

**In Search of Acceptability Criteria:
Database Consistency Requirements and
Transaction Correctness Properties**

K. Ramamritham and P.K. Chrysanthis

COINS Technical Report 92-54

Replaces COINS TR 91-92

July 1992

In Search of Acceptability Criteria: Database Consistency Requirements and Transaction Correctness Properties

Krithi Ramamritham
Dept. of Computer Science
University of Massachusetts
Amherst, MA 01003

Panos K. Chrysanthis
Dept. of Computer Science
University of Pittsburgh
Pittsburgh, PA 15260

Abstract

Whereas serializability captures *database consistency requirements* and *transaction correctness properties* via a single notion, recent research has attempted to come up with correctness criteria that view these two types of requirements independently. In this paper, we develop a taxonomy of various *correctness criteria* that focus on database consistency requirements and transaction correctness properties from the viewpoint of *what* the different dimensions of these two are. This taxonomy allows us to categorize correctness criteria that have been proposed in the literature. To help in this categorization, we have applied a uniform specification technique, based on ACTA, to express the various criteria. Such a categorization helps shed light on the similarities and differences between different criteria and to place them in perspective.

1 INTRODUCTION

Database consistency requirements capture correctness from the perspective of objects in the database – as transactions perform operations on the objects. On the other hand, *transaction correctness properties* capture correctness from the perspective of the structure and behavior of transactions. That is, they deal, for example, with the results of transactions and the interactions between transactions. Serializability [Eswaran, 1976] captures database consistency requirements and transaction correctness properties via a single notion: (1) The state of the database at the end

of a set of concurrent transactions is the same as the one resulting from some serial execution of the same set of transactions; (2) The results of transactions and the interactions among the set of transactions are the same as the results and interactions, had the transactions executed one after another in this serial order. As applications using databases become more complex, the *correctness criteria* that are *acceptable* to the application become more complex and hence harder to capture using a single correctness notion.

Recent research has attempted to come up with correctness criteria, or *acceptability criteria*, that view these two types of requirements independently. The search for more flexible correctness requirements is also motivated by the introduction of new transaction models that extend the traditional atomic transaction model. (See [Elmagarmid, 1991] for a description of some of the *extended transaction models*.) These extensions came about because the atomic transaction model in conjunction with serializability is found to be very constraining when applied in advanced applications such as design databases that function in distributed, cooperative, and heterogeneous environments.

Proposed correctness criteria range from the standard serializability notion to eventual consistency [Sheth, 1990]. Quasi-serializability [Du, 1989], predicatewise serializability [Korth, 1988b], etc., are points that lie within this range. *Eventual consistency* can be viewed as a "catch-all" term with different connotations: For example, requiring consistency "at a specific real-time", "within some time" or "after a certain amount of change to some data", or enforcing consistency "after a certain value of the data is reached", etc. Whereas serializability and its relaxations are, in general, application and transaction model independent criteria, eventual consistency, as the examples above show, are application and transaction model specific. It is not difficult to see that these relaxed correctness requirements are useful within a single database as well as in multi-database environments.

Whereas serializability works under the simple assumption that individual transactions maintain the consistency of the database, proposed correctness criteria require more from the transaction developers. In particular, a transaction may have to be aware of the functionality of other transactions, especially in a cooperative environment. This makes transaction development as well as management more difficult. Our goal in this paper is to understand the conceptual similarities and differences between different correctness criteria without getting into the practical implications of adopting them.

So, in this paper we examine database consistency constraints and transaction correctness properties from the viewpoint of *what* the different dimensions of these two types of correctness are. This taxonomy allows us to categorize existing proposals thereby shedding some light on the similarities and differences between the proposals and to place them in perspective. The categorization also helps us determine whether or not a correctness notion is transaction model specific or application specific. We will see that even though some of the correctness notions were motivated by specific transaction models or specific applications, they have broader applicability.

To help in this categorization, we apply a uniform specification technique to express the various correctness criteria that have been proposed. The technique is based

on the ACTA formalism [Chrysanthis, 1990, 1991] which heretofore has been used for the specification of and for reasoning about extended transactions. One of the key ingredients of ACTA is the idea of constraining the occurrence of *significant events* associated with transactions, *Begin*, *Abort*, and *Split*, for example. These constraints are expressed in terms of necessary and sufficient conditions for events to occur. These, in turn, relate to the ordering of events and the validity of relevant conditions. Such constraints can also facilitate the specification of database consistency requirements and transaction correctness properties. The ACTA formalism is introduced in Section 3.

The rest of the paper is structured as follows: Subsection 2.1 provides a taxonomy of database consistency requirements while 2.2 provides a taxonomy of transaction correctness properties. A specification of existing proposals as well as their categorization (based on the taxonomy) is the subject of Section 4. Section 5 concludes the paper with some discussions of the next step in this work.

2 A TAXONOMY OF CORRECTNESS CRITERIA

We study different dimensions of the two aspects of correctness, namely, consistency of database state and correctness of transactions, in order to develop a taxonomy of correctness criteria. For concreteness, we give examples as the taxonomy is developed.

2.1 DATABASE CONSISTENCY REQUIREMENTS

Database consistency requirements can be examined with respect to two issues with further divisions of each as discussed below.

- *Consistency Unit:*

- *Complete Database:*

- All the objects in the database have to be consistent locally as well as mutually consistent, i.e., they should satisfy all the database integrity constraints typically specified in the form of predicates on the state of the objects.*

- Example: Traditional serializability (SR) applied to atomic transactions [Bernstein, 1987].*

- *Subsets of the objects in the database:*

- * *Location-independent subsets:*

- The database is viewed as being made up of subsets of objects. The subsets are not necessarily disjoint and are specified as part of the database definition. Each object in the database is expected to be consistent locally but mutual consistency is required only for objects that are within the same subset.

- Example: Setwise serializability (SSR) applied to compound transactions [Sha, 1985] and Predicatewise serializability (PSR) applied to cooperative transactions [Korth, 1988].*

- * *Location-dependent subsets:*

- Each subset corresponds to one of the sites of a (distributed / heteroge-

neous) database. In addition to mutual consistency among objects in a subset (i.e., site), consistency among subsets is also required depending on which parts of a database are accessed by a transaction.

Example: Quasi-serializability (QSR) [Du, 1989] and its generalization [Mehrotra, 1991] applied to distributed transactions.

– *Individual Objects:*

Each object in the database is expected to be consistent locally.

Example: Linearizability [Herlihy, 1987] applied to objects accessed by concurrent processes.

● *Consistency Maintenance:*

This is related to the issues of *when* a consistency requirement is expected to hold and *how* consistency is restored if it does not hold.

– *When is a consistency requirement expected to hold?*

* *At activity boundaries:* (An activity denotes a unit of work)

• *When an operation completes:*

When an operation on an object completes, the necessary consistency specifications must hold.

Example: Concurrent processes accessing shared objects.

• *When a set of operations completes:*

When a set of operations performed by a transaction completes, the necessary consistency is expected to hold.

Example: Semantic atomicity [Garcia-Molina, 1983] and multilevel atomicity [Lynch, 1983]

• *When a transaction completes:*

Consistency is expected to hold upon a transaction's completion.

Example: Atomic transactions.

• *When a set of transactions completes:*

Consistency is expected to hold not when individual transactions complete but when a set of transactions completes.

Example: Cooperative transactions [Korth, 1988b], sagas [Garcia-Molina, 1987].

* *At specific points of time:*

Consistency is required only at/after specific points in time. This is an example of *temporal consistency* [Sheth, 1990].

Example: A bank account is expected to be made consistent, with respect to the debits and credits that occur on a given day, upon closing of business.

* *At specific states:*

Objects may be required to be mutually consistent only when a certain number of updates have been made to one of the objects or a state satisfying a certain predicate is reached.

Example: A centralized database of a department store chain may require updates only upon the completion of 100 sales at a particular store. Such requirements are referred to as *spatial consistency* in [Sheth, 1990].

– *If a consistency requirement does not hold at a point it is supposed to, how is it restored?*

- * *Restored immediately:*
This applies when the consistency requirement is required to hold at activity boundaries. The activity is allowed to complete only if the requirement holds, i.e., completion is delayed until consistency holds.
Example: SR, PSR, QSR, and cooperative serializability (CoSR), applied to atomic, nested, and distributed transactions.
- * *Restored in a deferred manner:*
This typically applies when consistency is expected in certain states or at certain times. When it is applied to consistency that is expected at an activity boundary, the activity is allowed to complete and restoration is begun subsequently.
 - *Eventually:*
Consistency between objects must be restored eventually.
Example: If mutual consistency is required between two objects and one is changed by a transaction, another can be *triggered* to make changes in the other object.
 - *By a certain time:*
A deadline may be imposed on the time by which consistency is restored.
Example: In real-time systems, the state of the controlled environment should be reflected in the internal state of the controlling system within a certain time so that appropriate and timely control can be exercised.

2.2 TRANSACTION CORRECTNESS PROPERTIES

As was mentioned in the introduction, serializability suffices as a correctness criterion for traditional atomic transactions since once *individual transactions are guaranteed to take one consistent database state to another consistent state*, serializability guarantees that a set of concurrent transactions when started in a consistent state take the database to another consistent state. So the only transaction correctness property of interest is: Each transaction when executed by itself must maintain database consistency. From this it follows that, under serializability, the output of a transaction reflects a consistent database state. However, more elaborate correctness properties have been proposed in the context of additional application requirements and newer transaction models. These transaction correctness properties can be discussed with respect to three criteria:

- *Correctness of transaction results:*
 - *Absolute:*
The output of transactions must reflect a consistent database state.
Example: SR applied to atomic transactions, QSR applied to distributed transactions.
 - *Relative:*
Outputs of a transaction are considered correct even if they do not reflect a consistent state of the object, as long as they are within a certain bound of the result that corresponds to the consistent state.
Example: Epsilon-serializability (ESR) [Pu, 1991] applied to Epsilon-transactions, approximate query processing [Hou, 1989].

- *Correctness of transaction structure:*
Correctness depends on the (structural) relationship between transactions. This typically translates into *commit*, *abort*, *begin*, and other types of dependencies [Chrysanthis, 1991] between transactions.
Example: Sagas [Garcia-Molina, 1987], multi-level serializability [Korth, 1990c].
- *Correctness of transaction behavior:*
 - *Data access related behavior:*
Transactions are required to perform operations on objects in a certain manner to be considered correct.
Example: patterns [Skarra, 1991].
 - *Temporal behavior:*
Transactions have start time and completion time (deadline) constraints.
Example: Transactions in real-time systems.

The taxonomy just presented shows how the various weakened versions of serializability can be viewed from the perspectives of database consistency and transaction correctness. We revisit these notions in Section 4.1 where they are formally specified and categorized along the different dimensions of the taxonomy. Sections 4.2 and 4.3 deal with the formal specification of more general correctness criteria that are not directly related to serializability but deal, for example, with transaction structure and behavior, specific states of objects, or specific times.

3 A QUICK INTRODUCTION TO THE ACTA FORMALISM

ACTA is a first-order logic based formalism. As mentioned earlier, the idea of significant events underlies ACTA's specifications. Section 3.1 discusses these events. Specifications involve constraints on the occurrence of individual significant events as well as on the history of occurrence of these events. Hence the notion of history and the necessary and sufficient conditions for the occurrence of significant events are introduced in Section 3.2. Finally, Section 3.3 shows how sharing of objects leads to transaction inter-relationships which in turn induces certain dependencies between concurrent transactions.

3.1 SIGNIFICANT EVENTS ASSOCIATED WITH TRANSACTIONS

During the course of their execution, transactions invoke operations on objects. Also, they invoke transaction management primitives. For example, atomic transactions are associated with three transaction management primitives: *Begin*, *Commit* and *Abort*. The specific primitives and their semantics depend on the specifics of a transaction model [Chrysanthis, 1991]. For instance, whereas the *Commit* by an atomic transaction implies that it is terminating successfully and that all of its effects on the objects should be made permanent in the database, the *Commit* of a subtransaction of a nested transaction implies that all of its effects on the objects should be made persistent and visible with respect to its parent and sibling subtransactions. Other transaction management primitives include *Spawn*, found

in the nested transaction model [Moss, 1981], *Split*, found in the split transaction model [Pu, 1988], and *Join*, a transaction termination event also found in the split transaction model.

DEFINITION 3.1: Invocation of a transaction management primitive is termed a *significant event*. A transaction model defines the significant events that transactions adhering to that model can invoke.

The set of events invoked by a transaction t is a partial order with ordering relation $<_t$ where $<_t$ denotes the temporal order in which the related events occur.

$ts(\epsilon)$ gives the time of occurrence of event ϵ according to a globally synchronized clock¹. Clearly, $ts(\beta)$ will be larger than $ts(\alpha)$ if α appears earlier in the partial order. Further, no two significant events *that relate to the same transaction* can occur with the same ts value.

3.2 HISTORY, PROJECTION OF THE HISTORY, AND CONSTRAINTS ON EVENT OCCURRENCES

The concurrent execution of a set of transactions T is represented by the *history* [Bernstein, 1987] of the events invoked by the transactions in the set T and indicates the (partial) order in which these events occur. The partial order of the operations in a history is consistent with the partial order of the events of each individual transaction t in T .

The *projection* of a history H is a subhistory that satisfies a given criterion. For instance,

- The projection of a history H with respect to a specific transaction t yields a subhistory with just the events invoked by t . This is denoted by H_t .
- The projection of a history H with respect to a specific time interval $[i, j]$ yields the subhistory with the events which occurred between i and j (inclusive) and is denoted by $H^{[i,j]}$.

When $i =$ system initiation time, we drop the first element of the pair. Thus $H^j = H^{[system_init_time, j]}$ denotes all the events that occur until time j .

Consistency requirements imposed on concurrent transactions executing on a database can be expressed in terms of the properties of the resulting histories.

The occurrence of an event in a history can be constrained in one of three ways: (1) An event ϵ can be constrained to occur *only after* another event ϵ' ; (2) An event ϵ can occur *only if* a condition c is true; and (3) a condition c can *require* the occurrence of an event ϵ .

DEFINITION 3.2: The predicate $\epsilon \rightarrow \epsilon'$ is true if event ϵ *precedes* event ϵ' in history H . It is false, otherwise. (Thus, $\epsilon \rightarrow \epsilon'$ implies that $\epsilon \in H$ and $\epsilon' \in H$.)

¹This is obviously an abstraction - the effects of realizing this by a set of closely synchronized clocks on individual nodes in a distributed system will not be discussed here.

DEFINITION 3.3: $(\epsilon \in H) \Rightarrow Condition_H$, where \Rightarrow denotes *implication*, specifies that the event ϵ can belong to history H *only if* $Condition_H$ is satisfied. In other words, $Condition_H$ is *necessary* for ϵ to be in H . $Condition_H$ is a predicate involving the events in H .

Consider $(\epsilon' \in H) \Rightarrow (\epsilon \rightarrow \epsilon')$. This states that the event ϵ' can belong to the history H *only if* event ϵ occurs before ϵ' .

DEFINITION 3.4: $Condition_H \Rightarrow (\epsilon \in H)$ specifies that if $Condition_H$ holds, ϵ should be in the history H . In other words, $Condition_H$ is *sufficient* for ϵ to be in H .

We now describe some common types of constraints.

1. $(Commit_{t_j} \in H) \Rightarrow ((Commit_{t_i} \in H) \Rightarrow (Commit_{t_i} \rightarrow Commit_{t_j}))$. This says that if both transactions t_i and t_j commit then the commitment of t_i precedes the commitment of t_j . This Commit-Dependency is indicated by $(t_j \text{ CD } t_i)$. In general, $((Commit_{t_i} \in H) \Rightarrow condition)$ specifies that *condition* should hold for t_i to commit.
2. $(Abort_{t_i} \in H) \Rightarrow (Abort_{t_j} \in H)$ i.e., if t_i aborts then t_j aborts, states the Abort-Dependency of t_j on t_i ($t_j \text{ AD } t_i$). In general, $(condition \Rightarrow (Abort_{t_j} \in H))$ specifies that if *condition* holds t_j aborts.
3. $(Begin_{t_j} \in H) \Rightarrow (Begin_{t_i} \rightarrow Begin_{t_j})$ states that transaction t_j cannot begin executing until transaction t_i has begun.

3.3 OBJECTS, OPERATIONS, AND CONFLICTS

A transaction accesses and manipulates the objects in the database by invoking operations specific to individual objects. It is assumed that an operation always produces an output (return value), that is, it has an outcome (condition code) or a result. The result of an operation on an object depends on the current state of the object. For a given state s of an object, we use $return(s, p)$ to denote the output produced by operation p , and $state(s, p)$ to denote the state produced after the execution of p .

DEFINITION 3.5: Invocation of an operation on an object is termed an *object event*. The type of an object defines the object events that pertain to it. We use $p_t[ob]$ to denote the object event corresponding to the invocation of the operation p on object ob by transaction t . Object events are also part of the history H .

DEFINITION 3.6: Let $H^{(ob)}$ denote the projection of the history with respect to the operations on ob . Two operations p and q *conflict* in a state produced by $H^{(ob)}$, denoted by $conflict(H^{(ob)}, p, q)$, iff

$$\begin{aligned} (state(H^{(ob)} \circ p, q) \neq state(H^{(ob)} \circ q, p)) & \vee \\ (return(H^{(ob)}, q) \neq (return(H^{(ob)} \circ p, q))) & \vee \\ (return(H^{(ob)}, p) \neq (return(H^{(ob)} \circ q, p))) & \end{aligned}$$

Two operations that do not conflict are *compatible*.

(\circ denotes functional composition; $H \circ p$ appends p to history H .) Thus, two operations conflict if their effects on the state of an object are not independent of their execution order (first clause) or their return values are not independent of their execution order (second and third clauses). From now on, we drop the first parameter of conflict, namely, $H^{(ob)}$.

4 SPECIFICATION AND CATEGORIZATION OF CORRECTNESS CRITERIA

In this section, we study various database consistency requirements and transaction correctness properties that have been proposed and place them in perspective, given the taxonomy of the previous section. *Broadly speaking*, Section 4.1 deals with transaction model and application independent correctness criteria (even though, as we will see, those who proposed them may have had a specific transaction model or application in mind), Section 4.2 discusses transaction model dependent but application independent criteria, and Section 4.3 examines transaction and application dependent consistency requirements. (For a complete axiomatic semantics of the various extended transaction models, the reader is referred to [Chrysanthis, 1991].)

4.1 TRANSACTION AND APPLICATION INDEPENDENT CRITERIA

In general, transaction and application independent correctness criteria are extensions to serializability. In this section, we first specify some of these extensions using the formalism described in the previous section and then use the specifications to show how the different extensions relate to each other. All of these criteria are based on the notion of *conflicts* and their preservation in equivalent histories. Thus, we do not discuss correctness criteria such as *view serializability* [Yannakakis, 1984] that are not as easy to realize.

DEFINITION 4.7: Let \mathcal{R} be a binary relation on a set of transactions T , $t_i, t_k \in T$, $t_i \neq t_k$. \mathcal{R}^* is the transitive-closure of \mathcal{R} ; i.e.,
 $(t_i \mathcal{R}^* t_k)$ if $[(t_i \mathcal{R} t_k) \vee \exists t_j \in T ((t_i \mathcal{R} t_j) \wedge (t_j \mathcal{R}^* t_k))]$.

4.1.1 Serializability

In traditional databases, serializability and, in particular, *conflict serializability*, is the well-accepted criterion for concurrency control.

Let \mathcal{C} be a binary relation on transactions in T .

Let H be the history of events relating to committed transactions in T .

DEFINITION 4.8: $\forall t_i, t_j \in T, t_i \neq t_j$,
 $(t_i \mathcal{C} t_j)$ if
 $\exists ob \exists p, q (\text{conflict}(p_{t_i}[ob], q_{t_j}[ob]) \wedge (p_{t_i}[ob] \rightarrow q_{t_j}[ob]))$.

DEFINITION 4.9: H is (conflict) serializable iff
 $\forall t \in T, \neg(t C^* t)$.

To illustrate the practical implications of these definitions, note that the C relation captures the fact that two transactions have invoked conflicting operations on the same object and the order in which they have invoked the conflicting operations. Consequently, the C relation captures direct conflicts between transactions in a history as well as their serialization order. The fact that a serialization order is acyclic is stated by requiring that there be no cycles in the C relation.

Note also that the above definitions do not involve any significant events. This reflects the fact that serializability per se does not constrain the occurrence of any significant event, e.g., a *Commit* event to happen only after another *Commit* event. (If the C relationship between transactions is acyclic, transactions in H can commit in any order.) That is, the *commit order* of transactions is not necessarily the same as their serialization order and hence, the commit order cannot be used to induce the serialization order. However, a commit order induced by a C relation is consistent with the serialization order. For example, consider the case of *rigorous histories* [Breitbart, 1991] such as the ones produced by the strict two-phase locking protocol [Eswaran, 1976]. In this case, if transactions t_i and t_j have a C relationship, i.e., they have invoked conflicting operations, a commit dependency [Chrysanthis, 1991] forms between t_i and t_j . (Conflicting operations may also produce abort dependencies between the invoking transactions; but an abort dependency implies a commit dependency.) By requiring that the C relation be acyclic, commit dependencies must also be acyclic. By inducing a commit dependency between every pair of transactions invoking conflict operations, the commit order specified by the commit dependencies is the same as the serialization order.

With respect to the taxonomy of Section 2, for serializability, the consistency unit is the complete database, consistency is required at transaction boundaries, and immediate consistency restoration is required. Absolute correctness of transaction results is expected. Atomic transactions and top-level transactions of nested transactions, for example, behave according to the serializability correctness criterion.

The semantics of the operations on the objects (for example, see [Badrinath, 1990][Herlihy, 1987][O'Neil, 1986]) can be used to define the *conflict* relationship between operations. Furthermore, different *degrees* of consistency [Gray, 1975] can be ensured by ignoring some of the conflicts. The resulting inconsistencies can be accommodated in applications that can cope with such inconsistencies or when these are masked by the structuring of the objects used by the applications. The former is the case in [Gray, 1975] and with ESR [Pu, 1991] (see [Ramamritham, 1991b] for a formal characterization of ESR). The latter is the case with abstract serializability – used in the context of multi-level transactions [Beeri, 1989][Moss, 1986][Martin, 1988][Badrinath, 1990].

4.1.2 Predicatewise Serializability

Predicatewise serializability has been proposed in [Korth, 1988, 1988b] as the correctness criterion for concurrency control in databases in which consistency constraints are in a conjunctive normal form. In such cases, consistency constraints can be maintained by requiring serializability only with respect to objects which

relate to a disjunctive clause.

Let $P = (P_1 \wedge P_2 \dots \wedge P_n)$ be the database consistency constraint. Suppose the disjunctive clause P_k relates to objects in $D_k \subseteq DB$, where DB is the database.

$\forall k \in \{1 \dots n\}$, let C_k be a binary relation on transactions in T .

Let H be the history of events relating to committed transactions in T .

DEFINITION 4.10: $\forall k \in \{1 \dots n\}, \forall t_i, t_j \in T, t_i \neq t_j,$
 $(t_i C_k t_j),$ if
 $\exists ob \in D_k \exists p, q (conflict(p_{t_i}[ob], q_{t_j}[ob]) \wedge (p_{t_i}[ob] \rightarrow q_{t_j}[ob])).$

DEFINITION 4.11: H is *predicatewise serializable* iff
 $\forall t \in T \forall D_k 1 \leq k \leq n \neg(t C_k^* t)$

In [Sha, 1985], each D_k is said to be an *atomic data set*. With respect to the taxonomy, for predicatewise serializability, an atomic data set [Sha, 1985] forms a consistency unit, consistency is required at transaction boundaries, and immediate consistency restoration is required. Absolute correctness of transaction results is expected. Compound transactions [Sha, 1985] behave according to the predicatewise serializability correctness criterion.

4.1.3 Cooperative Serializability

We define *cooperative serializability* (CoSR) with respect to a set of transactions which maintain some consistency properties. Transactions form cooperative transaction sets. A cooperative transaction set could be formed by the components of an extended transaction or transactions collaborating over some objects while maintaining the consistency of the objects. In such cases, consistency can be maintained if other transactions which do not belong to the set are serialized with respect to all the transactions in the set. In other words, the set of cooperative transactions becomes the unit of serializability.

Let T_c be a set of cooperative transactions, $T_c \subseteq T$.

Let C_c be a binary relation on transactions in T .

Let H be the history of events relating to committed transactions in T .

DEFINITION 4.12: $\forall t_i, t_j, t_k \in T, t_i \neq t_j, t_i \neq t_k, t_j \neq t_k \forall T_c \subseteq T$
 $(t_i C_c t_j),$ if
 $\exists ob \exists p, q ((t_i \notin T_c, t_j \notin T_c (conflict(p_{t_i}[ob], q_{t_j}[ob]) \wedge (p_{t_i}[ob] \rightarrow q_{t_j}[ob]))) \vee$
 $(t_i \notin T_c, t_j \in T_c, t_k \in T_c (conflict(p_{t_i}[ob], q_{t_k}[ob]) \wedge (p_{t_i}[ob] \rightarrow q_{t_k}[ob]))) \vee$
 $(t_i \in T_c, t_j \notin T_c, t_k \in T_c (conflict(p_{t_k}[ob], q_{t_j}[ob]) \wedge (p_{t_k}[ob] \rightarrow q_{t_j}[ob])))$

In this definition, the first clause expresses how a dependency between two transactions which do not belong to the same set is directly established when they invoke conflicting operations on a shared object. This is similar to the clause in the classical definition of (conflict) serializability. The other two clauses reflect the fact that when a transaction establishes a dependency with another transaction, the same dependency is established between all the transactions in their corresponding

cooperative transactions sets. These clauses can be viewed as expressions of the development of dependencies between transaction sets.

DEFINITION 4.13: H is *cooperative serializable* iff

$$\forall t \in T \neg(t C_c^* t)$$

With respect to the taxonomy of the previous section, for cooperative serializability, the consistency unit is the complete database, consistency is required when an ordinary transaction (not a member of a T_c) completes or a set of cooperating transactions complete, and immediate consistency restoration is required. Absolute correctness of transaction results is expected. The correctness requirement expressed informally in [Martin, 1990] corresponds to CoSR.

4.1.4 Quasi Serializability

Quasi Serializability has been proposed in [Du, 1989] as a correctness criterion for maintaining transaction consistency in multidatabases, e.g., heterogeneous distributed databases. In these systems, transactions can either execute on a single site (called *local* transactions), or can execute on multiple sites (called *global* transactions).

In QSR, the correctness of the execution of a set of global and local transactions is based on the notion of a *quasi serial history* which, unlike a serial history, specifies that only global transactions are executed serially. A history is *quasi serial* if (1) all local histories are (conflict) serializable, and (2) there exists a total order of all global transactions g_m and g_n where g_m precedes g_n in the order and all g_m 's operations precede g_n 's operations in all local histories in which they both appear. A quasi serializable history is equivalent to a quasi serial history.

Let G be the set of global transactions and g_n^i be a (sub)transaction of a global transaction g_n ($g_n \in G$) executing all the operations of g_n on site i^2 .

Let T_i be the set of transactions, both local transactions and global (sub)transactions, executing on site i . $T = (\cup_i T_i)$.

Let H be the history of events relating to committed transactions in T .

Let \mathcal{R} be a binary relation on a set of global transactions G .

DEFINITION 4.14: $\forall g_m, g_n \in G, g_m \neq g_n$,
 $(g_m \mathcal{R} g_n)$, if

$$\exists k, g_m^k, g_n^k (g_m^k C g_n^k) \vee \exists p, t_0 = g_m^k, t_p = g_n^k \forall i, 0 \leq i \leq p-2$$

$$\exists ob \exists p, q ((\text{conflict}(p_{t_i}[ob], q_{t_{i+1}}[ob]) \wedge (p_{t_i}[ob] \rightarrow q_{t_{i+1}}[ob])) \wedge$$

$$\exists ob' \exists p', q' (\text{conflict}(p'_{t_{i+1}}[ob'], q'_{t_{i+2}}[ob']) \wedge (p'_{t_{i+1}}[ob'] \rightarrow q'_{t_{i+2}}[ob'])) \wedge$$

$$(q_{t_{i+1}}[ob] \rightarrow p'_{t_{i+1}}[ob']))$$

where C is the binary relation defined in Definition 4.8.

The C relation captures the fact that two global transactions *directly conflict* in a local history when they invoked conflicting operations on a share object. Two

²That is, at most one (sub)transaction of a global transaction can execute on a particular site.

global transactions might also *indirectly conflict* in a local history even if they don't access any share objects. Indirect conflicts are introduced by other transactions that directly conflict with each other and with the global transactions. These, indirect conflicts between two transactions, particularly those introduced by local transactions, are captured by the second clause of the definition of \mathcal{R} . Note that this clause, and consequently \mathcal{R} , is not equivalent to the transitive-closure of \mathcal{C} which does not place any restriction on the execution ordering of the conflicting operations, but $\mathcal{R} \subset \mathcal{C}^*$.

DEFINITION 4.15: H is *quasi serializable* iff

1. $\forall i \forall t \in T_i \neg(t \mathcal{C}^* t)$, and
2. $\forall g \in G \neg(g \mathcal{R}^* g)$

It should be pointed that since the \mathcal{R} relation captures both direct and indirect conflicts between two global transactions in a history, the serializable execution of global transactions is in terms of both direct and indirect conflicts. Indirect conflicts between local transactions induced by conflicts of global transactions that execute on multiple sites are not captured by either clause, the reason being that QSR assumes no data dependency across sites.

With respect to the taxonomy of the previous section, for Quasi serializability, (site-based) subsets of the database objects form the consistency units (i.e., objects in each site form a subset) and consistency should hold when a transaction completes. Absolute correctness of transactions' results is expected. Quasi serializability has been proposed in the context of distributed (multi-database) transactions.

4.1.5 Relationship between different Serializability-based Correctness Criteria

Thus far in this section, we have specified four serializability-based correctness criteria using the ACTA formalism and classified them with respect to our taxonomy in Section 2. Here, we will use the formal definitions of the various criteria to relate them to each other.

According to predicatewise serializability, each D_k is associated with a C_k and hence conflicting transactions can be serialized differently with respect to different D_k . This is contrary to serializability which permits only a single system-wide serialization order involving conflicting transactions based on \mathcal{C} . However, if D_k is the complete database, then $\mathcal{C} = C_k$, and consequently, predicatewise serializability is equivalent to serializability.

In the case of cooperative serializability, transactions in different cooperative transaction sets may be related by the C_c relation (individual transactions not belonging to any cooperative set can be viewed as singleton cooperative transaction sets). Hence, if each cooperative transaction set has just one member, then $C_c = \mathcal{C}$ and, in this case, cooperative serializability is equivalent to serializability.

In the case of quasi-serializability, there are two distinct conditions under which QSR is equivalent to serializability. These correspond to the situations in which one of the two clauses of the definition of QSR is trivially true: (1) In the absence of global transactions, transactions in T_i are serialized based only on \mathcal{C} . (2) In the

absence of local transactions, transactions in T_i are serialized based only on \mathcal{R} , i.e., here $\mathcal{R} = \mathcal{C}$. Indirect conflicts due to local transactions are not possible whereas indirect conflicts due to global transactions are considered by \mathcal{C}^* .

Finally, we would like to point out that these different correctness criteria can be combined and/or adopted within a single database. For instance, it is easy to picture how cooperatively serializability can be used in conjunction with predicatewise serializability and how cooperatively serializability can be used in conjunction with even quasi-serializability.

4.2 TRANSACTION MODEL DEPENDENT AND APPLICATION INDEPENDENT CRITERIA

Transaction model dependent but application independent correctness criteria are typically related to the structure of transactions that conform to a particular model. (Note that specific transaction models may be more suited to specific applications.) As was mentioned earlier, different transaction models produce different transaction structures where the structure of an extended transaction defines its component transactions and the relationships between them. Dependencies can express these relationships and thus, can specify the links in the structure. For example, in hierarchically-structured nested transactions, the parent/child relationship is established at the time the child is *spawned*. This is expressed by a child transaction t_c establishing a weak-abort dependency (defined below) on its parent t_p ($(t_c \text{ WD } t_p)$) and by a parent establishing a commit dependency on its child ($(t_p \text{ CD } t_c)$). The weak-abort dependency guarantees the abortion of an uncommitted child if its parent aborts whereas the commit dependency prevents a child from committing after its parent has committed.

We now formally specify some of the dependencies that can occur in addition to the Commit Dependency, Abort Dependency, and Begin Dependency specified in Section 3.2.

Let t_i and t_j be two transactions and H be a finite history which contains all the events pertaining to t_i and t_j .

Weak-Abort Dependency ($t_j \text{ WD } t_i$): if t_i aborts and t_j has not yet committed, then t_j aborts. In other words, if t_j commits and t_i aborts then the commitment of t_j precedes the abortion of t_i in a history; i.e.,
 $(\text{Abort}_{t_i} \in H) \Rightarrow (\neg(\text{Commit}_{t_j} \rightarrow \text{Abort}_{t_i}) \Rightarrow (\text{Abort}_{t_j} \in H))$.

Strong-Commit Dependency ($t_j \text{ SCD } t_i$): if transaction t_i commits then t_j commits; i.e., $(\text{Commit}_{t_i} \in H) \Rightarrow (\text{Commit}_{t_j} \in H)$.

Termination Dependency ($t_j \text{ TD } t_i$): t_j cannot commit or abort until t_i either commits or aborts; i.e., $(\epsilon' \in H) \Rightarrow (\epsilon \rightarrow \epsilon')$
 where $\epsilon \in \{\text{Commit}_{t_i}, \text{Abort}_{t_i}\}$, and $\epsilon' \in \{\text{Commit}_{t_j}, \text{Abort}_{t_j}\}$.

Exclusion Dependency ($t_j \text{ ED } t_i$): if t_i commits and t_j has begun executing, then t_j aborts (both t_i and t_j cannot commit); i.e.,
 $(\text{Commit}_{t_i} \in H) \Rightarrow ((\text{Begin}_{t_j} \in H) \Rightarrow (\text{Abort}_{t_j} \in H))$.

Force-Commit-on-Abort Dependency ($t_j \text{ CAD } t_i$): if t_i aborts, t_j commits; i.e.,
 $(\text{Abort}_{t_i} \in H) \Rightarrow (\text{Commit}_{t_j} \in H)$.

Serial Dependency (t_j *SD* t_i): transaction t_j cannot begin executing until t_i either commits or aborts; i.e.,

$$(Begin_{t_j} \in H) \Rightarrow (\epsilon \rightarrow Begin_{t_i}) \text{ where } \epsilon \in \{Commit_{t_i}, Abort_{t_i}\}.$$

Begin-on-Commit Dependency (t_j *BCD* t_i): transaction t_j cannot begin executing until t_i commits; i.e., $(Begin_{t_j} \in H) \Rightarrow (Commit_{t_i} \rightarrow Begin_{t_j})$.

Begin-on-Abort Dependency (t_j *BAD* t_i): transaction t_j cannot begin executing until t_i aborts; i.e., $(Begin_{t_j} \in H) \Rightarrow (Abort_{t_i} \rightarrow Begin_{t_j})$.

Weak-begin-on-Commit Dependency (t_j *WCD* t_i): if t_i commits, t_j can begin executing after t_i commits; i.e.,

$$(Begin_{t_j} \in H) \Rightarrow ((Commit_{t_i} \in H) \Rightarrow (Commit_{t_i} \rightarrow Begin_{t_j})).$$

Let us look at further examples of structure-related transaction correctness properties. In the transaction model proposed in [Buchmann, 1990, Garcia-Molina, 1991] a parent can commit only if its *vital* children commit, i.e., a parent transaction has an abort dependency on its *vital* children t_v (t_p *AD* t_v). Child transactions may also have different dependencies with their parents if the transaction model supports various spawning or coupling modes [Dayal, 1990]. Sibling transactions may also be interrelated in several ways. For example, components of a *saga* [Garcia-Molina, 1987] can be paired according to a compensated-for/compensating relationship [Korth, 1990]. Relations between a compensated-for and compensating transactions as well as those between them and the saga can be specified via begin-on-commit dependency *BCD*, begin-on-abort dependency *BAD*, force-commit-on-abort dependency *CAD* and strong-commit dependency *SCD* [Chrysanthis, 1992]. In a similar fashion, dependencies that occur in the presence of alternative transactions and contingency transactions [Buchmann, 1990] can also be specified [Chrysanthis, 1992].

4.3 TRANSACTION AND APPLICATION DEPENDENT CRITERIA

We now focus on the required *behavior* of a transaction and hence on the requirements imposed by the application that employs that specific transaction. We distinguish between two types of behavior related properties:

1. Those that relate to constraints on a transaction's access to objects – beyond those mandated by the concurrency properties of the objects.
2. Those that relate to properties that deal with its other behavioral properties, such as, *when* a transaction can/must begin and *when* it can/must end. Spatial and temporal requirements are related to this type.

We elaborate upon the first type through an example. Consider a *page* object with the standard *read* and *write* operations, where read and write operations conflict. A read's return-value is dependent on a previous write, whereas a write's return-value is independent of a read or another write. In addition, consider transactions which have the ability to reconcile potential read-write conflicts: When a transaction t_i reads a page x and another transaction t_j subsequently writes x , t_i and t_j can commit in any order. However, if t_j commits before t_i commits, t_i must reread x in order to commit. This is captured by the following requirement:

$$(read_{t_i}[x] \rightarrow write_{t_j}[x]) \Rightarrow ((Commit_{t_j} \rightarrow Commit_{t_i}) \Rightarrow (Commit_{t_j} \rightarrow read_{t_i}[x])).$$

In this example, t_i has to reread the page x when, subsequent to the first read, the page is written and committed by t_j . In general, t_i may need to invoke an operation on the same or a different object. For instance, instead of x , t_i may have to read a *scratch-pad* object which t_i and t_j use to determine and reconcile potential conflicts. In general, the specification of correct transaction behavior can include the specification of operations that need to be controlled to produce correct histories as well as the specification of operations that *have* to occur in correct histories. These correspond to *conflicts* and *patterns* in [Skarra, 1991].

Let us now turn to other behavioral specifications, for example, those that concern the beginning and termination of transactions. Consider the following simple requirement which states that if *condition* is true then transaction t_j must begin.

$$condition \Rightarrow (Begin_{t_j} \in H)$$

condition can depend on the occurrence of an event, on the state of the database, and on time. As we will see, the above requirement can be used for the flexible enforcement of consistency, to trigger the propagation of changes, to react to consistency violations, and to notify changes. Thus, the above specification can be considered to be a specification for the automatic triggering of situation-dependent actions, e.g., for expressing the rules that govern the triggering of actions in an active database [Dittrich, 1991].

Suppose *condition* is related to the occurrence of some significant event within a transaction t_i . In this case, the additional structural relationships (for instance, the different *coupling modes* [Dayal, 1990]) between t_i and t_j can be specified via the dependencies discussed in Section 4.2.

If *condition* relates to the state of the database, what we have is related to *spatial consistency* discussed in [Sheth, 1990]. For instance, consider the following *condition*: "One hundred sales have occurred at this store since the master database at the store's headquarters was last updated."

If *condition* relates to time, for instance, if *condition* is "*time* > 8pm", we have a *temporal consistency* requirement.

Now let us consider situations where *constraints* are placed on the beginning of transactions. For example, a transaction t_i to compute daily interest can start after midnight but only after the day's withdrawals and deposits have been reflected in the account (say by a transaction t_j). This can be specified as

$$(ts(Begin_{t_i}) \geq 12am) \wedge (t_i \text{ BCD } t_j).$$

This is an example where a transaction has a time-based start-dependency as well as a begin-on-commit dependency on another transaction.

Let us consider another example. If a deposit is made by time x then the transaction that reflects it in the account should not be started until time y . This is specified by

$$((Commit_{t_i} \in H) \wedge (ts(Commit_{t_i}) < x)) \Rightarrow ((Begin_{t_j} \in H) \Rightarrow (ts(Begin_{t_j}) > y)).$$

Through several examples, we now consider requirements and constraints associated with the termination of transactions.

Sometimes, we may want to specify that some specific change of state (by one transaction) triggers [Dayal, 1990] another transaction (that perhaps fixes the inconsistency resulting from the first transaction). Clearly, this type of constraint is related to deferred consistency restoration. This can occur, for example, if we had two versions of a database, one which was complete and another (at perhaps a different site) which only contained data required at that site. The two are not required to be consistent at all times but changes done to the complete database are required to percolate to the other within a specified delay. If the changes should be reflected within d units of time, we have the following "temporal commit dependency":

$$((Commit_{t_i} \in H) \wedge (ts(Commit_{t_i}) = t)) \Rightarrow (Commit_{t_j} \in H^{t+d}).$$

This says that if t_i commits at time t , t_j should commit by time $t + d$. For another example, consider the following:

$$(temperature \geq threshold \wedge time = t) \Rightarrow (Commit_{t_j} \in H^{t+function(temperature)})$$

Here t_j could be a transaction that opens a valve to pass more coolant into the reactor whose temperature is above *threshold*. The length of time available to complete this transaction is a function of the current temperature. This is a form of triggered transaction but with specific time constraints imposed on its completion [Korth, 1990b]. Such time constrained activities occur in real-time databases [Ramamritham, 1992].

In some situations, it may be desirable to specify an interval $[l, u]$ such that t_j does not commit before l (the lower bound) but definitely commits before u (the upper bound). For example, consider deposits into a bank account. During the day, if a deposit is made before 3pm, it is just "logged" into a file but is reflected in the appropriate account between 10pm and 4am that night. Such constraints take the form

$$((Commit_{t_i} \in H) \wedge (ts(Commit_{t_i}) < t)) \Rightarrow (Commit_{t_j} \in H^{[l,u]})$$

The above conditions imposed on the initiation and termination of transactions can be viewed as generalizations of the preconditions and postconditions associated with specific transactions [Korth, 1988b].

For a final example of behavior related specifications, consider the situation in which it may be desirable to prevent a transaction t_i from aborting after a time t . This corresponds to the assumption that a transaction is implicitly committed if it has not aborted by a certain time [Rusinkiewicz, 1990]. For example, no bets can be canceled after a race is started and a lottery ticket cannot be refunded after a given time.

$$(Abort_{t_i} \notin H^t) \Rightarrow (Commit_{t_i} \in H^{t+1})$$

Other examples of behavior related requirements appear in [Ramamritham, 1991].

5 CONCLUSIONS

In this paper, we have examined different types of acceptability criteria and have attempted to provide a taxonomy with respect to *database consistency requirements* and *transaction correctness properties*. Given space limitations, we could examine, in detail, only a subset of the proposals that have been made to capture the correctness properties applicable to extended transaction models as well as those demanded by the newer database applications.

We have approached the problem of categorizing the different proposals by formally specifying them using the framework of ACTA. This allows us to clearly see where one proposal differs from another and what its relationship with serializability is.

We believe this taxonomy to be a good starting point in our endeavor to classify proposed correctness criteria and to compare and contrast them. It can be viewed as a common framework with respect to which one can study where a new correctness criterion fits and how it relates to existing criteria. In this regard, we expect the taxonomy to evolve as better understanding is gained about the correctness needs of emerging database applications.

Let us now examine some of the other implications of this work. In the context of a multi-database system, the specifications of database consistency and transaction correctness can be viewed as requirements on the coordinator of the blackboxes [Breitbart, 1990] that control individual databases. We would like to apply the reasoning capabilities of the ACTA formalism to study the properties of mechanisms, such as in [Sheth, 1991], for maintaining relaxed correctness properties of interdependent data. In the same context, we would like to investigate ways in which the ACTA primitives themselves can be used as part of these mechanisms [Rusinkiewicz, 1991, Sheth, 1992]. Similarly, in the context of active databases, we can see how the semantics of the rules that govern the triggering of actions can be formally specified. In addition, the relationships between the triggering action and the triggered action can also be expressed precisely using dependencies. Just as we were able to reason about extended transactions using ACTA [Chrysanthis, 1991], we see the formalization of different aspects of active databases as the starting point for addressing issues, such as, reasoning about the consequences of rule firings, changes to rules, and coupling modes.

Acknowledgements

The authors thank Alex Buchmann for lively discussions about notions of consistencies and correctness, and Nandit Soparkar, Amit Sheth and Greg Speegle for their comments on previous versions of this paper.

This material is based upon work supported by the N.S.F. under grants IRI-9109210 and CDA-8922572 and a grant from the University of Pittsburgh.

References

Badrinath B. R. and K. Ramamritham. (1990) Performance Evaluation of Semantics-based Multilevel Concurrency Control Protocols. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pages 163-172.

- Beeri, 1989 Beeri C., P. A. Bernstein, and N. Goodman. (1989) A Model for Concurrency in Nested Transaction Systems. *Journal of the ACM*, 36(2):230-269.
- Bernstein P. A., V. Hadzilacos, and N. Goodman. (1987) *Concurrency Control and Recovery in Database Systems*. Addison-Wesley.
- Breitbart Y. et al. (1990) Final Report of the Workshop on Multidatabases and Semantic Interoperability.
- Breitbart Y., D. Georgakopoulos, M. Rusinkiewicz and A. Silberschatz. (1991) On Rigorous Transaction Scheduling, *IEEE Transactions on Software Engineering*, 17(9):954-960.
- Buchmann. A, M. Hornick, E. Markatos, and C. Chronaki. (1990) Specification of a Transaction Mechanism for a Distributed Active Object System. In *Proceedings of the OOPSLA/ECOOOP 90 Workshop on Transactions and Objects*, pages 1-9.
- Chrysanthis P. K. and K. Ramamritham. (1990) ACTA: A Framework for Specifying and Reasoning about Transaction Structure and Behavior. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pages 194-203.
- Chrysanthis P. K. and K. Ramamritham. (1991) A Formalism for Extended Transaction Models. In *Proceedings of the seventeenth International Conference on Very Large Databases*, pages 103-112.
- Chrysanthis P. K. and K. Ramamritham. (1992) ACTA: The SAGA Continues. In A. K. Elmagarmid, editor, *Database Transaction Models for Advanced Applications*. Morgan Kaufmann, 1992.
- Dayal U., M. Hsu, and R. Ladin. (1990) Organizing Long-Running Activities with Triggers and Transactions. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pages 204-214.
- Dittrich K. R. and U. Dayal. (1991) Active Database Systems (Tutorial Notes). In *The Seventeenth International Conference on Very Large Databases*.
- Du W. and A. K. Elmagarmid. (1989) Quasi Serializability: a Correctness Criterion for Global Concurrency Control in InterBase. In *Proceedings of the Fifteenth International Conference on Very Large Databases*, pages 347-355.
- Elmagarmid A. (Editor). (1991) Special Issue on Unconventional Transaction Management. *Bulletin of the IEEE Technical Committee on Data Engineering*, 14(1).
- Eswaran K. P., J. N. Gray, R. A. Lorie, and I. L. Traiger. (1976) The Notion of Consistency and Predicate Locks in a Database System. *Communications of the ACM*, 19(11):624-633.
- Garcia-Molina H. (1983) Using semantic knowledge for transaction processing in a distributed database. *ACM Transactions on Database Systems*, 8(2):186-213.
- Garcia-Molina H., D. Gawlick, J. Klein, K. Kleissner, and K. Salem. (1991) Modeling Long-Running Activities as Nested Sagas. *Bulletin of the IEEE Technical Committee on Data Engineering*, 14(1):14-18.

- Garcia-Molina H. and Kenneth Salem. (1987) SAGAS. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pages 249–259.
- Gray J. N., R. A. Lorie, G. R. Putzulo, and I. L. Traiger. (1975) Granularity of locks and degrees of consistency in a shared database. In *Proceedings of the First International Conference on Very Large Databases*, pages 25–33.
- Herlihy M. P. and W. Weihl. (1988) Hybrid concurrency control for abstract data types. In *Proceedings of the Seventh ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pages 201–210.
- Herlihy M. P. and J. M. Wing. (1987) Axioms for Concurrent Objects. In *Proceedings of the Fourteenth ACM Symposium on Principles of Programming Languages*.
- Hou. W, G. Ozsoyoglu, and B. K. Taneja. (1989) Processing Aggregate Relational Queries with Hard Time Constraints. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pages 68–67.
- Korth H. F., W. Kim, and F. Bancilhon. (1988) On Long-Duration CAD Transactions. *Information Sciences*, 46(1-2):73–107.
- Korth H. F. and G. Speegle. (1988b) Formal Models of Correctness without Serializability. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pages 379–386.
- Korth H. F., E. Levy, and A. Silberschatz. (1990) Compensating Transactions: A New Recovery Paradigm. In *Proceedings of the Sixteenth International Conference on Very Large Databases*, pages 95–106.
- Korth H. F., N. Soparkar, and A. Silberschatz. (1990b) Triggered Real-Time Databases with Consistency Constraints. In *Proceedings of the Sixteenth International Conference on Very Large Databases*, pages 71–82
- Korth H. F. and G. Speegle. (1990c) Encapsulation of Transaction Management in Object Databases. In *Proceedings of the OOPSLA/ECOOP'90 Workshop on Transactions and Objects*, pages 27–32.
- Lynch N. A. (1983) Multilevel atomicity — A new correctness for database concurrency control. *ACM Transactions on Database Systems*, 8(4):484–502.
- Martin B. E. (1988) Scheduling protocols for nested objects. Technical Report CS-094, Department of Computer Science and Engineering, University of California, San Diego.
- Martin, 1990 Martin B E. and C. Pedersen. (1990) Long-Lived Concurrent Activities. Technical Report HPL-90-178, HP Laboratories.
- Mehrotra S., R. Rastogi, H. Korth, and A. Silberschatz. (1991) Non-serializable Executions in Heterogeneous Distributed Database Systems. In *Proceedings of the First International Conference on Parallel and Distributed Information Systems*.
- Moss J. E. B. (1981) *Nested Transactions: An approach to reliable distributed computing*. PhD thesis, Massachusetts Institute of Technology.

- Moss J. E. B., N. Griffeth, and M. Graham. (1986) Abstraction in recovery management. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pages 72–83.
- O’Neil P. E. (1986) The Escrow Transactional Method. *ACM Transactions on Database Systems*, 11(4):405–430.
- Pu C., G. Kaiser, and N. Hutchinson. (1988) Split-Transactions for Open-Ended activities. In *Proceedings of the Fourteenth International Conference on Very Large Databases*, pages 26–37.
- Pu C. and A. Leff. (1991) Replica Control in Distributed Systems: An Asynchronous Approach. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pages 377–386.
- Ramamritham K., P. K. Chrysanthis, and A. Buchmann. (1991) Degrees of Consistency – Specification and Maintenance. Technical report, Computer Science Department, University of Massachusetts.
- Ramamritham K. and C. Pu. (1991b) A Formal Characterization of Epsilon Serializability. (*Submitted for publication*).
- Ramamritham K. (1992) Real-Time Databases. (*To appear in*) *Int’l Journal of Distributed and Parallel Databases*.
- Rusinkiewicz M., A. Elmagarmid, Y. Leu, and W. Litwin. (1990) Extending the Transaction Model to Capture more meaning. *ACM Sigmod Record*, 19(1):3–7.
- Rusinkiewicz M., A. Sheth, and G. Karabatis. (1991) Specification of Dependencies for the Management of Interdependent Data. *IEEE Computer*, 12(12):46–54.
- Sha L. (1985) *Modular concurrency control and failure recovery— Consistency, Correctness and Optimality*. PhD thesis, Carnegie-Mellon University.
- Sheth A. and M. Rusinkiewicz. (1990) Management of Interdependent Data: Specifying Dependency and Consistency Requirements. In *Proceedings of the Workshop on the Management of Replicated Data*, pages 133–136.
- Sheth. A, Y. Leu, and A. Elmagarmid. (1991) Maintaining consistency of interdependent data in multidatabase systems. Technical Report CSD-TR-91-016, Computer Science Department, Purdue University.
- Sheth A., M. Rusinkiewicz, and G. Karabatis. (1992) Polytransactions: A Mechanism for Management of Interdependent Data. *book chapter in Transaction Models for Advanced Database Applications, Ed. A. Elmagarmid, Morgan-Kaufmann*.
- Skarra A. (1991) Localized Correctness Specifications for Cooperating Transactions in an Object-Oriented Database. *IEEE Bulletin on Office and Knowledge Engineering*, 4(1):79–106.
- Yannakakis M. (1984) Serializability by Locking. *Journal of the ACM*, 31(2):227–244.