* is  $S$ .

**Example 2.** Consider the transaction program  $T_1$  of

Example 1. Since  $T_1$  has a conditional statement there are two histories,  $X$  and  $Y$ , which can be generated when  $T_1$  is executed in isolation. We list the histories as sequences of operations:

$$\begin{aligned} X &= \langle u := a, v := b, c := f(c, v) \rangle, \\ Y &= \langle u := a, v := b, w := c, b := g(u, w) \rangle \end{aligned}$$

Let  $S = \{ a = 1, b = 0, c = 2 \}$  be database state, then  $S$  is an initial state for  $X$ .  $X(S) = S'$ , where  $S'(c) = f(2, 0)$ . Consider a *concurrent* execution of  $T_1$  and  $T_2$  of the previous example. We show two (out of the many possible) histories,  $Z$  and  $W$ , whose initial state is  $S$  given above. Each operation is prefixed with the name of the transaction that issued it.

$$\begin{aligned} Z &= \langle T_2 : a := 0, T_1 : u := a, T_2 : b := 1, \\ &\quad T_1 : v := b, T_1 : w := c, T_1 : b := g(u, w) \rangle \\ W &= \langle T_2 : a := 0, T_2 : b := 1, T_1 : u := a, \\ &\quad T_1 : v := b, T_1 : w := c, T_1 : b := g(u, w) \rangle \end{aligned}$$

Observe that  $Z(S) = W(S) = S''$ , where  $S'' = \{ a = 0, b = g(0, 2), c = 2 \}$ . Observe that  $Z \equiv W$ .  $\diamond$

A key notion in the treatment of compensation is *commutativity*. We say that two sequences of operations,  $X$  and  $Y$ , *commute*, if  $(X \circ Y) \equiv (Y \circ X)$ . Two operations *conflict* if they do not commute. Observe that defining operations as functions, regardless to whether they read or update the database, leads to a very simple definition of the key concept of commutativity.

Part of the orderings implied by the total order in which operations are composed to form a history are arbitrary, since only conflicting operations must be totally ordered. In essence, our equivalence notion (when restricted to database state) is similar to final-state equivalence [14]. However, in what follows, we shall need to equate histories that are not necessarily over the same set of transactions, which is in contrast to final-state equivalence (and actually to all familiar equivalence notions).

A *projection* of a history  $X$  on an entity  $e$  is a subsequence of  $X$ , that consists of the operations in  $X$  that updated  $e$ . We denote the projection of  $X$  on  $e$  as  $X_e$ . The same notation is used for a projection on a set of entities.

We impose very weak constraints on concurrent executions in order to exclude as few executions as possible from consideration. In this paper we consider the following types of histories:

- A history  $X$  is *serial* if for every two transactions  $T_i$  and  $T_j$  that appear in  $X$ , either all operations of  $T_i$  appear before all operations of  $T_j$  or vice versa.

- A history  $X$  is *serializable* (SR) if there exists a serial history  $Y$  such that  $X \equiv Y$ .
- Let  $C = c_1 \wedge \dots \wedge c_n$  be a predicate over the database state. For each conjunct  $c_i$  let  $d_i$  denote the set of database entities mentioned in  $c_i$ . A history  $X$  is *predicate-wise serializable* with respect to a predicate  $C$  ( $\text{PWSR}_C$ ) if for every set of entities  $d_i$  there exists a serial history  $Y$  such that  $X_{d_i} \equiv Y_{d_i}$ .
- A history  $X$  is *entity-wise serializable* (EWSR) if for every entity  $e$  there exists a serial history  $Y$  such that  $X_e \equiv Y_e$ .

The definition of PWSR histories is adapted from [9]. As we shall see shortly, EWSR histories are going to be quite useful in our work. The following lemma is given without proof.

**Lemma 1.** *Let  $C$  be a predicate that mentions all database entities, and let  $\text{ewsr}$ ,  $\text{pwsr}_C$ ,  $\text{sr}$  denote the set of EWSR histories, PWSR $_C$  histories, and SR histories, respectively. Then,  $\text{sr} \subset \text{pwsr}_C \subset \text{ewsr}$ .  $\square$*

We denote by  $X_T$  the sequence of operations of a transaction  $T$  in a history  $X$ , involving possibly other transactions. The same notation is used for sets of transactions. When  $X_T$  is projected on entity  $e$  the resulting sequence is denoted  $X_{T,e}$ .

## 4 Compensating Transactions

With the aid of the tools developed in the last section, we are in a position to define compensation more formally.

### 4.1 Specification Constraints

Although compensation is an application-dependent activity, there are certain guidelines to which every compensating transaction must adhere. After introducing some notation and conventions we present three specification constraints for defining compensating transactions. These constraints provide a very broad framework for defining concrete compensating transactions for concrete applications, and can be thought of as a generic specification for all compensating transactions.

We say that transaction  $T_j$  is *dependent upon* transaction  $T_i$  in a history if there exists an entity  $e$  such that

- $T_j$  reads  $e$  after  $T_i$  has updated  $e$ ;
- $T_i$  does not abort before  $T_j$  reads  $e$ ; and
- every transaction (if any) that updates  $e$  between the time  $T_i$  updates  $e$  and  $T_j$  reads  $e$ , is aborted before  $T_j$  reads  $e$ .

The above definition is adapted from [2].

A transaction  $T_i$ , which is the depended-upon transaction may be either a committed transaction, or an active transaction. In either case, if we want to support the undo of  $T_i$ , then the corresponding compensating transaction,  $CT_i$ , must be pre-defined. The key point is that admitting non-recoverable histories and supporting the undo of committed transactions is predicated on the existence of the compensatory mechanisms needed to handle undoing externalized transactions. In the rest of the paper,  $T$  denotes a compensated-for transaction,  $CT$  denotes the corresponding compensating transaction, and  $dep(T)$  denotes a set of transactions dependent upon  $T$ . This set of dependent transactions can be regarded as a set of related (sub)transactions that perform some coherent task.

**Constraint 1.** *For all histories  $X$ , if  $X_{T,e} \circ X_{CT,e}$  is a contiguous subsequence of  $X_e$ , then  $(X_{T,e} \circ X_{CT,e}) \equiv I$ , where  $I$  is the identity mapping.*  $\square$

The simplest interpretation of Constraint 1 is that for all entities  $e$  that were updated by  $T$  but read by no other transaction (since  $X_{CT,e}$  follows  $X_{T,e}$  in the history),  $CT$  amounts simply to undoing  $T$ . Consequently, if there are no transactions that depend on  $T$ , (i.e., no transaction reads  $T$ 's updated data entities), then  $CT$  is just the traditional  $undo(T)$ . The fact that  $CT$  does not always just undo  $T$  is crucial, since the effects of compensation depend on the span of history from the execution of the compensated-for transaction till its own initiation. If such a span exists, and  $T$  has dependent transactions, the effects of compensation may vary and can be very different from undoing  $T$ . For instance, compensation may include additional activity that is not directly related to undoing. A good example here is a cancellation of reservation in an airline reservation system which is handled as a compensating transaction that causes the transfer of pending reservation from a waiting list to the confirmed list.

There are certain operations on certain entities that cannot be undone, or even compensated-for, in the form of inverting the state. In [6] these type of operations and entities are termed *real* (e.g., dispensing money, firing a missile). For simplicity's sake, we omit discussion of such entities.

**Constraint 2.** *Given a history  $X$  involving  $T$  and  $CT$ , there must exist  $X'$  and  $X''$  subsequences of  $X$ , such that no transaction has operations both in  $X'$  and in  $X''$ , and  $X \equiv X' \circ X_{CT} \circ X''$ .*  $\square$

This constraint represents the atomicity of compensation. That is, a transaction should either see a database state affected by  $T$  (and not by  $CT$ ), or see a state following  $CT$ 's termination. More precisely, transactions

should not have operations that conflict with  $CT$ 's operations scheduled both before and after  $CT$ 's operations, or in between  $CT$ 's first and last operations. It is the responsibility of the concurrency control protocol to implement this constraint (see Section 7 for implementation discussion).

In what follows, we use the notation  $O_T$  and  $I_T$  to denote the output and input predicate of transaction  $T$ , respectively. The same notation is used for a set of transactions. These predicates are predicates over the database state.

**Constraint 3.** *Let  $Q$  be a predicate defined over the database state, if  $(O_{dep(T)} \Rightarrow Q) \wedge (I_T \Rightarrow Q)$  then  $O_{CT} \Rightarrow Q$ .*  $\square$

Constraint 3 is appropriate when  $Q$  is either a general consistency constraint, or a specific predicate that is established by  $dep(T)$  (that is, one of the collective tasks of the transactions in  $dep(T)$  was to make  $Q$  true). Informally, this constraint says that if  $Q$  was established by  $dep(T)$ , and is not violated by undoing  $T$ , then it should be preserved by  $CT$ . Observe that the assumption that  $Q$  holds initially (i.e.,  $I_T \Rightarrow Q$ ) is crucial since  $T$ 's effects are undone by  $CT$ , and hence, predicates established by  $T$  and preserved by  $dep(T)$  do not persist after the compensation. It is the responsibility of whoever defines  $CT$  to enforce Constraint 3.

Constraints 1 and 2 will be assumed to hold for all compensating transactions, hereafter. Constraint 3, which is more intricate and captures more of the semantics of compensation, will be discussed further in Section 6.

## 4.2 Types of Compensation

For some applications, it is acceptable that an execution of the dependent transaction, without the compensated-for and the compensating transactions, would produce different results than those produced by the execution with the compensation. On the other hand, other applications might forbid compensation unless the outcome of these two executions is the same. Next we make explicit the above criterion that distinguishes among types of compensation by defining the notion of compensation soundness.

**Definition 1.** *Let  $X$  be the history of  $T$ ,  $CT$ , and  $dep(T)$  whose initial state is  $S$ . Let  $Y$  be some history of only the transactions in  $dep(T)$  whose initial state is also  $S$ . The history  $X$  is sound, if  $X(S) = Y(S)$ .*  $\square$

The history  $Y$  can be any history of  $dep(T)$ . As far as the definition goes, different sets of (sub)transactions of  $dep(T)$  may commit in  $X$  and in  $Y$ , and conflicting operations may be ordered differently. The key point is

that  $X(S) = Y(S)$ . If a history is sound then compensation does not disturb the outcome of the dependent transactions. The database state after compensation is the same as the state after an execution of only the dependent transactions in  $dep(T)$ . All direct and indirect effects of the compensated-for transaction,  $T$ , have been erased by the compensation.

Transactions in  $dep(T)$  see different database states when  $T$  and  $CT$  are not executed, and therefore generate a history  $Y$  which can be totally different than the history  $X$ . This distinction between the histories  $X$  and  $Y$ , which is the essence of the important notion of soundness, would not have been possible had we viewed a transaction merely as sequence of operations rather than a program.

A delicate point arises with regard to soundness when  $S$  does not satisfy  $I_{dep(T)}$ . Such situations may occur when  $T$  establishes  $I_{dep(T)}$  for  $dep(T)$  in such a manner that  $dep(T)$  must follow  $T$  in any history. Hence, if  $T$  is compensated-for, there is no history of  $dep(T)$ ,  $Y$ , that can satisfy the soundness requirement. We model such situations by postulating that if  $I_{dep(T)}(S)$  does not hold, then  $Y(S)$  results in a special state that is not equal to any other state (the *undefined* state), and hence  $X$  is indeed not sound.

We illustrate Definition 1 by considering the following two histories over read and write operations (the notation  $r_i[e]$  denotes reading  $e$  by  $T_i$ , and similarly  $w_i[e]$  for write, and  $c_i$  for commit):

$$\begin{aligned} W &= \langle w_j[e], r_i[e], c_j, c_i \rangle \\ Z &= \langle w_j[e], r_i[e], w_i[e'], c_i \rangle \end{aligned}$$

The history  $W$  is recoverable. History  $Z$  is not recoverable. If however,  $CT_j$  is defined,  $T_j$  can still be aborted. Let us extend  $Z$  with the operations of  $CT_j$  and call the extended history  $Z'$ .  $Z'$  is sound provided that  $Z'_i$  would have been generated by  $T_i$ 's program, and the same value would have been written to  $e'$ , had  $T_i$  run in isolation starting with the same initial state as in  $Z'$ .

The key notion in the context of compensation, as we defined it, is *commutativity* of compensating operations with operations of dependent transactions. Significant attention has been devoted to the effects of commutative operations on concurrency control [8, 16, 1]. Our work parallels these results as it exploits commutativity with respect to recovery. In all of our theorems we prefer to impose commutativity requirements on  $CT$  rather than on  $T$ , since  $CT$  is less exposed to users, and hence constraining it, rather than constraining  $T$ , is preferable. Predicated on commutativity, the operations of the compensated-for transaction and the corresponding compensatory operations can be 'brought together', and then cancel each other's effects (by the enforcement of Constraint 1), thereby ensuring sound histories. The

following theorem formalizes this idea.

**Theorem 1.** *Let  $X$  be a history involving  $T, dep(T)$  and  $CT$ . If each of the operations in  $X_{dep(T)}$  commutes with each of the operations in  $X_{CT}$ , then  $X$  is sound.*  $\square$

We illustrate this theorem by the following simple example:

**Example 3.** Let  $T_i, T_j$  and  $CT_i$  be a compensated-for transaction, a dependent transaction and the compensating transaction, respectively. Let the programs of all these transactions include no condition statements (i.e., they are sequences of operations). We give a history  $X$ , in which each operation is prefixed by the name of the issuing transaction.  $X = \langle T_i : a := a + 2, T_j : u := b, T_j : a := a + u, CT_i : a := a - 2 \rangle$ . Clearly, every operation of  $T_j$  commutes with every operation of  $CT_i$  in  $X$ . Hence,  $X$  is sound, and the history that demonstrates soundness is simply  $Y = X_{T_j} = \langle T_j : u := b, T_j : a := a + u \rangle$ . As will become clear in Section 6, the fact that no condition statements appear in  $T_j$  is important.  $\diamond$

Our main emphasis in this paper is on more liberal forms of compensation soundness, where the results of executing the dependent transactions in isolation may be different from their results in the presence of the compensated-for, and the compensating transactions. One way of characterizing these weaker forms of soundness is by qualifying the set of entities for which the equality in Definition 1 holds. In Section 5.1, we define a type of compensating transaction that ensures sound compensation with respect to some set of entities. Alternatively, in Section 6 we investigate other weak forms of soundness that approximate (pure) soundness.

## 5 Examples and Applications

In this section, we present several examples to illustrate the various concept we have introduced so far. Throughout this section we use the symbols  $T, dep(T), CT, X,$  and  $S$  to denote a compensated-for transaction, its compensating transaction, the corresponding set of dependent transactions, the history of all these transactions, and the history's initial state, respectively.

### 5.1 A Generic Example

In this example we present a generic compensation definition. Let  $update(T, X)$  denote the set of database entities that were updated by  $T$  in history  $X$ . The same notation is used for a set of transactions.

**Definition 2.** *Let  $X(S) = S'$ , and  $X \equiv X' \circ X_{CT}$  (by Constraint 2). We define the generic compensating*

transaction  $CT$ , by characterizing  $S'$  for all entities  $e$ :

$$S'(e) = \begin{cases} S(e) & \text{if } e \notin \text{update}(\text{dep}(T), X) \\ (X'(S))(e) & \text{if } e \in \text{update}(\text{dep}(T), X) \\ \wedge e \notin \text{update}(T, X) \\ X_{\text{dep}(T), e}(S) & \text{if } e \in \text{update}(\text{dep}(T), X) \\ \wedge e \in \text{update}(T, X) \end{cases}$$

□

Before we proceed, we informally explain the meaning of this type of compensation. If no dependent transaction updates an entity that  $T$  updates,  $CT$  undoes  $T$ 's updates on that entity. The value of entities that were updated only by dependent transactions is left intact. The value of entities updated by both  $T$  and its dependents should reflect only the dependents' updates.

There is a certain subtlety in the second case of the definition which is illustrated next. Assume that  $T$  updated  $e$ . The modified  $e$  is read by a transaction in  $\text{dep}(T)$  and the value read determines how this transaction updates  $e'$ . After compensation, even though the initial value of  $e$  is restored (by the first case of the definition), the indirect effect it had on  $e'$  is left intact (by the second case of the definition). We use the above definition as a precise specification of what  $CT$  should accomplish.

To further illustrate the type of compensation just described, we give a concrete example. Consider an airline reservation system with the entity *seats* that denotes the total number of seats in a particular flight, entity *rs* that denotes the number of already reserved seats in that flight, and entity *reject* that counts the number of transactions whose reservations for that flight have been rejected. Let  $\text{reserve}(x)$  be a simplified seat reservation transaction for  $x$  seats defined as:

```
if (rs + x) <= seats then rs:=rs+x
                    else reject:= reject+1
```

The consistency constraint  $Q$  in this case is:  $Q(S)$  iff  $S(rs) \leq S(\text{seats})$ . Assume:

$$S = \{\text{seats} = 100, \text{rs} = 95, \text{rejects} = 10\},$$

$$T = \text{reserve}(5), \text{dep}(T) = \{\text{reserve}(3)\}$$

Let the history be  $X \equiv X_T \circ X_{\text{dep}(T)} \circ X_{CT}$  where  $CT$  is defined by Definition 2. We would like to have after  $X$ :  $S' = \{\text{rs} = 95, \text{rejects} = 11\}$ , that is,  $T$ 's reservations were made and later canceled by running  $CT$ , and  $\text{dep}(T)$ 's reservations were rejected. And that is exactly what we get by our definition. Observe how  $T$ 's reservations were canceled, but still its indirect impact on *rejects* persists (since  $T$  caused  $\text{dep}(T)$ 's reservations to be rejected).

Hence, this example demonstrates a history that is not sound but is nevertheless intuitively acceptable.

Had the transaction in  $\text{dep}(T)$  been executed alone, it would result in successful reservations. Notice how in this example the operation of  $CT$  can be implemented as inverse of  $T$ 's operation (addition and subtraction). The less interesting case, where there are enough seats to accommodate both  $T$  and  $\text{dep}(T)$ , also fits nicely. In this case  $CT$ 's subtraction on the entity *seats* commutes with  $\text{dep}(T)$ 's addition to this entity.

## 5.2 Storage Management Examples

The following example is from [13], though the notion of compensation is not used there. Consider transactions  $T_1$  and  $T_2$ , each of which adds a new tuple to a relation in a relational database. Assume the tuples added have different keys. A tuple addition is processed by first allocating and filling in a slot in the relation's tuple file, and then adding the key and slot number to a separate index. Assume that  $T_i$ 's slot updating ( $S_i$ ) and index insertion ( $I_i$ ) steps can each be implemented by a single page read followed by a single page write (written  $r_i[tp]$ ,  $w_i[tp]$  for a tuple file page  $p$ , and  $r_i[ip]$ ,  $w_i[ip]$  for an index file page  $p$ ).

Consider the following history of  $T_1$  and  $T_2$  regarding the tuple pages  $tq, tr$  and the index page  $ip$ :

$$\langle r_1[tq], w_1[tq], r_2[tr], w_2[tr], \\ r_2[ip], w_2[ip], r_1[ip], w_1[ip] \rangle$$

This is a serial execution of  $\langle S_1, S_2, I_2, I_1 \rangle$ , which is equivalent to the serial history of executing  $T_1$  and then  $T_2$ . Assume, now, that we want to abort  $T_2$ . The index insertion  $I_1$  has seen and used page  $p$ , which was written by  $T_2$  in its index insertion step. The only way to abort  $T_2$ , without aborting  $T_1$  is to compensate for  $T_2$ . Fortunately, we have a very natural compensation,  $CT_2$ , which is a delete key operation. Observe that a delete operation as compensation, satisfies Constraint 1, commutes with insertion of a tuple with a different key, and encapsulates composite compensation for the slot updating and index insertion. The resulting history is sound.

## 6 Approximating Soundness

In this section we introduce weak forms of compensation soundness, where the results of an execution that includes compensation only *approximate* the results of executing the dependent transactions in isolation.

Let us denote the history of transactions  $T$ ,  $\text{dep}(T)$  and  $CT$  as  $X$ , and the history without compensation, i.e., a history of only  $\text{dep}(T)$ , as  $Y$ . In an approximated form of soundness, the final state of  $X$  is only *related* to the final state of  $Y$ .



The relation should serve to constrain  $CT$ , and prevent it from violating consistency constraints and other desirable predicates established by  $dep(T)$ . Thus, the relation should enforce some ‘goodness’ properties, for instance: “if a consistency constraint predicate holds on the final state of  $Y$ , it should also hold on the final state of  $X$ .”

Achieving even approximated soundness is an intricate problem when the histories are non-serializable, as we allow them to be. The obstacle is, as mentioned before, that the programs of transactions in  $dep(T)$  see different database states when  $T$  and  $CT$  are not executed, and therefore may generate a history  $Y$  which can be totally different than the original history  $X$ . Hence,  $X$  and  $Y$  may not be related as required.

We state several theorems that formalize the interplay among the approximated soundness notion, concurrency control constraints, restrictions on programs of dependent transactions, and commutativity. Each theorem is followed by a simplified example that serves to illustrate at least part of the theorem’s premises and consequences. Proofs of the theorems can be found in [10]. Throughout this section, we assume that a compensating transaction complies with Constraints 1 and 2 of Section 4. We start with definitions of weaker forms of commutativity and weaker forms of compensation soundness.

**Definition 3.** *Two sequences of operations,  $X$  and  $Y$ , commute with respect to a relation  $\mathcal{R}$  on augmented states (in short, R-commute), if for all augmented states  $S$ ,  $(X \circ Y)(S) \mathcal{R} (Y \circ X)(S)$ .  $\square$*

Observe that when  $\mathcal{R}$  is the equality relation we have regular commutativity.

**Definition 4.** *Let  $X$  be a history of  $T$ ,  $dep(T)$ , and  $CT$  whose initial state is  $S$ , and let  $\mathcal{R}$  be a reflexive relation on augmented states. The history  $X$  is sound with respect to  $R$  (in short R-sound), if there exists a history  $Y$  of  $dep(T)$  whose initial state is  $S$  such that  $Y(S) \mathcal{R} X(S)$ .  $\square$*

Observe that regular soundness is a special case of R-soundness when  $\mathcal{R}$  is the equality relation. Since  $\mathcal{R}$  is reflexive, the empty history is always R-sound, regardless of the choice of  $R$ .

We motivate the above definitions by considering adequate relations  $\mathcal{R}$  in the context of R-commutativity and R-soundness. Let  $Q$  be a predicate on database states such that  $O_{dep(T)} \Rightarrow Q$ .  $Q$  can be regarded as either a consistency constraint, or a desired predicate that is established by  $dep(T)$  (similarly to the predicate  $Q$  in Constraint 3). Therefore, we would like to guarantee that compensation does not violate  $Q$ . Define  $\mathcal{R}$

(in the context of  $X, Y$  and  $S$ ) as follows:

$$Y(S) \mathcal{R} X(S) \text{ iff } (Q(Y(S)) \Rightarrow Q(X(S)))$$

An R-sound history with such  $\mathcal{R}$  has the advantageous property that predicates like  $Q$  are not violated by the compensation. Such R-sound histories yield states that approximate states yielded by sound histories in the sense that both states satisfy some desirable predicates. In the examples that follow the theorems, we use relations  $\mathcal{R}$  of that form.

**Definition 5.** *Let  $\mathcal{R}$  be a relation on states, and let  $v_e$  and  $v'_e$  denote values of an arbitrary entity  $e$ . We define the relations  $R_e$  on values of  $e$  for every entity  $e$  as follows:*

$$v_e \mathcal{R}_e v'_e \text{ iff } (\exists S', S'' : S'(e) = v_1 \wedge S''(e) = v_2 \wedge S' \mathcal{R} S'') \quad \square$$

**Definition 6.** *Let  $X$  be a history of  $T$ ,  $dep(T)$  and  $CT$  whose initial state is  $S$ , and let  $\mathcal{R}$  be a reflexive relation on augmented states. The history  $X$  is partially R-sound if there exists a history  $Y$  of  $dep(T)$  whose initial state is  $S$  such that  $(\forall e \in db : (Y(S))(e) \mathcal{R}_e (X(S))(e))$ .  $\square$*

**Definition 7.** *A program of a transaction is fixed if it is a sequence of operations that use no private entities as arguments.  $\square$*

If  $T$ ’s program is fixed then it has no conditional branches. Moreover,  $T$  cannot use local variables to store values for subsequent referencing. A sequence of operations, where each operation reads and updates a single database entity (without storing values in local variables) is a fixed transaction. A transaction that uses a single operation to give a raise to a certain employee recorded in a salary management database is an example for a fixed transaction.

**Theorem 2.** *Let  $X$  be a history of  $T, dep(T)$  and  $CT$  whose initial state is  $S$ . If the histories  $X_{dep(T)}$  and  $X_{CT}$  R-commute,  $X$  is EWSR, and all programs of transactions in  $dep(T)$  are fixed, then  $X$  is partially R-sound.  $\square$*

**Example 4.** Consider a database system with the following entities, parametric operations, and reflexive relation:

$$\begin{aligned} db &= \{a : integer, b : integer\}, \\ f(e) &:: \text{if } e > 2 \text{ then } e := e - 2, \\ g(e) &:: \text{if } e > 10 \text{ then } e := e - 10 \\ S' \mathcal{R} S'' &\text{ iff } (((S'(b) \geq 0 \wedge S'(a) \geq 10) \vee (S'(a) = 4)) \\ &\Rightarrow ((S''(b) \geq 0 \wedge S''(a) \geq 10) \vee (S''(a) = 4))) \end{aligned}$$

(The predicates on  $a$  are present only to demonstrate partial R-soundness). We emphasize that  $f$  and  $g$  are

(atomic) operations. The history  $X$  is as follows (there is no need to give the program of  $dep(T)$  since it is fixed):

$$X = \langle dep(T) : A := a + 2, T : f(a), T : g(b), \\ dep(T) : g(b), CT : a := a + 2, CT : b := b + 10 \rangle$$

Observe that  $X_{dep(T)}$  and  $X_{CT}$  do not commute but they do R-commute for the given relation  $\mathcal{R}$ . Let the initial state be  $S = \{a = 2, b = 15\}$ . We have that  $X(S) = \{a = 4, b = 15\}$ , whereas  $Y(S) = \{a = 4, b = 5\}$ , and indeed  $X$  is *partially* R-sound.  $\diamond$

The inherent problem with (the proofs of) compensation soundness is the fact that they equate two histories that are *not* over the same set of transactions, which is in contrast to all the equivalence notions in the traditional theory of concurrency control. The obstacle is that the history  $Y$  may be generated by different executions of the programs of  $dep(T)$ , and may be totally different from  $X_{dep(T)}$ , which is just a syntactic derivative of the history  $X$ . In Theorem 2, this problem was solved only because  $dep(T)$  was fixed. This obstacle can be removed by posing more assumptions, as is done next.

**Definition 8.** A transaction  $T$  is a serialization point in a history  $X$  if  $X \equiv X' \circ X_T \circ X''$ .  $\square$

Observe that no restrictions are imposed on  $X'$  and  $X''$ . Also notice that a compensating transaction is a serialization point, as implied by Constraint 2.

**Theorem 3.** Let  $X$  be a history of  $T, dep(T)$  and  $CT$ . Let  $Z$  be a history of the transactions in  $dep(T)$  and  $CT$  such that  $Z \equiv Z_{dep(T)} \circ Z_{CT}$ . If for all states  $S$  and for all histories  $Z$ , there exists a history  $Y$  of  $dep(T)$  such that  $(Z_{CT} \circ Y)(S) \mathcal{R} (Z_{dep(T)} \circ Z_{CT})(S)$ , then every history  $X$  where  $T$  is a serialization point is R-sound.  $\square$

Note that it is required that  $dep(T)$ 's programs be such that executing  $CT$  before  $dep(T)$  would result in a state that is related by  $\mathcal{R}$  to the state resulting when executing  $dep(T)$  first and then  $CT$ . Observe that this requirement is stronger than R-commutativity.

This theorem is quite useful since it specifies a concurrency control policy that guarantees R-soundness. Namely, we need to ensure that every potential compensated-for transaction be isolated (i.e.,  $T$  is a serialization point) in order to guarantee R-soundness in case of compensation.

**Example 5.** Consider the set entities of Example 4, with the addition of a private entity  $u$  that belongs to some transaction in  $dep(T)$ . Let the programs of  $T, dep(T), CT$ , and the relation  $\mathcal{R}$  be defined as fol-

lows:

$$T = a := a + 1, CT = a := a - 1, \\ dep(T) = \{u := a; \text{ if } u \geq 5 \text{ then } f(b) \text{ else } g(b)\} \\ S' \mathcal{R} S'' \text{ iff } (S'(b) \geq 0 \Rightarrow S''(b) \geq 0)$$

Even though  $dep(T)$ 's history can branch differently when run alone and in the presence of  $T$  and  $CT$ , the two different histories produce final states that are related by  $\mathcal{R}$ .  $\diamond$

**Definition 9.** A program of a transaction is linear if it is a sequence of operations.  $\square$

Programs are sequences, but we allow operations to read multiple entities, that is, use local variables. Therefore, programs may not be fixed. An example for a linear transaction program is a program that gives a raise to all employees, where the raise based on some aggregated computation (for instance 10% of the minimum salary).

**Definition 10.** Let  $\mathcal{R}$  be a reflexive relation on augmented states. An operation  $f$  that updates  $e$  preserves  $\mathcal{R}$ , if  $(\forall e' \in adb : (S(e') \mathcal{R}_{e'} S'(e')) \Rightarrow (f(S) \mathcal{R}_e f(S')))$   $\square$

**Theorem 4.** Let  $X$  be a history of  $T, dep(T)$  and  $CT$  whose initial state is  $S$ . If the histories  $X_{dep(T)}$  and  $X_{CT}$  R-commute,  $X$  is EWSR, the programs of all transactions in  $dep(T)$  are linear,  $\mathcal{R}$  is transitive, and the operations of  $dep(T)$  preserve  $\mathcal{R}$ , then  $X$  is *partially* R-sound.  $\square$

**Example 6.** Consider the set entities of Example 4, with the addition of a private entity  $u$  that belongs to some transaction in  $dep(T)$ . We use the relation  $S' \mathcal{R} S''$  iff  $((S'(b) \geq S'(a)) \Rightarrow (S''(b) \geq S''(a)))$ . The history  $X$  is as follows:

$$X = \langle T : a := a + 1, dep(T) : u := a, \\ dep(T) : b := u + 10, CT : a := a - 1 \rangle$$

Observe that  $X_{CT}$  and  $X_{dep(T)}$  R-commute (but do not commute),  $dep(T)$  is linear (but not fixed), and  $X$  is (partially) R-sound.  $\diamond$

Finally, based on Lemma 1 from Section 2, we derive the following corollary.

**Corollary 1.** Theorems 2 and 4 hold when  $X$  is PWSR<sub>C</sub> or SR instead of EWSR.  $\square$

The requirements from the dependent transactions in Theorems 2, 3, and 4 are quite severe. Besides the R-commutativity requirement imposed on the operations of the dependent transactions, there are restrictions on the shape of the programs (e.g., fixed or linear programs) in each of the theorems' premises. Clearly, in practical systems, there are many transactions that do

not stand up to any of these criteria. The practical ramification of this observation is that externalization of uncommitted data items should be done in a controlled manner if a degree of soundness is of importance. That is, uncommitted data should be externalized only to transactions that do satisfy the requirements specified in the premises of the theorems. In the context of locks, locks should be released only to qualified transactions, that is, those transactions that do satisfy the requirements. Other transactions must be delayed and are subject to the standard concurrency control and recovery policies.

## 7 Implementation Issues

In this section we discuss several implementation issues that need to be considered in order for compensation to be of practical use. We envision that a compensating transaction would be driven by a scan of the log starting from the first record of the compensated-for transaction and up to its own begin-transaction log record. It is important to provide convenient on-line access to the log information for these purposes. Without a suitable logging architecture, these accesses might translate to I/O traffic that would interrupt the sequential log I/O that is performed on behalf of executing transactions. In addition, log records should contain enough semantic information in order to guide the execution of the compensating transaction. Therefore, it is likely that some form of operation logging will be used [7].

There are some subtle ramifications on concurrency control which are discussed next in the context of locking. We have required that  $CT$ 's execution is serializable with respect to other concurrent transactions. (Constraint 2). Also, if it is reasonably assumed that  $update(CT, X) \subseteq update(T, X)$ , then this leads to the conclusion that the compensating transaction and the dependent transactions should follow a 2-Phase Locking protocol [2] with respect to entities in  $update(T, X)$ . Otherwise, it is possible to violate Constraint 2. A viable strategy that might simplify matters for the implementation can be as follows. Once  $CT$  is invoked, the entities in  $update(T, X)$  should be identified by analyzing the log and then  $CT$  should exclusively lock all entities in this set. After performing the necessary updates,  $CT$  can release these locks.

The recovery issues of compensating transactions themselves must be also considered. As was noted earlier, we should disallow a compensating transaction to be aborted either externally (by user, or an application), or internally (e.g., as a deadlock resolution victim). Still, there is the problem of system failures. We think that the preferred way to handle this problem is to resume uncompleted compensating transactions rather

than undoing them. To accomplish this, we need to resume a compensating transaction from a point where its internal state was saved along with the necessary concurrency control information. We emphasize that the principle for recovery of compensating transaction is that once a begin-transaction record of  $CT$  appears in stable storage,  $CT$  must be completed. An implementation along the lines of the ARIES system [12] can support the persistence of compensating transactions across system crashes. In ARIES, undo activity is logged using Compensating Log Records (CLRs). It is guaranteed that actions are not undone more than once, and that undo actions are not undone even if the undo of a transaction is interrupted by a system crash.

## 8 Related Work

The idea of compensating transactions as a semantically-rich recovery mechanism is mentioned, or at least referred to, in several papers. However, to the best of our knowledge, a formal and comprehensive treatment of the issue and its ramifications is lacking.

Strong motivation for our work can be found in Gray's early paper [6]. The notion of compensation (countersteps) is mentioned in the context of histories that preserve consistency without being serializable in [4, 3].

Compensating transactions are also mentioned in the context of a *saga*, a long-duration transaction that can be broken into a collection of subtransactions that can be interleaved in any way with other transactions [5]. A saga must execute all its subtransactions, hence compensating transactions are used to amend partial execution of sagas. In [5] and in [4] the idea that a compensating transaction cannot voluntarily abort itself is introduced.

A noteworthy approach, which can be classified as a simple type of compensation, is employed in the XPRS system [15]. There, a notion of *failure commutativity* is defined for complete transactions. Two transactions failure commute if they commute, and if they can both succeed then a unilateral abort by either transaction cannot cause the other to abort. Transactions that are classified as failure commutative can run concurrently without any conflicts. Handling the abort of such a transaction is done by a log-based special undo function, which is a special case of compensation as we define it.

In [1], semantics of operations on abstract data types are used to define *recoverability*, which is a weaker notion than commutativity. Conflict relations are based on recoverability rather than commutativity. Consequently, concurrency is enhanced since the potential for conflicts is reduced. When an operation is recoverable with respect to an uncommitted operation, the former operation can be executed; however a commit depen-

dependency is forced between the two operations. This dependency affects the order in which the operations should commit, if they both commit. If either operation aborts, the other can still commit, thereby avoiding cascading aborts. This work is more conservative than ours in the sense that it narrows the domain of interest to serializable histories.

## 9 Conclusions

In this paper, we have argued that exposing uncommitted data is very useful for many database applications employing long-duration, nested and/or collaborative transactions. Compensating transactions are proposed as the means for recovery management in the presence of early externalization. Several types of compensation soundness criteria were introduced and were found to be predicated on notions of commutativity. Even the approximated forms of soundness can be used to guarantee that compensation results in desirable consequences and does not abrogate dependent transactions' outcome. A semantically-rich model that is adequate for dealing with non-serializable and non-recoverable histories was set up, and was offered as a viable tool for the understanding of these intricate histories and compensation issues.

We believe that future database applications will require the rethinking of the traditional transaction model that is founded on serializability and permanence of commitment. Contemporary applications in the domains of CAD and CASE exemplify our belief. The work presented in this paper is a step towards the establishment of this new model.

## References

- [1] B. R. Badrinath and K. Ramamirham. Semantic-based concurrency control: Beyond commutativity. In *Proceedings of the Third International Conference on Data Engineering, Los Angeles, 1987*.
- [2] P. A. Bernstein, V. Hadzilacos, and N. Goodman. *Concurrency Control and Recovery in Database Systems*. Addison-Wesley, Reading, MA, 1987.
- [3] A. A. Farrag and M. T. Ozso. Using semantic knowledge of transactions to increase concurrency. *ACM Transactions on Database Systems*, 14(4):503-525, December 1989.
- [4] H. Garcia-Molina. Using semantic knowledge for transaction processing in a distributed database. *ACM Transactions on Database Systems*, 8(2):186-213, June 1983.
- [5] H. Garcia-Molina and K. Salem. Sagas. In *Proceedings of ACM-SIGMOD 1987 International Conference on Management of Data, San Francisco, pages 249-259, 1987*.
- [6] J. N. Gray. The transaction concept: Virtues and limitations. In *Proceedings of the Seventh International Conference on Very Large Databases, Cannes, pages 144-154, 1981*.
- [7] T. Haerder and A. Reuter. Principles of transaction oriented database recovery — a taxonomy. *ACM Computing Surveys*, 15(4):289-317, December 1983.
- [8] H. F. Korth. Locking primitives in a database system. *Journal of the ACM*, 30(1):55-79, January 1983.
- [9] H. F. Korth, W. Kim, and F. Bancillon. On long duration CAD transactions. *Information Sciences*, 46:73-107, October 1988.
- [10] H. F. Korth, E. Levy, and A. Silberschatz. A formal approach to recovery by compensating transactions. Technical Report TR-90-14, The University of Texas at Austin, Computer Sciences Department, 1990.
- [11] H. F. Korth and G. Speegle. Formal model of correctness without serializability. In *Proceedings of ACM-SIGMOD 1988 International Conference on Management of Data, Chicago, pages 379-388, June 1988*.
- [12] C. Mohan, D. Haderle, B. Lindsay, H. Pirahesh, and P. Schwarz. ARIES: A transaction recovery method supporting fine-granularity locking and partial rollbacks using write-ahead logging. Technical Report RJ 6649 (63960), IBM Research, 1989. To appear in *ACM Transactions on Database Systems*.
- [13] J. E. B. Moss, N. D. Griffeth, and M. H. Graham. Abstractions in recovery management. In *Proceedings of ACM-SIGMOD 1986 International Conference on Management of Data, Washington, pages 72-83, 1986*.
- [14] C. Papadimitriou. *The Theory of Database Concurrency Control*. Computer Science Press, Rockville, Maryland, 1986.
- [15] M. R. Stonebraker, R. H. Katz, D. A. Patterson, and J. K. Ousterhout. The design of XPRS. In *Proceedings of the Fourteenth International Conference on Very Large Databases, Los Angeles, pages 318-330, 1988*.
- [16] W. E. Weihl. Commutativity-based concurrency control for abstract data types. *IEEE Transactions on Computers*, C-37(12):1488-1505, December 1988.