A Family of Speculative Concurrency Control Algorithms for Real-Time Databases

Azer Bestavros
(best@cs.bu.edu)

Spyridon Braoudakis
(sb@cs.bu.edu)

Computer Science Department
Boston University
Boston, MA 02215

November 30, 1992

Abstract

Speculative Concurrency Control (SCC) was proposed in [Best92b] as a new concurrency control approach especially suited for real-time database applications. SCC relies on the use of redundancy to ensure that serializable schedules are discovered and adopted as early as possible, thus increasing the likelihood of the timely commitment of transactions with strict timing constraints. Using SCC, several *shadow* transactions execute on behalf of a given uncommitted transaction so as to protect against the hazards of *blockages* and *restarts*, which are characteristics of Pessimistic and Optimistic Concurrency Control algorithms, respectively.

We present SCC-nS, a generic algorithm that characterizes a family of SCC-based algorithms. Under the SCC-nS algorithm, shadows executing on behalf of a transaction are either *optimistic* or *speculative*. Optimistic shadows execute under an assumed serialization order, which requires them to wait for *no* other conflicting transactions. They execute unhindered until they are either aborted or committed. Speculative shadows are more conservative. They execute under an assumed serialization order, which requires them to wait for *some* conflicting transactions to commit. In this paper, we provide a description of the SCC-nS algorithm, establish its correctness by showing that it only admits serializable histories, and demonstrate its superiority for RTDBMS through numerous examples. Three members of the SCC-nS family (namely SCC-1S, SCC-2S and SCC-MS) are singled out and contrasted. SCC-1S and SCC-MS represent two extremes in a spectrum of choices regarding the total amount of spared redundancy in the system. SCC-1S is notable for its minimal use of redundancy, whereas SCC-MS is notable for its liberal use thereof.

1 Introduction

Traditional concurrency control algorithms can be broadly classified as either *pessimistic* or *optimistic*. Pessimistic Concurrency Control (PCC) algorithms [Eswa76, Gray76] avoid any concurrent execution of transactions as soon as *potential* conflicts between these transactions are detected. On the contrary, Optimistic Concurrency Control (OCC) algorithms [Boks87, Kung81] allow such transactions to proceed at the risk of having to restart them in case these suspected conflicts *materialize*.

For real-time database applications where transactions execute under strict timing constraints, maximum concurrency (or throughput) ceases to be an expressive measure of performance. Rather, the number of transactions completed before their set deadlines becomes the decisive performance measure [Buch89]. Most real-time concurrency control schemes considered in the literature [Abbo88, Agra87, Stan88, Huan89, Sing88, Sha88, Sha91] are based on Two-Phase Locking (2PL), which is a PCC strategy. Despite its widespread use in commercial systems, 2PL has some properties such as the possibility of deadlocks and long and unpredictable blocking times that damage its appeal for real-time environments, where the primary performance criterion is meeting time constraints and not just preserving consistency requirements. Over the last few years, several alternatives to 2PL for RTDBMS have been explored, proposed, and investigated [Kort90, Hari90b, Hari90a, Huan91, Kim91, Lin90, Son92].

In a recent study [Best92b], Bestavros proposed a categorically different approach to concurrency control for RTDBMS. His approach relies on the use of redundant computation to start on alternative schedules, once conflicts that threaten the consistency of the database are detected. These alternative schedules are adopted *only if* the suspected inconsistencies materialize; otherwise, they are abandoned. Due to its nature, this approach has been termed *Speculative Concurrency Control* (SCC). This paper examines a family of SCC algorithms and their implementations.

SCC algorithms use redundancy to combine the advantages of both PCC and OCC algorithms, while avoiding their disadvantages. On the one hand, SCC resembles PCC in that potentially harmful conflicts are detected as early as possible, allowing a head-start for alternative schedules, and thus increasing the chances of meeting the set timing constraints – should these alternative schedules be needed (due to restart as in OCC). On the other hand, SCC resembles OCC in that it allows conflicting transactions to proceed concurrently, thus avoiding unnecessary delays (due to blocking as in PCC) that may jeopardize their timely commitment.

Because of their reliance on redundant computations, SCC algorithms require the availability of enough capacity in the system. While abundant resources are usually not to be expected in conventional database systems, they may be more common in real-time database system environments [Fran85]. Throughout this paper, we make the assumption that an abundance of computing resources is, indeed, available. The SCC algorithms that we propose in this paper represent a host of choices in terms of the required amount of redundant computations. We argue that these algorithms are superior to any existing real-time concurrency control algorithms, even in the absence of any *spare* computing resources.

The remainder of this paper is organized as follows. In section 2, we review some of the problems encountered with traditional concurrency control in RTDBMS, and we overview the basic idea behind the SCC-based approach. In section 3, SCC-nS, a generic SCC algorithm is described, its correctness is established and its superiority for RTDBMS is demonstrated through numerous examples. In section 4, three members of the SCC-nS family (namely SCC-1S, SCC-2S, and SCC-MS) are singled out and contrasted. In section 5, we conclude with a description of our current and future research work.

2 Concurrency Control for RTDBMS

A disadvantage of classical OCC when used in RTDBMS is that transaction conflicts are not detected until the validation phase, at which time it might be too late to restart. This may have a negative impact on the number of timing constraint violations. PCC two-phase locking algorithms do not suffer from this problem because they detect potential conflicts as they occur.

The Broadcast Commit variant (OCC-BC) [Mena82, Robi82] of the classical OCC remedies this problem partially. When a transaction commits, it notifies all concurrently running, conflicting transactions about its commitment. All those conflicting transactions are immediately restarted. The broadcast commit method detects conflicts earlier than the basic OCC algorithm resulting in less wasted resources and earlier restarts.

To illustrate this point, consider the following example. Assume that we have two transactions T_1 and T_2, which (among others) perform some conflicting actions. In particular, T_2 reads item x after T_1 has updated it. Adopting the basic OCC algorithm means restarting transaction T_2 when it enters its validation phase because it conflicts with the already committed transaction T_1 on data item x. This scenario is illustrated in figure 1. Obviously, the likelihood of the restarted transaction T_2 meeting its timing constraint decreases considerably.

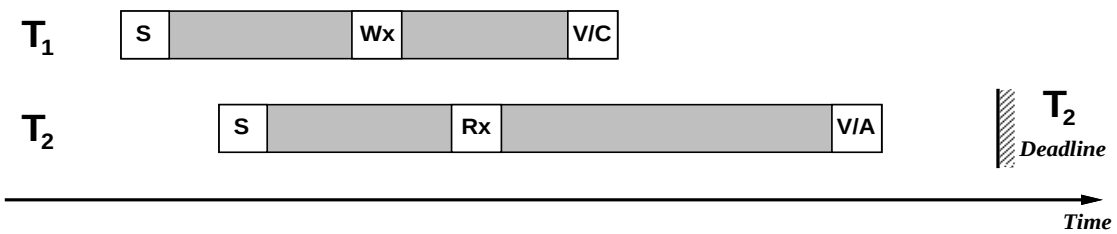

Figure 1: Transaction management under the basic OCC algorithm.

The OCC-BC algorithm avoids waiting unnecessarily for a transaction's validation phase in order to restart it. In particular, a transaction is aborted if any of its conflicts with other transactions in the system becomes a materialized consistency threat. This is illustrated in figure 2.

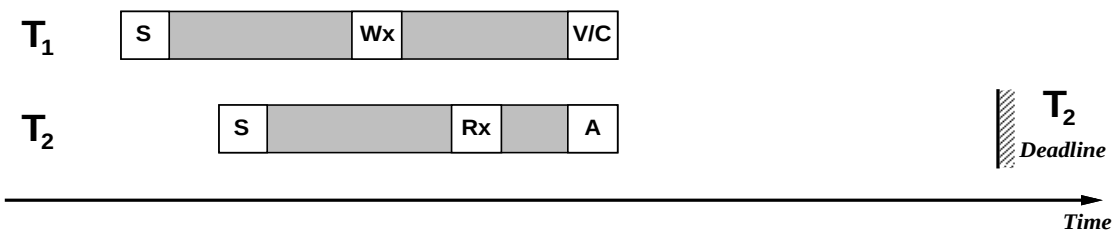

Figure 2: Transaction management under the OCC-BC algorithm.

The SCC-based Approach:

The SCC approach proposed in [Best92b] goes one step further in utilizing information about conflicts. Instead of waiting for a potential consistency threat to materialize and *then* taking a corrective measure, an SCC algorithm uses redundant resources to start on *speculative* corrective measures as soon as the conflict in question develops. By starting on such corrective measures as early as possible, the likelihood of meeting any set timing constraints will be greatly enhanced. Figure 3 and figure 4 show two possible scenarios that may develop depending on the time needed for transaction T_2 to reach its validation phase. Each one of these scenarios corresponds to a different serialization order.

In figure 3, T_2 reaches its validation phase before T_1. T_2 will be validated[1] and committed without any need to disturb T_1. Therefore, this schedule will be serializable with transaction T_2 preceding transaction T_1. Obviously, once T_2 commits, the shadow transaction T_2' has to be aborted.

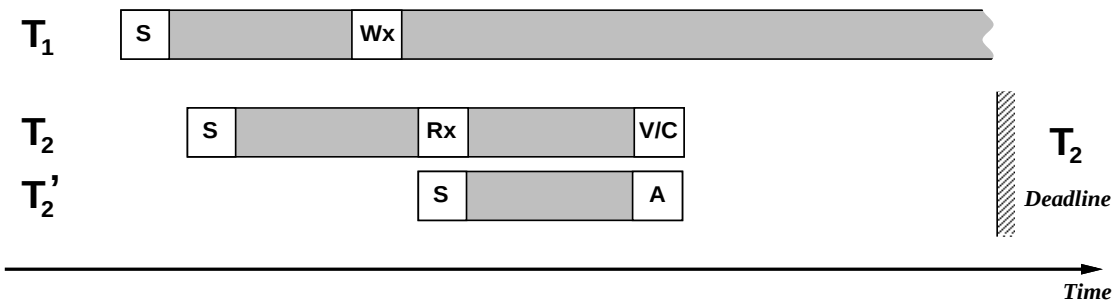

Figure 3: Schedule with an undeveloped potential conflict.

However, if transaction T_1 reaches its validation phase first, then transaction T_2 cannot continue to execute due to the (now visible) conflict over x; T_2 must abort. With OCC-BC algorithms, T_2 would have had to restart when T_1 commits. This might be too late if T_2's deadline is close. The SCC protocol, instead of restarting T_2, simply aborts T_2 and adopt its shadow transaction T_2'. This scenario is illustrated in figure 4.

[1] since T_2's write-set does not intersect T_1's read-set (assuming that there are no conflicting actions other than the reading and writing of x).

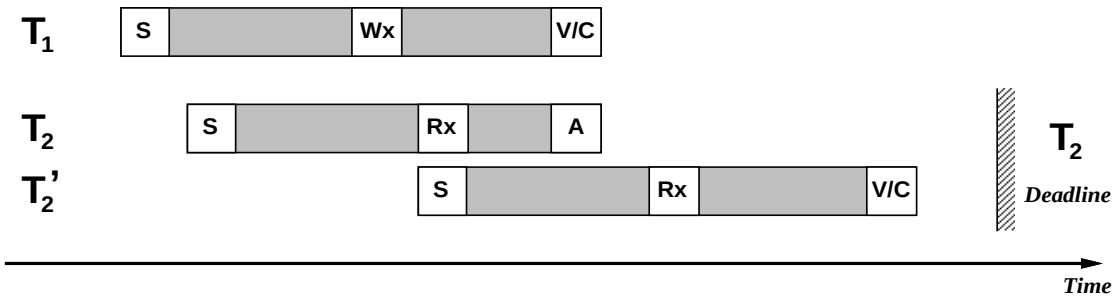

Figure 4: Schedule with a developed conflict.

3 The Generic SCC-nS Algorithm

The most general SCC algorithm requires a tremendous amount of redundancy to account for *every* possible speculated serialization order of the uncommitted transactions in the system. In this section, we present a class of SCC algorithms that operate under a *limited redundancy* assumption. In particular, we present a generic SCC algorithm which does not allow more than n shadows to execute on behalf of any given uncommitted transaction in the system.

3.1 Preliminaries

A transaction T_i consists of a sequence of actions $a_{i1}, a_{i2}, \ldots a_{im}$, where each a_{ij}, $j = 1, 2, \ldots m$, is either a *read* or a *write* operation on one of the shared objects of the database. Each transaction in the system is assumed to preserve the consistency of these shared objects. Therefore, *any* sequential (or serializable) execution of any collection of transactions will also preserve the consistency of the database [Papa79, Bern87].

Given a concurrent execution of transactions, action a_{ir} of transaction T_i conflicts with action a_{js} of transaction T_j, if they access the same object *and* either a_{ir} is a read operation and a_{js} is a write operation (*read-write* conflict), or a_{ir} is a write operation and a_{js} is a read operation (*write-read* conflict).

Write-write conflicts (when both a_{ir} and a_{js} actions are write operations) are treated using the Thomas' Write Rule (TWR). At validation, when all database updates are made permanent, all write requests are buffered by the data manager and serialized according to their transaction order. A timestamp can be assigned to every committing transaction for that purpose. With the TWR every write request arriving out of order (late) is being *ignored* rather than being *rejected* [Bern87]. In other words, all write requests are granted, whether or not the targeted data object is being updated by another uncommitted transaction.

As we have hinted before, SCC-based algorithms allow several shadows (processes or tasks) to execute concurrently on behalf of the same transaction. Each one of these processes corresponds to a different *speculated serialization order*. For a transaction T_r, each one of these processes is called

a *shadow* of T_r. In this paper, a shadow can be in one of two modes: *optimistic* or *speculative*. Each transaction T_r has, at any point in its execution, exactly one optimistic shadow T_r^o. In addition, T_r may have i speculative shadows T_r^i, for $i = 0, \ldots, n-1$. Accordingly, each transaction can have *at most* n shadows executing on its behalf (possibly concurrently) at any point in its lifetime.

For each transaction T_r we keep a variable $SpecNumber(T_r)$, which counts the number of the speculative shadows currently executing on behalf of T_r. With each shadow T_r^i of a transaction T_r – whether optimistic, or speculative – we maintain a set $ReadSet(T_r^i)$, which records pairs (X, t_x), where X is an object read by T_r^i, and t_x represents the order[2] in which this operation was performed. We use the notation: $(X, _) \in ReadSet(T_r^i)$ to mean that shadow T_r^i read object X.

For each speculative shadow T_r^i in the system, we maintain a set $WaitFor(T_r^i)$, which contains pairs of the form (T_u, X), where T_u is an uncommitted transaction and X is an object of the shared database. $(T_u, X) \in WaitFor(T_r^i)$ implies that T_r^i must wait for T_u before being allowed to Read object X. We use $(T_u, _) \in WaitFor(T_r^i)$ to denote the existence of at least one tuple (T_u, X) in $WaitFor(T_r^i)$, for some object X.

3.2 Algorithm Overview

Under the SCC-nS algorithm, shadows executing on behalf of a transaction are either *optimistic* or *speculative*. Optimistic shadows execute unhindered, whereas speculative shadows are maintained so as to be ready to replace a defunct optimistic shadow, if such a replacement is deemed necessary.

Optimistic shadow behavior:

For a transaction T_r, the optimistic shadow T_r^o executes with the assumption that it will commit *before* all the other uncommitted transactions in the system with which it conflicts. Obviously, this is an *optimistic assumption*. T_r^o records any conflicts found during its execution, and proceeds uninterrupted until one of these conflicts materializes (due to the commitment of a competing transaction), in which case T_r^o is aborted – or else until its validation phase is reached, in which case T_r^o is committed.[3]

Speculative shadow behavior:

Each speculative shadow T_r^s executes with the assumption that it will finish before the materialization of any detected conflict with any other uncommitted transaction, except for one particular conflict which is *speculated* to materialize before the commitment of T_r. Thus, T_r^s remains blocked on the shared object X, on which this conflict has developed, waiting to read the value that the conflicting transaction, T_u will assign to X when it commits. If this speculated assumption becomes true, (*e.g.* T_u commits before T_r enters its validation phase), T_r^s will be unblocked and *promoted* to become T_r's optimistic shadow, replacing the old optimistic shadow which will have to be aborted, since it made the wrong assumption with respect to the serialization order.

[2]This can be a special read timestamp, implemented by maintaining for each shadow T_r^i in the system a counter that is atomically incremented every time a read operation is performed by T_r^i.

[3]Only the optimistic shadow of a transaction can be committed.

At any point during the execution of our algorithm, the first k speculative shadows of a transaction T_r account for the first k detected conflicts in which T_r participated. These may not be the first k conflicts that transaction T_r will develop during the course of its execution. To illustrate this point, consider the condition depicted in figure 5. Transaction T_1 may detect at some point in its execution a conflict over some object X, which it had read earlier. In particular, when the read operation for object X was requested by the optimistic shadow T_1^o, there was no conflict to be detected. Such a conflict appeared later when transaction T_3 requested to update that same object X.

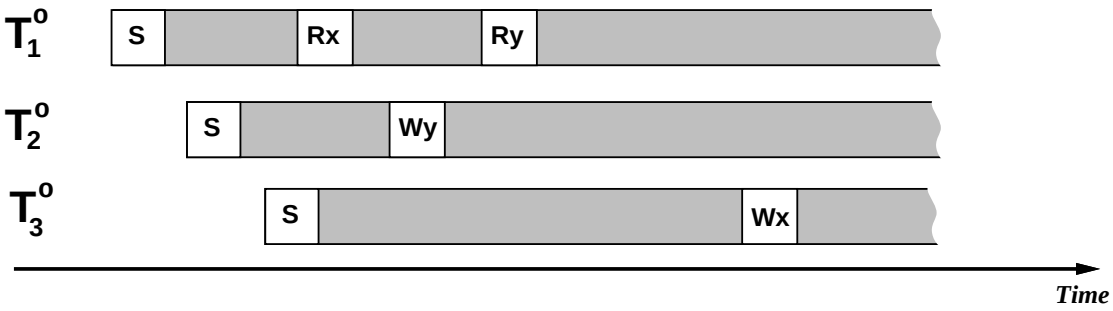

Figure 5: T_1 detects the (T_3, X) conflict only after T_3 writes object X.

The *shadow replacement algorithm* we are using in this paper is one of several algorithms that could be adopted. In [Best92c] some alternatives to this policy are discussed and evaluated. In particular, information about deadlines and priorities of the conflicting transactions can be utilized so as to account for the *most probable* serialization orders.

It is very important to realize that the imposed limit of at most $n-1$ speculative shadows per transaction does not prohibit a transaction T_r from developing more than $n-1$ conflicts at any point during its lifetime. Rather, this limit is on the number of potential hazards that our algorithm will be ready to *optimally* deal with (by using the speculative shadows). Every *extra* hazard that develops after this limit is reached will be accounted for only *suboptimally*[4] (since no such speculative shadow will be available). In that sense, we can view the aforementioned description as encompassing a hierarchy of algorithms. Going down a level in this hierarchy (by reducing n) can compromise only performance not correctness.

[4]We can still use the presense of other speculative shadows to improve those decisions (see the Commit Rule below).

3.3 Description of the SCC-nS Algorithm

Let $\mathcal{T} = T_1, T_2, T_3, \ldots, T_m$ be the set of uncommitted transactions in the system. Furthermore, let $\mathcal{T}^{\mathcal{O}}$, and $\mathcal{T}^{\mathcal{S}}$ be, respectively the sets of optimistic, and speculative shadows executing on behalf of the transactions in the set $\mathcal{T}$. We use the notation $\mathcal{T}_r^{\mathcal{S}}$ to denote the set of speculative shadows executing on behalf of transaction T_r. The SCC-nS algorithm is described as a set of five rules, which we describe below.

Start Rule:

The *Start Rule*, is followed whenever a new transaction T_r is submitted for execution, in which case an optimistic shadow T_r^o is created. In the absence of any conflicts this shadow will run to completion (the same way as with the OPT-BC algorithm).

Read Rule:

The *Read Rule* is activated whenever a read-after-write conflict is detected. The processing that follows is straightforward. In particular, if the maximum number of speculative shadows of the transaction in question, say T_r, is not exhausted, a new speculative shadow T_r^s is created (by forking it off T_r^o) to account for the newly detected conflict. Otherwise, in the absence of any new speculative shadow for transaction T_r, this potential conflict will have to be ignored at this point. The Commit Rule (see below) deals with the corrective measures that need to be taken, should this conflict materializes.

Write Rule:

The *Write Rule* is activated whenever a write-after-read conflict is detected. A number of issues must be dealt with; these are discussed below.

Speculative shadows cannot be forked off as before from the transaction's optimistic shadow. This is because the conflict is detected on some other transaction's write operation. Therefore, since its optimistic shadow already read that database object, we must either create a new copy of this transaction or choose another point during its execution from which we can fork it off. For performance reasons, this second choice was adopted. The algorithm makes use of the function *BestShadow* (discussed later) to find the most appropriate speculative shadow, if such a shadow indeed exists. In the absence of such a shadow a restarted copy of the transaction is created. Figure 6 illustrates this point. When the new conflict (T_2, X) is detected, the speculative shadow T_1^3 is forked off T_1^1 to accommodate it. Notice that if a copy of T_1 was instead created, all the operations before R_y (reading the database object Y) would have had to be repeated. T_1^2, even though in a later stage, is not an appropriate shadow to fork off because, like the optimistic shadow, it already read object X.

When the new conflict implicates transactions that already conflict with each other, some adjustments may be necessary. In figure 7, the speculative shadow T_1^j of transaction T_1, accounting for the conflict (T_2, Z), must be aborted as soon as the new conflict, (T_2, X), involving the same two transactions is detected. Since T_1 read object X before object Z, (T_2, X) is the *first* conflict

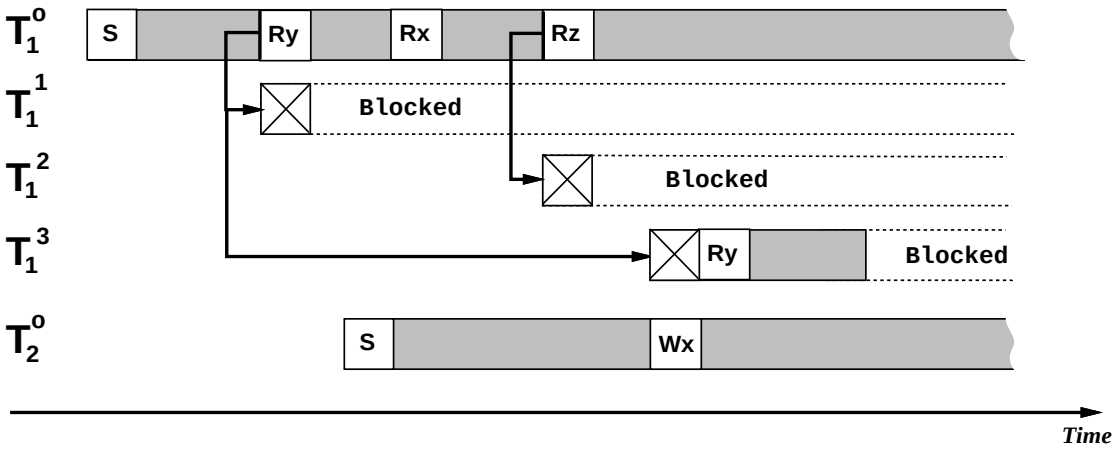

Figure 6: (Write Rule 1.1.1) T_1^3 is forked off the *BestShadow* (T_1, X), T_1^1.

between those two transactions. Therefore, the speculative shadow accounting for the possibility that transaction T_2 will commit before transaction T_1 must block before the read operation on X is performed. Speculative shadow T_1^k is forked off T_1^1 for that purpose. All other speculative shadows of T_1 remain unaffected.

The number of speculative shadows maintained by SCC-nS (namely $n-1$) might not be enough to account for all the conflicts that develop during a transaction's lifetime. The selection of the conflicts to be accounted for by speculative shadows is an interesting problem with many possible solutions [Best92c]. In this paper we have adopted a particular solution that requires the speculative shadows of SCC-nS to account for the *first* $k \leq n-1$ conflicts (whether read-after-write or write-after-read) encountered by a transaction. Because such conflicts are not necessarily detected *in order*, a *shadow replacement* might be necessary. To illustrate this point, consider the scenario depicted in figure 8, where the assumption that the first two conflicts in which transaction T_1 participated (by accessing objects Y, and Z, respectively), is revised when transaction T_2 writes object X. In particular, the newly detected conflict (T_2, X) becomes the first conflict of T_1. If it is the case that T_1 is restricted so as not to have more that two speculative shadows at any point during its execution, then a shadow replacement is necessary. T_1^2, the *latest* shadow of T_1 has to be aborted, and a new speculative shadow, T_1^3, accounting for the new (T_2, X) conflict should replace it. The *LastShadow* function (explained below) is used to find this *latest* speculative shadow.

Blocking Rule:

The *Blocking Rule* is used to control when a speculative shadow T_r^i must be blocked. This rule assures that T_r^i is blocked the *first* time it wishes to read an object X in conflict with any transaction that T_r^i must wait for according to its speculated serialization order.

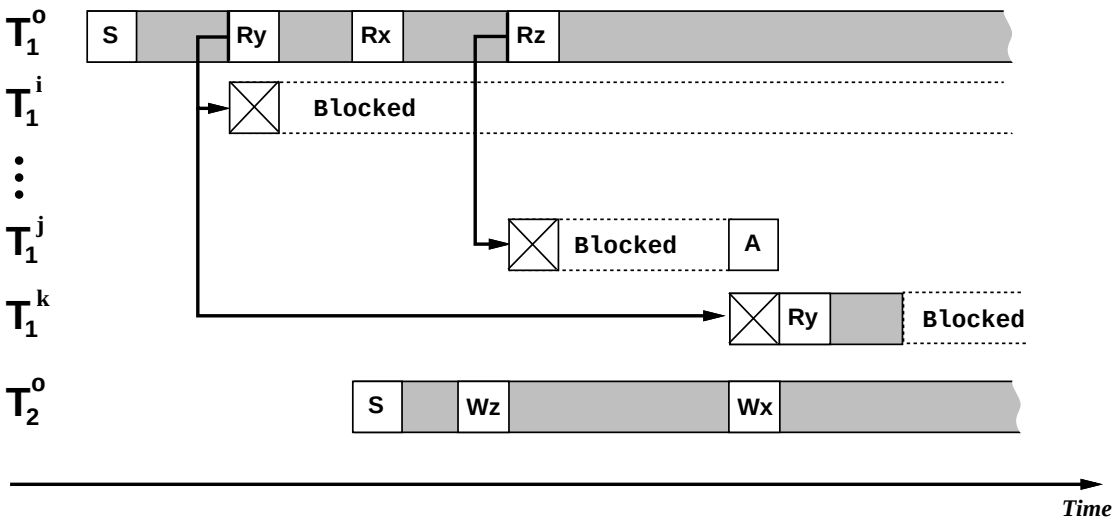

Figure 7: (Write Rule 1.2) T_1^j, which accounts for the (T_2, Z) conflict, is aborted and replaced by T_1^k when an *earlier* conflict, (T_2, X), with T_2 is detected.

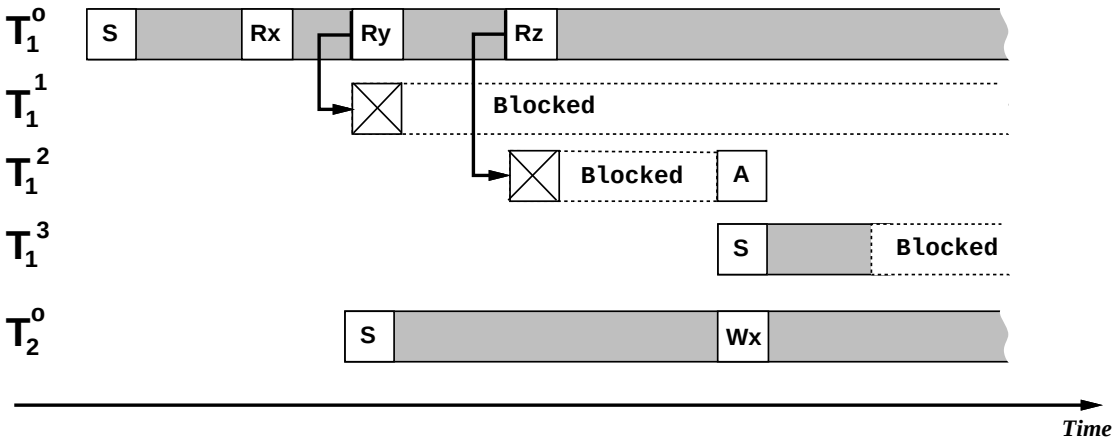

Figure 8: (Write Rule 2.1) The detection of the (T_2, X) conflict causes the abortion of the *LastShadow*(T_1), namely T_1^2, and its replacement by T_1^3.

Commit Rule:

Whenever it is decided to commit an optimistic shadow[5] T_r^o on behalf of a transaction T_r, the *Commit Rule* is activated. First, all other shadows of T_r become obsolete and are aborted. Next, all transactions conflicting with T_r must be dealt with. For each conflicting transaction T_u there are two cases that may occur: either there is a speculative shadow, T_u^i, waiting for T_r's commitment, or not.

The first case is illustrated in figure 9, where the speculative shadow T_1^2 of transaction T_1 – having anticipated (assumed) the correct serialization order – is promoted to become the new optimistic shadow of transaction T_1, replacing the old optimistic shadow which had to be aborted. Speculative shadow T_1^3, which like the old optimistic shadow exposed itself by reading the old value of object X had to be aborted as well. Alternatively, the speculative shadow T_1^1, which did not read object X, remains unhindered.

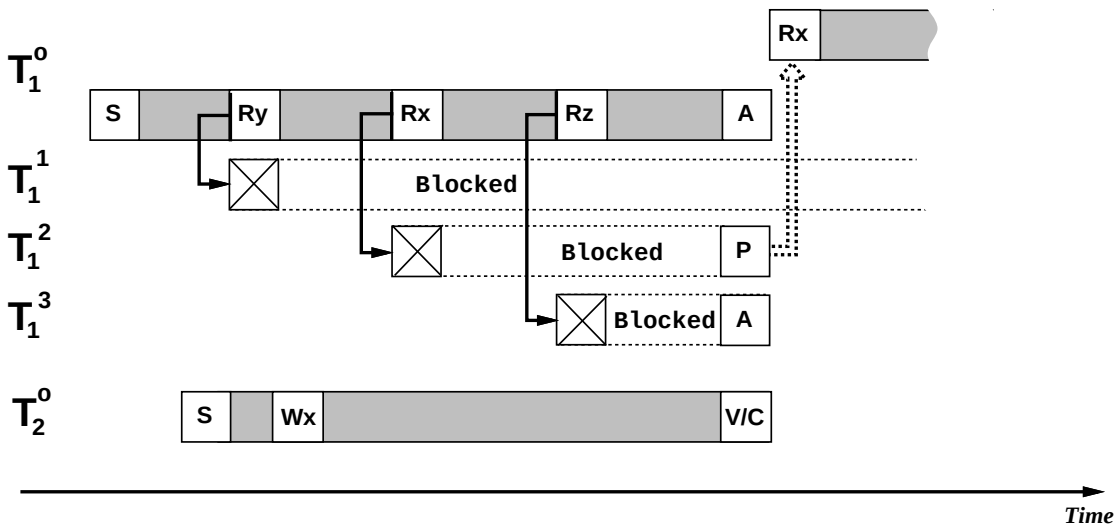

Figure 9: (Commit Rule 2.) T_1^2, accounting for the developed conflict (T_2, X), is promoted to replace the optimistic shadow of T_1. T_1^3 is aborted because it read object X, while T_1^1 remains unaffected.

The second case is illustrated in figure 10, where the commitment of the optimistic shadow T_2^o on behalf of transaction T_2 was not accounted for by any speculative shadow.[6] In this case, a shadow is forked off the *LastShadow*(T_1) to become the new optimistic shadow of transaction T_1. This, even though not optimal, is the best we can do in the absence of a speculative shadow accounting for the (T_2, Z) conflict. A complete and formal description of the SCC-nS algorithm can be found in figure 11.

[5] Recall that only optimistic shadows are allowed to commit.

[6] Figure 10 makes the implicit assumption that transaction T_1 is limited to having at most two speculative shadows

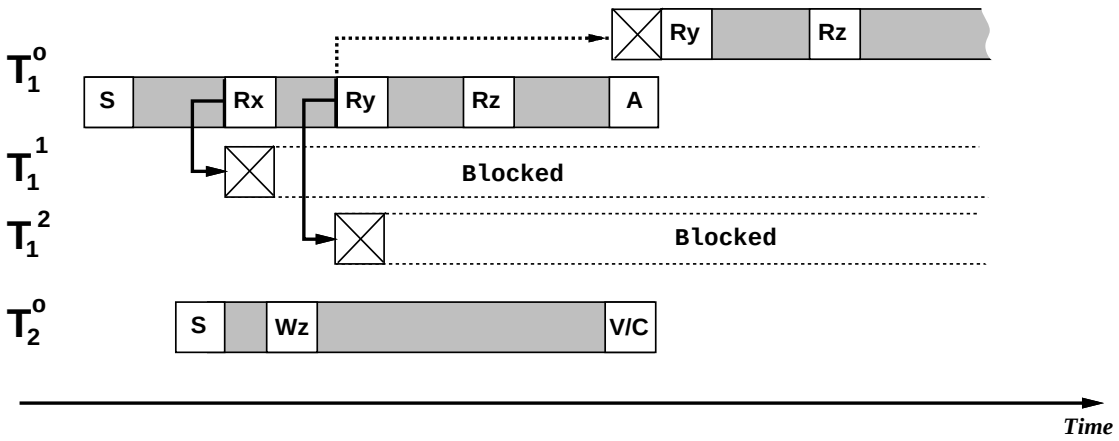

Figure 10: (Commit Rule 3.) When the unaccounted conflict (T_2, Z) materializes, a new optimistic shadow for T_1 is forked off the *LastShadow*(T_1), T_1^2.

As we mentioned above, the algorithm makes use of two functions: *LastShadow*, and *BestShadow*. *LastShadow* is a function from the set of uncommitted transactions $\mathcal{T}$ to the set of speculative shadows $\mathcal{T}^{\mathcal{S}}$. It takes for input a transaction T_r, and returns the *latest* speculative shadow T_r^{last} of T_r in order of read conflict. *BestShadow* is a function from the cross-product of uncommitted transactions and database objects, to the set of speculative shadows $\mathcal{T}^{\mathcal{S}}$. It takes as input a transaction T_r and a database object X read by its optimistic shadow T_r^o. It returns the speculative shadow T_r^{best} of T_r, which did not read object X and accounts for the *latest* conflict (T_u, Y) in which T_r participates. Should such a speculative shadow does not exist, T_r^{best} corresponds to the starting point in the execution of T_r. Figure 12 provides a formal definition of these functions.

3.4 Correctness of the SCC-nS Algorithm

Having described its basic concepts, we now present an informal proof of correctness for our algorithm. First, we define the notions of *history* and *serialization graph* (SG) introduced in [Bern87]. A *history* H is a partial order of operations that represents the execution of a set of transactions $\mathcal{T}$. Any two conflicting operations in H must be comparable. The *serialization graph* for history H, denoted by SG(H), is a directed graph whose nodes, $T_i \in \mathcal{T}$, are committed transactions in H. Its edges are all $T_i \leftarrow T_j$, for $i \neq j$, such that one of T_i's operations precedes and conflicts with one of T_j's operations in H. A history H is serializable if its serialization graph SG(H) is acyclic [Bern87].

---

at any point during its execution.

Theorem: *Every history H produced by the SCC-nS algorithm is serializable.*

Proof: Let T_1 and T_2 be two committed transactions in a history H produced by the SCC-nS algorithm. We argue that if there is an edge $T_1 \leftarrow T_2$ in SG(H), then the commitment of transaction T_1 precedes that of transaction T_2, denoted by $T_1 \prec T_2$, in the serialization order. Because of the edge $T_1 \leftarrow T_2$ the two transactions must have some conflicting operation. Three cases must then be examined:

• T_1 writes an object X that is read by T_2:

By the Read Rule, if T_2 requests to read X after it has been written by T_1, a speculative shadow[7] T_2^s is forked off the optimistic shadow T_2^o, which accounts for transaction T_1 committing before transaction T_2. The conflict (T_1, X) is then appended to the *WaitFor*(T_2^s) set. Alternatively, if the read request of T_2 on object X comes before the write request on X by T_1, a speculative shadow T_2^s is forked off the *BestShadow*(T_2), and (T_1, X) is appended in the *WaitFor*(T_2^s) set by the Write Rule. In either case, T_2^s will not reach its validation phase before T_1, because *WaitFor*(T_2^s) cannot be empty while T_1 is still in progress. No other shadow of T_2 can commit until T_1 is committed. Otherwise, if shadow T_2^i commits before T_1, the potential conflict on object X will not develop (see figure 3), invalidating the edge $T_1 \leftarrow T_2$ in the SG(H). Hence, T_1 must commit before T_2, and therefore $T_1 \prec T_2$ in the serialization order.

• T_1 reads an object X that is written by T_2:

T_1 must commit before T_2, because the edge $T_1 \leftarrow T_2$ in the SG(H) necessitates that the potential conflict on object X did not develop. Therefore, again $T_1 \prec T_2$ in the serialization order.

• T_1 and T_2 write object X:

Suppose $T_2 \prec T_1$. Then T_2 enters its validation phase before T_1. T_2's write operation on X is sent to the data manager first. It will either be processed before T_1's write operation on X, or it will be discarded when the data manager receives T_1's write operation on X (TWR). Therefore, T_1's write operation on X is never processed before that of T_2's. Then the conflict – implying the edge $T_1 \leftarrow T_2$ in the SG(H) – is impossible; a contradiction.

The proof of the theorem is by contradiction. In particular, suppose that there is a cycle $T_1 \leftarrow T_2 \leftarrow \cdots \leftarrow T_n \leftarrow T_1$ in SG(H). By the above argument, it must then be the case that $T_1 \prec T_2 \prec \cdots \prec T_n \prec T_1$, which leads to a contradiction. Therefore no cycle can exist in SG(H) and thus the SCC-nS algorithm only produces serializable histories.□

3.5 An Example

In this section, we will demonstrate (using a simple example) the distinctive nature of our algorithm, which makes it especially suited for RTDBMS. More precisely, given a sequence of transaction requests we will show that the SCC-nS algorithm produces a history that is distinct from those produced by either 2PL or OCC-BC. The example is depicted in figure 13.

[7] If the maximum allowed number of speculative shadows for T_1 is not exhausted. Otherwise, the development of this conflict will be addressed by the Write Rule.

A. The Start Rule: When the execution of a new transaction T_r is requested, an optimistic shadow $T_r^o \in \mathcal{T}^{\mathcal{O}}$ is created and executed.

B. The Read Rule: Whenever an optimistic shadow T_r^o wishes to read an object X, then:

for all T_u^o in $\mathcal{T}^{\mathcal{O}}$, such that T_u^o wrote object X **do**
1. **if** $((SpecNumber(T_r) < n-1) \wedge (\forall T_r^i \in \mathcal{T}_r^{\mathcal{S}}, (T_u, _) \notin WaitFor(T_r^i)))$ **then**{
1.1 A new speculative shadow T_r^j is forked off T_r^o;
1.2 $WaitFor(T_r^j) \leftarrow \{(T_u, X)\}$;
1.3 $SpecNumber(T_r) \leftarrow SpecNumber(T_r) + 1$};

C. The Write Rule: Whenever an optimistic shadow T_u^o wishes to write an object X, then:

for all T_r^o in $\mathcal{T}^{\mathcal{O}}$, such that T_r^o read object X **do**
1. **if** $(SpecNumber(T_r) < n-1)$ **then**{
1.1 **if** $(\forall T_r^i \in \mathcal{T}_r^{\mathcal{S}}, (T_u, _) \notin WaitFor(T_r^i))$ **then**{
1.1.1 A new speculative shadow T_r^j is forked off $BestShadow(T_r, X)$;
1.1.2 $WaitFor(T_r^j) \leftarrow \{(T_u, X)\}$;
1.1.3 $SpecNumber(T_r) \leftarrow SpecNumber(T_r) + 1$
1.2 }**else if** $(\exists T_r^k \in \mathcal{T}_r^{\mathcal{S}}, \exists Y : ((X, _) \in ReadSet(T_r^k) \wedge (T_u, Y) \in WaitFor(T_r^k)))$ **then**{
1.2.1 T_r^k is aborted and replaced by T_r^m which is forked off $BestShadow(T_r, X)$;
1.2.2 $WaitFor(T_r^m) \leftarrow \{(T_u, X)\}$}};
2. }**else if** $(SpecNumber(T_r) = n-1)$ **then**
2.1 **if** $(\exists T_r^k \in \mathcal{T}_r^{\mathcal{S}} : (X, _) \in ReadSet(T_r^k))$ **then**
2.1.1 Abort $LastShadow(T_r)$;
2.1.2 A new speculative shadow T_r^m is forked off $BestShadow(T_r, X)$;
2.1.3 $WaitFor(T_r^m) \leftarrow \{(T_u, X)\}$};

D. The Blocking Rule: A standby shadow T_r^i is blocked at the *earliest point* at which it wishes to Read an object X that is written by any transaction T_u, such that $(T_u, X) \in WaitFor(T_r^i)$.

E. The Commit Rule: Whenever it is decided to commit an optimistic shadow T_r^o on behalf of a transaction T_r, then:

1. $\forall T_r^i \in \mathcal{T}_r^{\mathcal{S}}, T_r^i$ is aborted;
2. **for** all $T_u \in \mathcal{T}$, such that $(\exists T_u^i \in \mathcal{T}_u^{\mathcal{S}} : (T_r, X) \in WaitFor(T_u^i))$ **do**{
2.1 T_u^o is aborted;
2.2 T_u^i is promoted to become the new optimistic shadow of T_u;
2.3 $SpecNumber(T_u) \leftarrow SpecNumber(T_u) - 1$;
2.4 **for** all $T_u^j \in \mathcal{T}_u^{\mathcal{S}}$, such that $(X, _) \in ReadSet(T_u^j)$ **do**{
2.4.1 T_u^j is aborted;
2.4.2 $SpecNumber(T_u) \leftarrow SpecNumber(T_u) - 1$ }};
3. **for** all $T_u \in \mathcal{T}$, such that $(\exists X : T_r^o$ wrote object $X \wedge$
$(\nexists T_u^i \in \mathcal{T}_u^{\mathcal{S}} : (T_r, X) \in WaitFor(T_u^i)))$ **do**{
3.1 T_u^o is aborted;
3.2 A new optimistic shadow T_u^o is forked off $LastShadow(T_u)$};

Figure 11: The Generic SCC-nS Algorithm.

(a) $\underline{LastShadow()} : \mathcal{T} \rightarrow \mathcal{T}^{\mathcal{S}}$, such that $T_r \in \mathcal{T} \longmapsto T_r^{last} \in \mathcal{T}^{\mathcal{S}}$ **iff**
$(\exists X : (X, t_x) \in ReadSet(T_r^o)) \wedge ((\exists T_u \in \mathcal{T} : (T_u, X) \in WaitFor(T_r^{last})) \wedge (\forall Y : ((Y, t_y) \in ReadSet(T_r^o) \wedge (\exists T_v \in \mathcal{T}, \exists T_r^i \in \mathcal{T}_r^{\mathcal{S}} : (T_v, Y) \in WaitFor(T_r^i)))) \Longrightarrow t_y \leq t_x)$.

(b) $\underline{BestShadow()} : (\mathcal{T}, object) \rightarrow \mathcal{T}^{\mathcal{S}}$, such that $(T_r, X) \in (\mathcal{T}, Object) \longmapsto T_r^{best} \in \mathcal{T}^{\mathcal{S}}$ **iff**
$(X, t_x) \in ReadSet(T_r^o) \wedge (X, t_x) \notin ReadSet(T_r^{best}) \wedge (\exists T_u \in \mathcal{T}, \exists Y : ((Y, t_y) \in ReadSet(T_r^o) \wedge (T_u, Y) \in WaitFor(T_r^{best}))) \wedge (\forall Z : ((Z, t_z) \in ReadSet(T_r^o) \wedge (\exists T_v \in \mathcal{T}, \exists T_r^i \in \mathcal{T}_r^{\mathcal{S}} : ((T_v, Z) \in WaitFor(T_r^i) \wedge (X, t_x) \notin ReadSet(T_r^i)))) \Longrightarrow t_z \leq t_y)$.

Figure 12: Definitions of the *LastShadow*, and *BestShadow* functions.

There are three transactions in this example: T_1, T_2, and T_3, arriving in that order. Among others the transactions request the following operations in that order: R_x^2, W_y^3, W_x^1, and R_y^2 – with the superscript denoting the transaction's number. Two speculative shadows T_2^1, and T_2^2, on behalf of transaction T_2 are generated to account for the two conflicts (T_1, X), and (T_2, Y), in which T_2 participates. Transaction T_1 is committed first. At that time the speculative shadow T_2^1 is promoted to become the new optimistic shadow of transaction T_2. It is executed to completion and committed before its deadline. Finally, transaction T_3 is validated and committed. Therefore, the serialization order produced by the SCC-nS algorithm is: $T_1 \prec T_2 \prec T_3$.

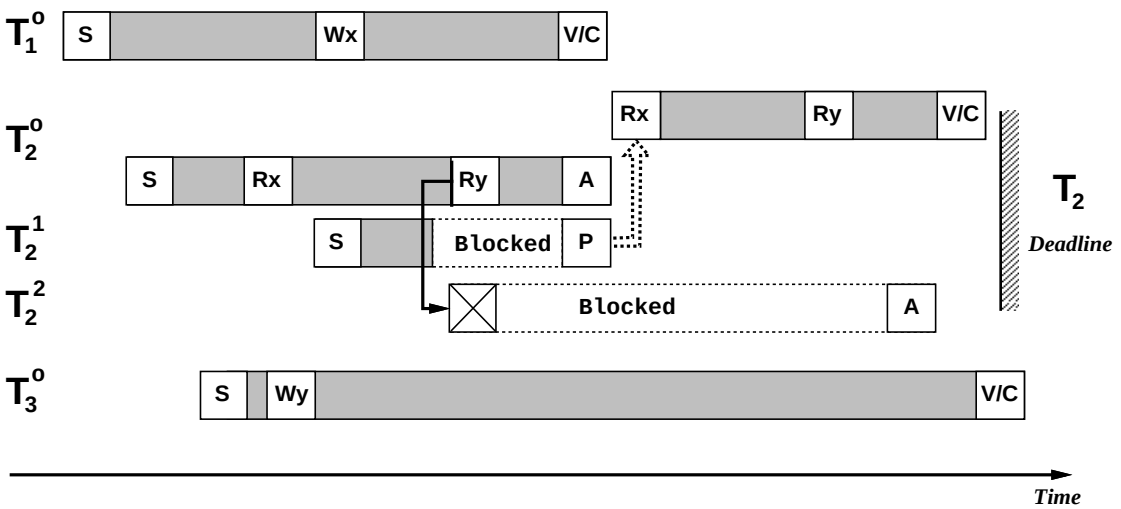

Figure 13: The serialization order, $T_1 \prec T_2 \prec T_3$, cannot be produced by either 2PL, or OCC-BC.

The same sequence of transaction requests using 2PL would have resulted in transaction T_3 finishing first, followed by transaction T_2, and transaction T_1 being the last to finish. This results in the serialization order: $T_3 \prec T_2 \prec T_1$. T_2 must precede T_1 in the serialization order because T_2's

read request for object X comes before T_1's write request for the same object. Similarly, T_3 must precede T_2 because T_3's write request precedes T_2's read request for object Y. This serialization order will clearly result in missing transaction T_2's deadline, and more than likely, it will cause transaction T_1 to miss its own deadline as well.

Using the OPT-BC protocol the same sequence will produce the serialization order: $T_1 \prec T_3 \prec T_2$. Transaction T_1 reaches its validation phase first and having no conflicts with the other two transactions[8] is committed. At this point – following the Broadcast Commit variant – transaction T_2 is restarted because it read object X, which is written by the committing transaction (T_1). This causes transaction T_2 to be delayed enough to allow T_3 to enter its validation phase before T_2. Transaction T_3 is, thus, committed, which results in transaction T_2 being restarted again (due to the conflict on object Y this time). This second restart severely damages transaction T_2's chances of meating its deadline.

4 Three members of the SCC-nS family

In this section, we consider three SCC-based algorithms: SCC-1S, SCC-2S, and SCC-MS. The first represents a specialization of the general case discussed in the previous section, which uses the minimum possible amount of redundancy. The second can be seen as the simplest form of a hybrid algorithm, allowing each transaction to have one optimistic and one pessimistic (speculative) shadow. Finally, the third represents the most flexible of this family of SCC algorithms. SCC-MS and SCC-1S illustrate the two extremes with regard to the level of the *computation redundancy* they tolerate and the *real-time performance* they accomplish.

4.1 The One-Shadow SCC Algorithm (SCC-1S)

In this case, every uncommitted transaction in the system has only an optimistic shadow. Neither a speculative nor a pessimistic shadow is present. The optimistic shadow for each T_i, then, runs under the assumption that it will be the first (among all the other transactions with which T_i conflicts) to commit. Therefore, it executes without incurring any blocking delays. The SCC-1S algorithm, thus, resembles the OCC-BC algorithm in that optimistic shadows of transactions continue to execute either until they are validated and committed, or until they are aborted (by a validating transaction). This represents the one extreme regarding the amount of redundant computations that SCC algorithms introduce. At their lowest extent, when no redundant computations are allowed, they identify with the optimistic paradigm. The more redundancy they are allowed to use, the better their real-time performance.

[8] Assuming that there are no conflicting actions other than the ones shown in figure 13.

4.2 The Two-Shadow SCC Algorithm (SCC-2S)

The SCC-2S allows a maximum of two shadows per uncommitted transaction to exist in the system at any point in time: an optimistic shadow and a speculative shadow. The speculative shadow of a transaction T_i, called here the *pessimistic* shadow T_i^p (in contrast with the optimistic shadow) is subject to blocking and restart. It is kept ready to replace the optimistic shadow T_i^o, should such a replacement be necessary. T_i^p runs under the *pessimistic* assumption that it will be the last (among all the other transactions with which T_i conflicts) to commit.

The SCC-2S like the SCC-1S algorithm resembles the OCC-BC algorithm in that optimistic shadows of transactions continue to execute either until they are validated and committed or until they are aborted (by a validating transaction). The difference, however, is that SCC-2S keeps a *pessimistic* shadow for each executing transaction to be used if that transaction must abort. The pessimistic shadow is basically a replica of the optimistic shadow, except that it is blocked at the *earliest* point where a read-write conflict is detected between the transaction it represents and any other uncommitted transaction in the system. Should this conflict materialize into a consistency threat, the pessimistic shadow is promoted to become the new optimistic shadow, and execution is *resumed* (instead of being *restarted* as would be the case with OCC-BC) from the point where the potential conflict was discovered. The detailed algorithm, as well as illustrative examples of its use can be found in [Best92b].

4.3 The Multi-Shadow SCC Algorithm (SCC-MS)

This is an SCC-based algorithm, which allows the redundancy level for individual transactions to differ and vary dynamically. Each transaction T_r has, at each point of its execution, one optimistic shadow T_r^o, and i speculative shadows T_r^i, where i is the number of detected potential conflicts in which T_r participates.

This variant is more powerful than the generic SCC algorithm presented above. Its superior performance results from its flexibility to deal with *any* transaction conflicts. Contrary to the generic SCC algorithm, it does not fix a priori the number of speculative shadows that each transaction in the system is allowed to have at any point in its lifetime. Thus, every time that a new conflict is encountered, a new speculative shadow is created, to accommodate it. Moreover, each individual transaction can have a different degree of redundancy, in the number of shadows it can originate. This flexibility, of course, is gained at the expense of an increased amount of redundant computations that are allowed in the system. See figure 14 for the details of the SCC-MS algorithm.

A. The Start Rule: When the execution of a new transaction T_r is requested, an optimistic shadow $T_r^o \in \mathcal{T}^{\mathcal{O}}$ is created and executed.

B. The Read Rule: Whenever an optimistic shadow T_r^o wishes to read an object X, then:

for all T_u^o in $\mathcal{T}^{\mathcal{O}}$, such that T_u^o wrote object X, **do**
1. **if** $(\forall T_r^i \in \mathcal{T}_r^{\mathcal{S}}, (T_u, _) \notin WaitFor(T_r^i))$ **then**{
1.1 A new speculative shadow T_r^j is forked off T_r^o;
1.2 $WaitFor(T_r^j) \leftarrow \{(T_u, X)\}$;}

C. The Write Rule: Whenever an optimistic shadow T_u^o wishes to write an object X, then:

for all T_r^o in $\mathcal{T}^{\mathcal{O}}$, such that T_r^o read object X **do**
1. **if** $(\forall T_r^i \in \mathcal{T}_r^{\mathcal{S}}, (T_u, _) \notin WaitFor(T_r^i))$ **then**{
1.1 A new speculative shadow T_r^i is forked off $BestShadow(T_r, X)$;
1.2 $WaitFor(T_r^i) \leftarrow \{(T_u, X)\}$;
2. }**else if** $(\exists T_r^k \in \mathcal{T}_r^{\mathcal{S}}, \exists Y : ((X, _) \in ReadSet(T_r^k) \wedge (T_u, Y) \in WaitFor(T_r^k)))$ **then**{
2.1 T_r^k is aborted and replaced by T_r^m which is forked off $BestShadow(T_r, X)$;
2.2 $WaitFor(T_r^m) \leftarrow \{(T_u, X)\}\}$;

D. The Blocking Rule: A standby shadow T_r^i is blocked at the *earliest point* at which it wishes to Read an object X that is written by any transaction T_u, such that $(T_u, X) \in WaitFor(T_r^i)$.

E. The Commit Rule: Whenever it is decided to commit an optimistic shadow T_r^o on behalf of a transaction T_r, then:

1. $\forall T_r^i \in \mathcal{T}_r^{\mathcal{S}}, T_r^i$ is aborted;
2. **for** all $T_u \in \mathcal{T}$, such that $(\exists T_u^i \in \mathcal{T}_u^{\mathcal{S}} : (T_r, X) \in WaitFor(T_u^i))$ **do**{
2.1 T_u^o is aborted;
2.2 T_u^i is promoted to become the new optimistic shadow of T_u;
2.3 **for** all $T_u^j \in \mathcal{T}_u^{\mathcal{S}}$, such that $(X, _) \in ReadSet(T_u^j)$ **do**
2.3.1 T_u^j is aborted};

Figure 14: The Multi-Shadow SCC Algorithm.

5 Conclusion

SCC-based algorithms offer a new dimension (namely redundancy) that can be used effectively in RTDBMS. SCC-based algorithms use redundancy to ensure that serializable schedules are discovered and adopted as early as possible, thus increasing the likelihood of the timely commitment of transactions with strict timing constraints. Using SCC, several shadow transactions execute on behalf of a given uncommitted transaction so as to protect against the hazards of blockages and restarts, which are characteristics of Pessimistic and Optimistic Concurrency Control algorithms, respectively.

In this paper, we presented a generic algorithm SCC-nS, which characterizes a family of algorithms that differ in the total amount of redundancy they introduce. We described the SCC-nS algorithm both informally and formally. We established its correctness by showing that it only admits serializable histories. We demonstrated its superiority for RTDBMS through numerous examples. Three members of the SCC-nS family (namely SCC-1S, SCC-2S, and SCC-MS) were singled out and contrasted. SCC-1S does not introduce any additional redundancy and is shown to be equivalent to the OCC-BC algorithm of [Mena82, Robi82]. SCC-2S allows exactly one additional *pessimistic* shadow in the system and is shown to outperform OCC-BC with respect to the timely commitment of transactions. SCC-MS introduces as many shadows as necessary to account for all possible *pair-wise* conflicts between uncommitted transactions. This is in contrast to the general algorithm described in [Best92b], where conflicts involving more than two transactions are also considered.

A number of interesting research problems related to SCC-based algorithms are currently under investigation. In particular, the criteria to be used for redundancy and resource management (*e.g.* shadow replacement policies) can be related to the priority and urgency of the uncommited transactions in the system. Another interesting problem is to characterize (via simulation) the performance of the SCC-nS family. This requires the adoption of appropriate performance metrics suitable for RTDBMS (*e.g.* a characterization of the distribution of tardy tasks given a particular load distribution [Best92a]).

References

[Abbo88] Robert Abbott and Hector Garcia-Molina. "Scheduling real-time transactions: A performance evaluation." In *Prooceedings of the 14th International Conference on Very Large Data Bases*, Los Angeles, Ca, 1988.

[Agra87] R. Agrawal, M. Carey, and M. Linvy. "Concurency control performance modeling: Alternatives and implications." *ACM Transaction on Database Systems*, 12(4), December 1987.

[Bern87] A. Bernstein, A. Philip, V. Hadzilacos, and N. Goodman. *Concurrency Control And Recovery In Database Systems.* Addison-Wesley, 1987.

[Best92a] Azer Bestavros. "Performance measures for real-time database management systems." Technical Report (In progress), Computer Science Department, Boston University, Boston, MA, July 1992.

[Best92b] Azer Bestavros. "Speculative concurrency control." Technical Report TR-16-92, Computer Science Department, Boston University, Boston, MA, July 1992.

[Best92c] Azer Bestavros and Spyridon Braoudakis. "Speculative concurrency control algorithms for real-time databases: An alternative expression of transaction priority." Technical Report (In progress), Computer Science Department, Boston University, Boston, MA, September 1992.

[Boks87] C. Boksenbaum, M. Cart, J. Ferrié, and J. Francois. "Concurrent certifications by intervals of timestamps in distributed database systems." *IEEE Transactions on Software Engineering*, pages 409–419, April 1987.

[Buch89] A. P. Buchmann, D. C. McCarthy, M. Hsu, and U. Dayal. "Time-critical database scheduling: A framework for integrating real-time scheduling and concurrency controls." In *Proceedings of the 5th International Conference on Data Engineering*, Los Angeles, California, February 1989.

[Eswa76] K. P. Eswaran, J. N. Gray, R. A. Lorie, and I. L. Traiger. "The notions of consistency and predicate locks in a database system." *Communications of the ACM*, 19(11):624–633, November 1976.

[Fran85] Peter Franaszek and John Robinson. "Limitations of concurrency in transaction processing." *ACM Transactions on Database Systems*, 10(1), March 1985.

[Gray76] J. N. Gray, R. A. Lorie, G. R. Putzolu, and I. L. Traiger. "Granularity of locks and degrees of consistensy in a shared data base." In G. M. Nijssen, editor, *Modeling in Data Base Management Systems*, pages 365–395. North-Holland, Amsterdam, The Netherlands, 1976.

[Hari90a] Jayant R. Haritsa, Michael J. Carey, and Miron Linvy. "Dynamic real-time optimistic concurrency control." In *Prooceedings of the 11th Real-Time Systems Symposium*, December 1990.

[Hari90b] Jayant R. Haritsa, Michael J. Carey, and Miron Linvy. "On being optimistic about real-time constraints." In *Prooceedings of the 1990 ACM PODS Symposium*, April 1990.

[Huan89] J. Huang, J. A. Stankovic, D. Towsley, and K. Ramamritham. "Experimental evaluation of real-time transaction processing." In *Proceedings of the 10th Real-Time Systems Symposium*, December 1989.

[Huan91] Jiandong Huang, John A. Stankovic, and Don Towslwy Krithi Ramamritham. "Experimental evaluation of real-time optimistic concurrency control schemes." In *Prooceedings of the 17th International Conference on Very Large Data Bases*, Barcelona, Spain, September 1991.

[Kim91] Woosaeng Kim and Jaideep Srivastava. "Enhancing real-time dbms performance with multiversion data and priority based disk scheduling." In *Prooceedings of the 12th Real-Time Systems Symposium*, December 1991.

[Kort90] Henry Korth. "Triggered real-time databases with consistency constraints." In *Proceedings of the 16th International Conference on Very Large Data Bases*, Brisbane, Australia, 1990.

[Kung81] H. Kung and John Robinson. "On optimistic methods for concurrency control." *ACM Transactions on Database Systems*, 6(2), June 1981.

[Lin90] Yi Lin and Sang Son. "Concurrency control in real-time databases by dynamic adjustment of serialization order." In *Proceedings of the 11th Real-Time Systems Symposium*, December 1990.

[Mena82] D. Menasce and T. Nakanishi. "Optimistic versus pessimistic concurrency control mechanisms in database management systems." *Information Systems*, 7(1), 1982.

[Papa79] Christos Papadimitriou. "The serializability of concurrent database updates." *Journal of the ACM*, 26(4):631–653, October 1979.

[Robi82] John Robinson. *Design of Concurrency Controls for Transaction Processing Systems*. PhD thesis, Carnegie Mellon University, Pittsburgh, PA, 1982.

[Sha88] Lui Sha, R. Rajkumar, and J. Lehoczky. "Concurrency control for distributed real-time databases." *ACM, SIGMOD Record*, 17(1):82–98, 1988.

[Sha91] Lui Sha, R. Rajkumar, Sang Son, and Chun-Hyon Chang. "A real-time locking protocol." *IEEE Transactions on Computers*, 40(7):793–800, 1991.

[Sing88] Mukesh Singhal. "Issues and approaches to design real-time database systems." *ACM, SIGMOD Record*, 17(1):19–33, 1988.

[Son92] S. Son, S. Park, and Y. Lin. "An integrated real-time locking protocol." In *Prooceedings of the IEEE International Conference on Data Engineering*, Tempe, AZ, February 1992.

[Stan88] John Stankovic and Wei Zhao. "On real-time transactions." *ACM, SIGMOD Record*, 17(1):4–18, 1988.