

Sound Transaction-based Reduction without Cycle Detection

Vladimir Levin¹ Robert Palmer²
Shaz Qadeer¹ Sriram K. Rajamani¹

Microsoft¹
The University of Utah²

vladlev@microsoft.com, rpalmer@cs.utah.edu,
qadeer@microsoft.com, sriram@microsoft.com

April 2005
Technical Report
MSR-TR-2005-40

Partial order reduction is widely used to alleviate state space explosion in model checkers for concurrent programs. Traditional approaches to partial order reduction are based on *ample* sets. Natural ample sets can be computed for threads that communicate with each other predominantly through message queues. For threads that communicate with shared memory using locks for synchronization, Lipton's theory of reduction provides a promising way to aggregate several fine-grained transitions into larger transactions. In traditional partial order reduction, actions that are not in the ample set are delayed, thus avoiding the redundant exploration of equivalent interleaving orders. Delaying the execution of actions indefinitely can lead to loss of soundness. This is called the *ignoring problem*. The usual solution to the ignoring problem is by Cycle Detection. Explicit state model checkers usually use Depth First Search, and when a cycle is detected, disallow using a reduced ample set that closes the cycle.

The ignoring problem exists in transaction-based reduction as well. We present a novel solution to the ignoring problem in the context of transaction-based reduction. We designate certain states as *commit points* and track the exploration to discover whether the reduced exploration guarantees a path from each commit point to a state where the transaction is completed. If such a path does not exist, we detect this at the time a commit point is popped from the stack, and schedule all threads at the commit point. This paper presents our algorithm, called Commit Point Completion (CPC). We have implemented both CPC and Cycle Detection in the Zing model checker, and find that the CPC algorithm performs better.

Microsoft Research
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052
<http://www.research.microsoft.com>

1 Introduction

Partial order methods have been widely used as an optimization in building model checkers for concurrent software [1–6]. Traditional partial order reduction methods are based on the notion of *independence* between actions. Two actions α and β are independent if (1) they do not disable one another and (2) if both actions are enabled in a state s , then executing them in the order α followed by β from s , or in the order β followed by α from s , leads to the same resulting state. Partial order reduction algorithms explore a subset of enabled actions in each state called the *ample* set. The set of all actions enabled in a state s is denoted $Enabled(s)$ and the ample set of actions in a state s is denoted $Ample(s)$. Obviously, $Ample(s) \subseteq Enabled(s)$. For partial order reduction to be sound, ample sets need to be chosen in such a way that a transition that is dependent on a transition in $Ample(s)$ cannot execute without a transition in $Ample(s)$ occurring first (see condition **C1** in [1] page 148). Choosing a minimal ample set satisfying **C1** is a very hard problem. In practice, ample sets are formed from local actions, and from restricted versions of send and receive actions, such as: sending to a queue, with the sender having exclusive rights of sending to the queue, and receiving from a queue, with the receiver having exclusive rights of receiving from the queue [3]. If the system consists of threads interacting via shared memory, Lipton’s theory of reduction [7] provides an alternate way to do partial order reduction. Reduction views a transaction as a sequence of actions $a_1, \dots, a_m, x, b_1, \dots, b_n$ such that each a_i is a right mover and each b_i is a left mover. A right mover is an action that commutes to the right of every action by another thread; a left mover is an action that commutes to the left of every action by another thread. Thus, to detect transactions we need to detect right and left movers. Most programs consistently use mutexes to protect accesses to shared variables, we can exploit this programming discipline to infer left and right movers:

- The action `acquire(m)`, where m is a mutex, is a right mover.
- The action `release(m)` is a left mover.
- An action that accesses only a local variable or shared variable that is consistently protected by a mutex is both a left mover and a right mover.

A transaction is a sequence of right movers, followed by a *committing* action that is not a right mover, followed by a sequence of left movers. A transaction can be in two states: pre-commit or post-commit. A transaction starts in the pre-commit state and stays in the pre-commit state as long as right movers are being executed. When the committing action is executed, the transaction moves to the post-commit state. The transaction stays in the post-commit state as long as left movers are being executed until the transaction completes. In addition to being able to exploit programmer-imposed discipline such as protecting each shared variable consistently with the same lock, transaction-based reduction allows extra optimizations such as summarization [8].

Ignoring Problem. All partial order reduction algorithms work by delaying the execution of certain actions, thus avoiding the redundant exploration of

```

int g = 0;

void T1() {
L0:  g = 1;
L1:  skip;
L2:  while(true){
L3:    skip;
    }
L4:  return;
}

void T2() {
M0:  assert(g == 0);
M1:  return;
}

P = { T1() } || { T2() }

```

Fig. 1. Ignoring problem

equivalent executions. For instance, if thread t_1 executes an action from state s_1 that reads and writes only local variables, then thread t_2 does not need to be scheduled to execute in s_1 , and t_2 's scheduling can be delayed without losing soundness. For any interleaving that starts from s_1 and ends in a state where some thread t goes wrong, there exists an equivalent interleaving where the execution of t_2 is delayed at s_1 . However, unless we are careful, the scheduling of thread t_2 can be delayed indefinitely resulting in loss of soundness. This situation is called the *ignoring problem* in partial order reduction.

Consider the example in Figure 1. The initial state of this program has two threads t_1 and t_2 starting to execute functions T1 and T2 respectively. The program has one global variable g , which has an initial value 0. A typical model checking algorithm first schedules t_1 to execute the statement at line L0, which updates the value of g to 1. Let us call this state s_1 . Since the next statement executed by thread t_1 from s_1 reads and writes only local variables of t_1 (namely its program counter) and does not read or write the global variables, partial order reduction algorithms delay execution of thread t_2 at state s_1 . Continuing, the while loop in lines L2 and L3 also reads and writes only the local variables of t_1 and thus execution of t_2 can be delayed during the execution of these statements as well. However, since reached states are stored, and a newly generated state is not re-explored if it is already present in the set of reached states, a fix-point is reached after executing the loop in T_1 once. Thus, the execution of t_2 is delayed indefinitely, and the reduction algorithm can be unsound, and say that the assertion in line M0 is never violated.

Most partial order reduction algorithms “fix” the ignoring problem by detecting cycles, and disallowing the actions of a thread to be ample when a cycle is “closed” (see condition **C3**, pages 150 and 158 in [1]). Since explicit-state model checkers usually use depth first search (DFS), cycle detection can be performed by detecting whether a newly generated state is already present in the DFS stack. In the SPIN model checker this is implemented using a bit in the hash table entry for reached states. This bit indicates whether the newly generated successor state is currently also on the depth first search stack.

Cycle detection is neither necessary nor sufficient for transaction-based reduction. Consider the variant of our current example in Figure 2. Here, we have

```

int g = 0;

void T1() {
L0:   g = 1;
L1:   skip;
L2:   if (*) {
L3:       while(true){
L4:           skip;
L5:       }
L6:   return;
}

void T2() {
M0:   assert(g == 0);
M1:   return;
}

P = { T1() } || { T2() }

```

Fig. 2. Cycle detection is not necessary for transaction-based reduction

introduced a nondeterministic choice in line L2 of procedure T1. In one branch of the nondeterministic choice, we have a while-loop with statements reading and writing only local variables of thread t_1 (lines L3-L4). The other branch of the nondeterministic choice just terminates the procedure. In this case, even without doing any cycle detection, since one branch of the nondeterministic choice terminates, a partial order reduction algorithm can schedule thread t_2 after procedure T1 terminates, and thus the assertion violation in line M0 can be detected. If we consider a variant of this example, where the entire “if” statement (from line L2 to L6 is replaced by `assume(false)` at line L2, some other mechanism in addition to cycle detection is needed to schedule the thread t_2 after t_1 executes the statement L1.

In the current literature on transaction-based reduction, the ignoring problem is addressed indirectly by disallowing certain types of infinite executions, such as those consisting of only internal hidden actions, within each thread (see Condition **C** from Section 4.2 in [9] which forbids the transaction from having infinite executions after committing, but without completing, and well-formedness assumption **Wf-iffinite-invis** from Section 4 in [10]). These assumptions do not hold in practice. In particular, when we analyze models that arise from abstractions (such as predicate abstraction) of programs, it is common to have loops with non-deterministic termination conditions, which violate the above assumptions. Thus, a more direct and computationally effective solution to the ignoring problem is required for wide applicability of transaction-based reduction. This paper presents a novel solution to this problem.

CPC Algorithm. We propose a new technique called *Commit Point Completion* (CPC) to solve the ignoring problem without cycle detection. We keep track of the state immediately after the committing action is executed, called the *commit point*. When a committed transaction completes, we simply mark the commit point as completed. When an unmarked commit point is about to be popped from the DFS stack, we schedule all threads from that state. Our insight is that we can delay the decision to forcibly end a transaction up to the time when commit point is about to be popped from the stack, avoiding taking such a decision pre-maturely when cycles are closed.

```

Mutex m;
int x = 0; /* all accesses to x will be guarded by m*/
int y = 0; /* accesses to y are not guarded */

void T1() {
L0: acq(m);
L1: y := 42;
L2: x := 1;
L3: rel(m);
L4: while (true)
    { skip; }
}
P = { T1() } || { T2() } || { T3() }

void T2() {
M0: acq(m);
M1: assert(x == 0);
M2: rel(m);
}

void T3() {
N0: y = 10;
}

```

Fig. 3. CPC algorithm in the presence of left movers

In the example from Figure 1 the state immediately after t_1 executes the statement at line L0 is a commit point. Due to the non-terminating while loop, the transaction that is committed here never completes. Thus, when this commit point is about to be popped from the DFS stack, it is unmarked, and the CPC algorithm schedules thread t_2 from this state, and the assertion violation in line M0 is detected. The example from Figure 2, has an identical commit point. However, since one nondeterministic branch completes the transaction, the commit point gets marked. Thus, when the commit point gets popped from the DFS stack, the other thread t_2 is not scheduled. Note that the assertion failure at M0 is detected even without scheduling thread t_2 from the commit point, because t_2 will be scheduled by the reduction algorithm after the transaction in t_1 completes on one of the nondeterministic branches.

The above description of the CPC algorithm is simplistic. In the presence of left movers there may be more than one commit point for a transaction, and all of these commit points need to reach a state where the transaction completes to ensure sound reduction. For example, consider the example shown in Figure 3. In this example, there are two global variables x and y and one mutex m . All accesses to x are protected by mutex m , and are thus both movers. Accesses to y are unprotected, and are hence non-movers. Acquires of mutex m are right movers and releases are left movers as mentioned earlier. Thus, when thread T1 executes the assignment to y at label L1, its transaction commits, since the access to y is a non-mover. The resulting state, where y has just been assigned 42 and the program counter of the thread T1 is at L2 is a commit point. Due to the infinite while-loop at L4 this committed transaction never completes, and the CPC algorithm can schedule threads at the above commit point when it is about to be popped from the stack. However, for us to detect the assertion violation at line M1 of thread T2, another commit point needs to be established in T1 after the assignment to x at line L2. We handle this case by designating every state in a committed-transaction obtained by executing a “pure” left mover (i.e, a transaction that is a left mover but not a both-mover) as a commit point.

Thus, in T1, the state after executing the release at line L3 is also designated as a commit point, and the algorithm schedules T2 when this state is about to be popped, leading to the assertion violation.

We have implemented the CPC algorithm in the Zing model checker at MSR. Section 5 presents experimental results that compare the CPC algorithm with a Cycle Detection algorithm for various Zing programs. The results clearly demonstrate that the CPC algorithm generally explores far fewer states than the Cycle Detection algorithm.

Outline. The rest of the paper is organized as follows. Section 2 introduces notations for describing multithreaded programs precisely. Section 3 gives an abstract framework for sound transaction-based reduction. Section 4 presents the CPC algorithm and a statement of its correctness. This section contains the core new technical results of the paper. Section 5 presents experimental results from the implementation of the CPC algorithm in the Zing model checker. Section 6 compares the CPC algorithm with related work, and Section 7 concludes the paper.

2 Multithreaded programs

The store of a multithreaded program is partitioned into the global store *Global* and the local store *Local* of each thread. We assume that the domains of *Local* and *Global* are finite sets. The set *Local* of local stores has a special store called *wrong*. The local store of a thread moves to *wrong* on failing an assertion and thereafter the failed thread does not make any other transitions.

$$\begin{aligned}
 t, u &\in \text{ \textit{Tid}} = \{1, \dots, n\} \\
 i, j &\in \text{ \textit{Choice}} = \{1, 2, \dots, m\} \\
 g &\in \text{ \textit{Global}} \\
 l &\in \text{ \textit{Local}} \\
 ls &\in \text{ \textit{Locals}} = \text{ \textit{Tid}} \rightarrow \text{ \textit{Local}} \\
 \text{ \textit{State}} &= \text{ \textit{Global}} \times \text{ \textit{Locals}}
 \end{aligned}$$

A multithreaded program (g_0, ls_0, T) consists of three components. g_0 is the initial value of the global store. ls_0 maps each thread id $t \in \text{ \textit{Tid}}$ to the initial local store $ls_0(t)$ of thread t . We model the behavior of the individual threads using two transition relations:

$$\begin{aligned}
 T_G &\subseteq \text{ \textit{Tid}} \times (\text{ \textit{Global}} \times \text{ \textit{Local}}) \times (\text{ \textit{Global}} \times \text{ \textit{Local}}) \\
 T_L &\subseteq \text{ \textit{Tid}} \times \text{ \textit{Local}} \times \text{ \textit{Choice}} \times \text{ \textit{Local}}
 \end{aligned}$$

The relation T_G models system visible thread steps. The relation $T_G(t, g, l, g', l')$ holds if thread t can take a step from a state with global store g and local store l , yielding (possibly modified) stores g' and l' . The relation T_G has the property that for any t, g, l , there is at most one g' and l' such that $T_G(t, g, l, g', l')$. We use functional notation and say that $(g', l') = T_G(t, g, l)$ if $T_G(t, g, l, g', l')$. Note that in the functional notation, T_G is a partial function from $\text{ \textit{Tid}} \times (\text{ \textit{Global}} \times \text{ \textit{Local}})$ to $(\text{ \textit{Global}} \times \text{ \textit{Local}})$. The relation T_L models thread local thread steps. The relation $T_L(t, l, i, l')$ holds if thread t can move its local store from l to l' on choice i . The nondeterminism in the behavior of a thread

is captured by T_L . This relation has the property that for any t, l, i , there is a unique l' such that $T_L(t, l, i, l')$.

The program starts execution from the state (g_0, ls_0) . At each step, any thread may make a transition. The transition relation $\rightarrow_t \subseteq State \times State$ of thread t is the disjunct of the system visible and thread local transition relations defined below. For any function h from A to B , $a \in A$ and $b \in B$, we write $h[a := b]$ to denote a new function such that $h[a := b](x)$ evaluates to $h(x)$ if $x \neq a$, and to b if $x = a$.

$$\frac{T_G(t, g, ls(t), g', l')}{(g, ls) \rightarrow_t (g', ls[t := l'])} \quad \frac{T_L(t, ls(t), i, l')}{(g, ls) \rightarrow_t (g, ls[t := l'])}$$

The transition relation $\rightarrow \subseteq State \times State$ of the program is the disjunction of the transition relations of the various threads:

$$\rightarrow = \exists t. \rightarrow_t$$

3 Transactions

Transactions occur in multithreaded programs because of the presence of right and left movers. Inferring which actions of a program are right and left movers is a problem that is important but orthogonal to the contribution of this paper. In this section, we assume that right and left movers are available to us as the result of a previous analysis (see, e.g. [11]).

Let $RM, LM \subseteq T_G$ be subsets of the transition relation T_G with the following properties for all $t \neq u$:

1. If $RM(t, g_1, l_1, g_2, l_2)$ and $T_G(u, g_2, l_3, g_3, l_4)$, there is g_4 such that $T_G(u, g_1, l_3, g_4, l_4)$ and $RM(t, g_4, l_1, g_3, l_2)$.
2. If $T_G(u, g_1, l_1, g_2, l_2)$ and $RM(t, g_2, l_3, g_3, l_4)$, then for all g', l' ($T_G(t, g_1, l_3, g', l') \Rightarrow RM(t, g_1, l_3, g', l')$).
3. If $T_G(u, g_1, l_1, g_2, l_2)$ and $LM(t, g_2, l_3, g_3, l_4)$, there is g_4 such that $LM(t, g_1, l_3, g_4, l_4)$ and $T_G(u, g_4, l_1, g_3, l_2)$.
4. If $T_G(u, g_1, l_1, g_2, l_2)$ and $LM(t, g_1, l_3, g_3, l_4)$, there is g_4 such that $LM(t, g_2, l_3, g_4, l_4)$.

The first property states that a right mover of thread t commutes to the right of a transition of a different thread u . The second property states that if a right mover of thread t is enabled in the post-state of a transition of another thread u , and thread t is enabled in the pre-state, then the transition of thread t is a right mover in the pre-state. The third property states that a left mover of thread t commutes to the left of a transition of a different thread u . The fourth property states that a left mover that is enabled in the pre-state of a transition by another thread is also enabled in the post-state.

Our analysis is parameterized by the values of RM and LM and only requires that they satisfy these four properties. The larger the relations RM and LM ,

the longer the transactions our analysis infers. Therefore, these relations should be as large as possible in practice.

In order to minimize the number of explored interleaving orders and to maximize reuse, we would like to **infer** transactions that are as long as possible (i.e., they are maximal with respect to a given thread). To implement this inference, we introduce in each thread a boolean local variable to keep track of the phase of that thread's transaction. Note that this instrumentation is done automatically by our analysis, and not by the programmer. The phase variable of thread t is true if thread t is in the right mover (or pre-commit) part of the transaction; otherwise the phase variable is false. We say that the transaction *commits* when the phase variable moves from true to false. The initial value of the phase variable for each thread is *false*.

$$\begin{aligned} p, p' \in \text{Boolean} &= \{\text{false}, \text{true}\} \\ \ell, \ell' \in \text{Local}^\# &= \text{Local} \times \text{Boolean} \\ \ell s, \ell s' \in \text{Locals}^\# &= \text{Tid} \rightarrow \text{Local}^\# \\ \text{State}^\# &= \text{Global} \times \text{Locals}^\# \end{aligned}$$

Let $\text{Phase}(t, (g, \ell s))$, the phase of thread t in state $(g, \ell s)$ be the second component of $\ell s(t)$.

The initial value of the global store of the instrumented program remains g_0 . The initial value of the local stores changes to ℓs_0 , where $\ell s_0(t) = \langle \ell s_0(t), \text{false} \rangle$ for all $t \in \text{Tid}$. We instrument the transition relations T_G and T_L to generate a new transition relation $T^\#$.

$$T^\# \subseteq \text{Tid} \times (\text{Global} \times \text{Local}^\#) \times \text{Choice} \times (\text{Global} \times \text{Local}^\#)$$

$$T^\#(t, g, \langle l, p \rangle, i, g', \langle l', p' \rangle) \stackrel{\text{def}}{=} \begin{cases} \vee T_G(t, g, l, g', l') \wedge \\ p' = (\text{RM}(t, g, l, g', l') \wedge (p \vee \neg \text{LM}(t, g, l, g', l'))) \\ \vee T_L(t, l, i, l') \wedge g = g' \wedge p' = p \end{cases}$$

In the definition of $T^\#$, the relation between p' and p reflects the intuition that if p is true, then p' continues to be true as long as it executes right mover transitions. The phase changes to false as soon as the thread executes a transition that is not a right mover. Thereafter, it remains false as long as the thread executes left movers. Then, it becomes true again as soon as the thread executes a transition that is a right mover and not a left mover. A transition from T_L does not change the phase. We overload the transition relation \rightarrow_t defined in Section 2 to represent transitions in the instrumented transition relation. Similar to the functional notation defined for T_G in Section 2, we sometimes use functional notation for $T^\#$.

Given an instrumented transition relation $T^\#$, we define three sets for each thread t : $\mathcal{R}(t), \mathcal{L}(t), \mathcal{N}(t) \subseteq \text{State}^\#$. These sets respectively define when a thread is executing in the right mover part of a transaction, the left mover part of a transaction, and outside any transaction. These three sets are a *partition* of $\text{State}^\#$ defined as follows:

$$\begin{aligned}
- \mathcal{R}(t) &= \{ (g, \ell s) \mid \exists l. \ell s(t) = \langle l, true \rangle \wedge l \notin \{ls_0(t), wrong\} \}. \\
- \mathcal{L}(t) &= \left\{ (g, \ell s) \mid \begin{array}{l} \exists l. \ell s(t) = \langle l, false \rangle \wedge l \notin \{ls_0(t), wrong\} \wedge \\ (\exists i, g', l'. LM(t, g, l, g', l') \vee T_L(t, l, i, l')) \end{array} \right\}. \\
- \mathcal{N}(t) &= State^\# \setminus (\mathcal{R}(t) \cup \mathcal{L}(t)).
\end{aligned}$$

The definition of $\mathcal{R}(t)$ says that thread t is in the right mover part of a transaction if and only if the local store of t is neither its initial value nor *wrong* and the phase variable is true. The definition of $\mathcal{L}(t)$ says that thread t is in the left mover part of a transaction if and only if the local store of t is neither its initial value nor *wrong*, the phase variable is false, and there is an enabled transition that is either a left mover or thread-local. Note that since the global transition relation is deterministic, the enabled left mover is the only enabled transition that may access a global variable. Since $(\mathcal{R}(t), \mathcal{L}(t), \mathcal{N}(t))$ is a partition of $State^\#$, once $\mathcal{R}(t)$ and $\mathcal{L}(t)$ have been picked, the set $\mathcal{N}(t)$ is implicitly defined.

A transaction $p \rightarrow_i^* p'$ has the following properties.

1. $p \in \mathcal{N}(t)$, and
2. – Either
 - (a) for $0 < j < n$ $p_0, \dots, p_j \in \mathcal{R}(t) \vee \mathcal{L}(t)$, and
 - (b) $p' = p_n \in \mathcal{L}(t) \vee \mathcal{N}(i)$.
- Or
 - (a) for $0 < j \leq n$ $p_j \in \mathcal{R}(t)$.

A transition relation Y *right-commutes* with a transition relation Z if $Y \circ Z \subseteq Z \circ Y$, and Y *left-commutes* with Z if $Z \circ Y \subseteq Y \circ Z$. Any transaction of thread u in which $p' \in \mathcal{R}(u)$ can be commuted to the right of transactions of other threads. Therefore we consider sequences of transactions.

$$p = p_1 \xrightarrow{+}_{t(1)} p_2 \xrightarrow{+}_{t(2)} p_3 \cdots p_k \xrightarrow{+}_{t(k)} p_{k+1} = q_1 \xrightarrow{+}_{u(1)} q_2 \xrightarrow{+}_{u(2)} q_3 \cdots q_l \xrightarrow{+}_{u(l)} q_{l+1} = q$$

$$\underbrace{(p_2 = p_{2,1} \xrightarrow{+}_{t(2)} \cdots \xrightarrow{+}_{t(2)} p_{2,x} = p_3)}_{(p_2 = p_{2,1} \xrightarrow{+}_{t(2)} \cdots \xrightarrow{+}_{t(2)} p_{2,x} = p_3)} \quad \underbrace{(q_2 = q_{2,1} \xrightarrow{+}_{u(2)} \cdots \xrightarrow{+}_{u(2)} q_{2,x} = q_3)}_{(q_2 = q_{2,1} \xrightarrow{+}_{u(2)} \cdots \xrightarrow{+}_{u(2)} q_{2,x} = q_3)}$$

Fig. 4. A sequence of transactions.

A sequence of states from Figure 4 is called a *sequence of transactions* if

- for all $1 \leq m \leq k$, if $p_m = p_{m,1} \xrightarrow{+}_{t(m)} \cdots \xrightarrow{+}_{t(m)} p_{m,x} = p_{m+1}$, then
 - (1) $p_{m,1} \in \mathcal{N}(t(m))$,
 - (2) $p_{m,2}, \dots, p_{m,x-1} \in \mathcal{R}(t(m)) \vee \mathcal{L}(t(m))$,
 - (3) $p_{m,x} \in \mathcal{L}(t(m)) \vee \mathcal{N}(t(m))$.
- for all $1 \leq m \leq l$, if $q_m = q_{m,1} \xrightarrow{+}_{u(m)} \cdots \xrightarrow{+}_{u(m)} q_{m,x} = q_{m+1}$, then
 - (1) $q_{m,1} \in \mathcal{N}(u(m))$,
 - (2) $q_{m,2}, \dots, q_{m,x} \in \mathcal{R}(u(m))$.

Intuitively, for every i , $p_i \xrightarrow{+}_{t(i)} p_{i+1}$ is a committed transaction and for every j , $q_j \xrightarrow{+}_{u(j)} q_{j+1}$ is an uncommitted transaction.

For any state set $X \subseteq State^\#$ and transition relation $Y \subseteq State^\# \times State^\#$, by X/Y we mean the transition relation obtained by restricting Y to pairs whose first component is in X . Similarly, by $Y \setminus X$ we mean the restriction of Y to pairs whose second component is in X .

Theorem 1 says that we can represent any arbitrarily interleaved execution sequence with a sequence in a canonical form based on transactions.

Theorem 1 (Reduction). *For all t , let $\mathcal{R}(t)$, $\mathcal{L}(t)$, and $\mathcal{W}(t)$ be sets of states, and \rightarrow_t be a transition relation. Suppose for all t ,*

1. $\mathcal{R}(t)$, $\mathcal{L}(t)$, and $\mathcal{W}(t)$ are pairwise disjoint,
2. $(\mathcal{L}(t)/\rightarrow_t \setminus \mathcal{R}(t))$ is false,

and for all $u \neq t$,

3. $(\rightarrow_t \setminus \mathcal{R}(t))$ right-commutes with \rightarrow_u ,
4. $(\mathcal{L}(t)/\rightarrow_t)$ left-commutes with \rightarrow_u ,
5. if $p \rightarrow_t q$, then $p \in \mathcal{R}(u) \Leftrightarrow q \in \mathcal{R}(u)$, $p \in \mathcal{L}(u) \Leftrightarrow q \in \mathcal{L}(u)$, and $p \in \mathcal{W}(u) \Leftrightarrow q \in \mathcal{W}(u)$.

Let $\mathcal{N}(t) = \neg(\mathcal{R}(t) \vee \mathcal{L}(t))$, $\mathcal{N} = \forall t. \mathcal{N}(t)$, and $\rightarrow = \exists t. \rightarrow_t$.

Suppose $p \in \mathcal{N}$ and $p \rightarrow^* q$. **Then** there exist $k, l \geq 0$ and a transition sequence

$$\begin{aligned} p &= p_1 \xrightarrow{+}_{t(1)} p_2 \xrightarrow{+}_{t(2)} p_3 \cdots p_k \xrightarrow{+}_{t(k)} p_{k+1} \\ &= q_1 \xrightarrow{+}_{u(1)} q_2 \xrightarrow{+}_{u(2)} q_3 \cdots q_l \xrightarrow{+}_{u(l)} q_{l+1} = q \end{aligned}$$

with the following properties:

- for all $1 \leq m \leq k$, if $p_m = p_{m,1} \rightarrow_{t(m)} \cdots \rightarrow_{t(m)} p_{m,x} = p_{m+1}$, then
 1. $p_{m,1} \in \mathcal{N}(t(m))$,
 2. $p_{m,2}, \dots, p_{m,x-1} \in \mathcal{R}(t(m)) \vee \mathcal{L}(t(m))$, and
 3. $p_{m,x} \in \mathcal{L}(t(m)) \vee \mathcal{N}(t(m))$.
- for all $1 \leq m \leq l$, if $q_m = q_{m,1} \rightarrow_{u(m)} \cdots \rightarrow_{u(m)} q_{m,x} = q_{m+1}$, then
 1. $q_{m,1} \in \mathcal{N}(u(m))$, and
 2. $q_{m,2}, \dots, q_{m,x} \in \mathcal{R}(u(m))$.

Proof. We will prove our theorem by induction on the length of the sequence from p to q . Suppose $p \in \mathcal{N}$ and $p \rightarrow^* q$. For the base case we have $k = l = 0 \Rightarrow p = q$. For the inductive case we will assume that the inductive hypothesis holds (i.e., the sequence exists) for $p \rightarrow^* q$. We will consider the added step $q \rightarrow_t q'$. Now we will perform a case analysis on q .

- $q \in \mathcal{N}(t)$. We show the following two statements by mutual induction:
 - $q_m \in \mathcal{N}(t)$ for all $1 \leq m \leq l + 1$
 - $t \neq u(m)$ for all $1 \leq m \leq l$

For the base case, we have $q_{l+1} = q \in \mathcal{N}(t)$. There are two inductive cases. Suppose $q_{m+1} \in \mathcal{N}(t)$. From the definition of the sequence we know that $q_{m+1} \in \mathcal{R}(u(m))$. Now if $t = u(m)$ then $q_{m+1} \in \mathcal{N}(t) \cap \mathcal{R}(u(m))$ which would result in a contradiction. Therefore we have $t \neq u(m)$. We also have $q_m \in \mathcal{N}(t)$ by the definition of the sequence. We can conclude that $p_{k+1} = q_1 \in \mathcal{N}(t)$.

Since $u(1), \dots, u(l)$ are all different from t , we commute all the actions performed by these threads to the right of the action by thread t such that thread t executes at state $p_{k+1} \rightarrow_t q'$. Resulting in one of the following sequences:

- $q' \in \mathcal{N}(t)$

$$\begin{aligned} p &= p_1 \xrightarrow{+}_{t(1)} p_2 \xrightarrow{+}_{t(2)} p_3 \cdots p_k \xrightarrow{+}_{t(k)} p_{k+1,1} \rightarrow_t p_{k+1,2} \\ &= q'_1 \xrightarrow{+}_{u(1)} q'_2 \xrightarrow{+}_{u(2)} q'_3 \cdots q'_l \xrightarrow{+}_{u(l)} q'_{l+1} = q' . \end{aligned}$$

- $q' \in \mathcal{L}(t)$

$$\begin{aligned} p &= p_1 \xrightarrow{+}_{t(1)} p_2 \xrightarrow{+}_{t(2)} p_3 \cdots p_k \xrightarrow{+}_{t(k)} p_{k+1,1} \rightarrow_t p_{k+1,2} \\ &= q'_1 \xrightarrow{+}_{u(1)} q'_2 \xrightarrow{+}_{u(2)} q'_3 \cdots q'_l \xrightarrow{+}_{u(l)} q'_{l+1} = q' . \end{aligned}$$

- $q' \in \mathcal{R}(t)$

$$\begin{aligned} p &= p_1 \xrightarrow{+}_{t(1)} p_2 \xrightarrow{+}_{t(2)} p_3 \cdots p_k \xrightarrow{+}_{t(k)} p_{k+1} \\ &= q_{1,1} \rightarrow_t q_{1,2} \xrightarrow{+}_{u(1)} q'_2 \xrightarrow{+}_{u(2)} q'_3 \cdots q'_l \xrightarrow{+}_{u(l)} q'_{l+1} = q' . \end{aligned}$$

– $q \in \mathcal{L}(t)$. By induction we show that

- $q_m \in \mathcal{L}(t)$ for all $1 \leq m \leq l+1$
- $t \neq u(m)$ for all $1 \leq m \leq l$

For the base case, we have $q_{l+1} = q \in \mathcal{L}(t)$. There are two inductive cases. Suppose $q_{m+1} \in \mathcal{L}(t)$. Since $q_{m+1} \in \mathcal{R}(u(m))$, we have $t \neq u(m)$. Therefore $q_m \in \mathcal{L}(t)$. In particular, we get that $p_{k+1} = q_1 \in \mathcal{L}(t)$. There must be some m such that $t(m) = t$. Otherwise $p = p_1 \in \mathcal{L}(t)$ which violates our assumptions. Consider the greatest m such that $t(m) = t$. Then $t(m+1), \dots, t(k)$ and $u(1), \dots, u(l)$ are all different from t . (Recall that $(\rightarrow_t / \mathcal{L}(t))$ left-commutes with $\rightarrow_{t(m)}$ for all $t(m) \neq t$.) We commute the action performed by thread t to the left of all actions performed by these threads to get the execution sequence

$$\begin{aligned} p &= p_1 \xrightarrow{+}_{t(1)} p_2 \cdots p_m \xrightarrow{+}_{t(m)} \\ &\quad p'_{m+1} \xrightarrow{+}_{t(m+1)} p'_{m+2} \cdots p'_k \xrightarrow{+}_{t(k)} p'_{k+1} \\ &= q'_1 \xrightarrow{+}_{u(1)} q'_2 \xrightarrow{+}_{u(2)} q'_3 \cdots q'_l \xrightarrow{+}_{u(l)} q'_{l+1} = q' . \end{aligned}$$

Since $q \in \mathcal{L}(t)$, we have $q' \in \mathcal{L}(t) \vee \mathcal{N}(t)$. Therefore $p'_{m+1} \in \mathcal{L}(t) \vee \mathcal{N}(t)$.

– $q \in \mathcal{R}(t)$. We first prove by contradiction that $u(1), \dots, u(l)$ are all distinct from each other. Suppose $1 \leq a, b \leq l$ are such that $u(a) = u(b)$ and $u(m) \neq u(a)$ for all $a < j < b$. Then we know that $q_{a+1} \in \mathcal{R}(u(a))$. Therefore

$q_b \in \mathcal{R}(u(a))$ which is a contradiction since $q_b \in \mathcal{N}(u(b))$. From which we conclude that any thread that moved from $q_1 \cdots q_{l+1}$ only moved once.

We now prove by contradiction that $t = u(m)$ for some m such that $1 \leq m \leq l$. If not, then $p_{k+1} = q_1 \in \mathcal{R}(t)$. We now perform a case analysis.

1. Suppose there is no m' such that $t(m') = t$. Then $p = p_1 \in \mathcal{R}(t)$ which is a contradiction since $p \in \mathcal{N}$.
2. Suppose m' is the greatest such that $t(m') = t$. Then $p_{m'+1} \in \mathcal{R}(t)$, which is a contradiction since $p_{m'+1} \in \mathcal{L}(t(m')) \vee \mathcal{N}(t(m'))$.

From which we can conclude that thread t did move in the sequence from $q_1 \cdots q_{l+1}$.

Therefore $u(1), \dots, u(m-1), u(m+1), \dots, u(l)$ are all different from $t = u(m)$. We first commute the actions performed by $u(m+1), \dots, u(l)$ to the right of the action performed by t to get the execution sequence

$$\begin{aligned} p &= p_1 \xrightarrow{+}_{t(1)} p_2 \xrightarrow{+}_{t(2)} p_3 \cdots p_k \xrightarrow{+}_{t(k)} p_{k+1} \\ &= q_1 \xrightarrow{+}_{u(1)} q_2 \cdots q_j \xrightarrow{+}_{u(m)} q'_{m+1} \cdots q'_l \xrightarrow{+}_{u(l)} q'_{l+1} = q' \end{aligned}$$

If $q'_{m+1} \in \mathcal{R}(u(m))$ then we are done. Otherwise, we commute actions performed by threads $u(1), \dots, u(m-1)$ to the right of all actions performed by thread $u(m)$ to get the execution sequence

$$\begin{aligned} p &= p_1 \xrightarrow{+}_{t(1)} p_2 \xrightarrow{+}_{t(2)} p_3 \cdots p_k \xrightarrow{+}_{t(k)} p_{k+1} \\ &= q'_1 \xrightarrow{+}_{u'(1)} q'_2 \cdots q'_j \xrightarrow{+}_{u'(j)} q'_{j+1} \cdots q'_l \xrightarrow{+}_{u'(l)} q'_{l+1} = q' \end{aligned}$$

where $u'(1) = u(m), u'(2) = u(1), \dots, u'(m) = u(m-1), u'(m+1) = u(m+1), \dots, u'(l) = u(l)$.

Since q is in $\mathcal{R}(t), \mathcal{L}(t)$, or $\mathcal{N}(t)$ our proof is complete.

Theorem 2 says if a sequence of transactions results in a state $q \in \mathcal{N}$ then every transaction that committed also completed.

Theorem 2. *Suppose the antecedents to Theorem 1 are satisfied. Let $\rightarrow = \exists t. (\forall u \neq t. \mathcal{N}(u)) / \rightarrow_t$. If $\forall t. p \in \mathcal{N}(t) \wedge p \rightarrow^* q \wedge \forall t. q \in \mathcal{N}(t)$ then $p \rightarrow^* q$*

Proof. By theorem 1 we have that there exists a transition sequence

$$\begin{aligned} p &= p_1 \xrightarrow{+}_{t(1)} p_2 \xrightarrow{+}_{t(2)} p_3 \cdots p_k \xrightarrow{+}_{t(k)} p_{k+1} \\ &= q_1 \xrightarrow{+}_{u(1)} q_2 \xrightarrow{+}_{u(2)} q_3 \cdots q_l \xrightarrow{+}_{u(l)} q_{l+1} = q \end{aligned}$$

with the properties stated in that theorem. Since $\forall t, q \in \mathcal{N}(t)$ we have $l = 0$ and $p_{k+1} = q_1 = q$. We will now show by induction that for $0 \leq m \leq k+1$ and for all $t, p_m \in \mathcal{N}(t)$. The base case is that $p_m = p_{k+1} \in \mathcal{N}(t)$ for all t . The inductive case is as follows. Suppose $p_{m+1} \in \mathcal{N}(t)$ for all t . Since $p_m \xrightarrow{*}_{t(m)} p_{m+1}$ we have $p_m \in \mathcal{N}(t)$ for all $t \neq t(m)$. We have as a property of the sequence that $p_m = p_{m,1} \in \mathcal{N}(t(m))$ so for all $t, p_m \in \mathcal{N}(t)$. Therefore for all $m \leq k+1$ and for all $t, p_m \in \mathcal{N}(t)$. Therefore $p \rightarrow^* q$.

Theorem 3 says that it is unnecessary to explore any transactions where the final state is in $\mathcal{R}(i)$ to find a state that is in \mathcal{W} .

Theorem 3. *Suppose the antecedents to Theorem 1 are satisfied. Let $\mathcal{W} = \exists t. \mathcal{W}(t)$. If $p \in \mathcal{N} \wedge p \rightarrow^* q \wedge q \in \mathcal{W}$ then $\exists q'$ such that $p \rightarrow^* q' \rightarrow^* q$, for all t , $q' \in \mathcal{L}(t) \vee q' \in \mathcal{N}(t)$ and $q' \in \mathcal{W}$.*

Proof. Since $q \in \mathcal{W}$ we have $\exists t. q \in \mathcal{N}(t)$. By induction we have already proved that $q_{l+1} \in \mathcal{N}(t) \Rightarrow p_{k+1} \in \mathcal{N}(t) \wedge u(m) \neq t$ for all m . By the same argument we conclude that $p_{k+1} \in \mathcal{W}(t)$. Let $q' = p_{k+1}$ and we are done.

Theorem 4 gives a sufficient subset of transaction sequences to insure soundness.

Theorem 4. *Suppose the antecedents to Theorem 1 are satisfied. Let $\mathcal{W} = \exists t. \mathcal{W}(t)$. If $p \in \mathcal{N} \wedge p \rightarrow^* q \wedge q \in \mathcal{W}$. Then $p \rightarrow^* q$ or $p \rightarrow^* q' \rightarrow^* q \wedge (\forall t. q' \in \mathcal{L}(t) \vee q' \in \mathcal{N}(t)) \wedge q' \in \mathcal{W}$.*

Proof. The representative sequence from p to q from Theorem 1 takes one of two forms. If for all t , $q \in \mathcal{N}(t)$ then all transitions satisfy \rightarrow in which case $q = q'$. Otherwise, by theorem 3 we have $p \rightarrow^* q'$ via the sequence

$$p = p_1 \xrightarrow{t(1)} p_2 \xrightarrow{t(2)} p_3 \cdots p'_k \xrightarrow{t(m)} p'_{k+1} = q'$$

where $q' \in \mathcal{W}(t)$. For each m , $1 \leq m \leq k+1$, $p_{m,x} \in \mathcal{L}(t(m)) \vee p_{j,x} \in \mathcal{N}(t(m))$. Therefore $q' \in \mathcal{W}$ is visited by one of these sequences and we are done.

We will find the following Lemmas useful in the proof of Theorem 5.

Lemma 1 says that the mover status of a transition of thread u does not change when a right mover transition of thread t is commuted past the transition of thread u .

Lemma 1. *If $RM(t, g_1, l_1, g_2, l_2)$ and $T_G(u, g_2, l_3, g_3, l_4)$, then there exists g_4 such that*

1. $T_G(u, g_1, l_3, g_4, l_4)$
2. $RM(u, g_1, l_3, g_4, l_4) \Leftrightarrow RM(u, g_2, l_3, g_3, l_4)$
3. $LM(u, g_1, l_3, g_4, l_4) \Leftrightarrow LM(u, g_2, l_3, g_3, l_4)$

Proof. We will prove each item separately.

1. Follows immediately from Assumption 1 on RM and LM .
2. In the \Rightarrow direction, assume $RM(u, g_1, l_3, g_4, l_4)$. We have assumed that $RM, LM \subseteq T_G$. From Assumption 1 on RM and LM we have $T_G(t, g_4, l_1, g_3, l_2)$. We can again apply Assumption 1 on RM and LM to get that there exists some g_5 such that $T_G(t, g_1, l_1, g_5, l_2)$ and $RM(u, g_5, l_3, g_3, l_4)$. We have also assumed that for any state (g, ℓ) there is at most one state (g', ℓ') such that $T_G(t, g, \ell(t), g', \ell'(t))$. Therefore we know that $T_G(t, g_1, l_1, g_2, l_2)$ is the only transition enabled for t at (g_1, l_1) . Therefore $g_5 = g_2$ and $RM(u, g_2, l_3, g_3, l_4)$. In the \Leftarrow direction, assume $RM(u, g_2, l_3, g_3, l_4)$. Then by Assumption 2 on RM and LM and (1) from above we have $RM(u, g_1, l_3, g_4, l_4)$.

3. In the \Rightarrow direction, assume $LM(u, g_1, l_3, g_4, l_4)$. Then by assumption 4 on RM and LM there exists some g_5 such that $LM(u, g_2, l_3, g_5, l_4)$. We have assumed that $RM, LM \subseteq T_G$. We have also assumed that for any state (g, ℓ_s) there is at most one state (g', ℓ_s') such that $T_G(t, g, \ell_s(t), g', \ell_s'(t))$. Therefore we know that $T_G(u, g_2, l_3, g_5, l_4)$ is the only transition enabled for u at (g_2, l_3) . Therefore $g_5 = g_2$ and $LM(u, g_2, l_3, g_3, l_4)$. In the \Leftarrow direction, assume $LM(u, g_2, l_3, g_3, l_4)$. By Assumption 3 on RM and LM we have that there exists some g_5 such that $LM(u, g_1, l_3, g_5, l_4)$ and $T_G(t, g_5, l_1, g_3, l_2)$. Again by the deterministic nature of T_G we have $g_5 = g_4$ and we are done.

Lemma 2 says that the mover status of a transition of thread u does not change when a left mover transition of thread t is commuted past the transition of thread u .

Lemma 2. *If $T_G(u, g_1, l_1, g_2, l_2)$ and $LM(t, g_2, l_3, g_3, l_4)$, then there exists g_4 such that*

1. $T_G(u, g_4, l_1, g_3, l_2)$
2. $RM(u, g_4, l_1, g_3, l_2) \Leftrightarrow RM(u, g_1, l_1, g_2, l_2)$
3. $LM(u, g_4, l_1, g_3, l_2) \Leftrightarrow LM(u, g_1, l_1, g_2, l_2)$

Proof. We will prove each item separately.

1. Follows immediately from Assumption 3 on RM and LM .
2. For the \Rightarrow direction, assume $RM(u, g_4, l_1, g_3, l_2)$. We have $LM(t, g_1, l_3, g_4, l_4)$ and $RM(u, g_4, l_1, g_3, l_2)$. By Assumption 2 on RM and LM we can conclude $RM(u, g_1, l_1, g_2, l_2)$. For the \Leftarrow direction, assume $RM(u, g_1, l_1, g_2, l_2)$. We have $RM(u, g_1, l_1, g_2, l_2)$ and $LM(t, g_2, l_3, g_3, l_4)$. By Assumption 1 on RM and LM there exists some g_5 such that $T_G(t, g_1, l_3, g_5, l_4)$ and $RM(u, g_5, l_1, g_3, l_2)$. We have assumed that $RM, LM \subseteq T_G$. We have also assumed that for any state (g, ℓ_s) there is at most one state (g', ℓ_s') such that $T_G(t, g, \ell_s(t), g', \ell_s'(t))$. Therefore we know that $T_G(t, g_1, l_3, g_5, l_4)$ is the only transition enabled for t at (g_1, l_3) . Therefore $g_5 = g_4$ and we are done.
3. For the \Rightarrow direction, assume $LM(u, g_4, l_1, g_3, l_2)$. We have $LM(t, g_1, l_3, g_4, l_4)$ and $LM(u, g_4, l_1, g_3, l_2)$. By Assumption 3 on RM and LM there exists some g_5 such that $LM(u, g_1, l_1, g_5, l_2)$ and $T_G(t, g_5, l_3, g_3, l_4)$. From the deterministic nature of T_G we have $g_5 = g_2$ and $LM(u, g_1, l_1, g_2, l_2)$. For the \Leftarrow direction, assume $LM(u, g_1, l_1, g_2, l_2)$. By Assumption 3 on RM and LM we have $LM(t, g_1, l_3, g_4, l_4)$ and $T_G(u, g_4, l_1, g_3, l_2)$. Assumption 4 on RM and LM gives us that there exists a g_5 such that $LM(u, g_4, l_1, g_5, l_2)$. From the deterministic nature of T_G we have $g_5 = g_3$ and $LM(u, g_4, l_1, g_3, l_2)$.

Theorem 5 shows that our construction satisfies the antecedents of Theorem 1.

Theorem 5. *Let $P = (g_0, \ell_{s_0}, T^\#)$ be the instrumented multithreaded program. For all $t \in Tid$, let $(\mathcal{R}(t), \mathcal{L}(t))$ and $\mathcal{N}(t)$ be the partition obtained from P as defined above. For all $t \in Tid$, let $\mathcal{W}(t) = \{(g, \ell_s) \mid \exists p. \ell_s(t) = \langle \text{wrong}, p \rangle\}$. Then,*

the antecedents of Theorem 1 are satisfied. Moreover, for any state $(g', ls') \in \mathcal{W}(t)$ that is reachable from (g_0, ls_0) , there is another state $(g'', ls'') \in \mathcal{W}(t)$ that is reachable from (g_0, ls_0) by a sequence of transactions.

Proof. We proceed by showing our construction satisfies the antecedents for Theorem 1.

1. For all t , $\mathcal{R}(t)$, $\mathcal{L}(t)$, and $\mathcal{W}(t)$ are pairwise disjoint. We note that $\mathcal{W}(t) \subseteq \mathcal{N}(t)$ so it is sufficient to show that $\mathcal{R}(t)$, $\mathcal{L}(t)$, and $\mathcal{N}(t)$ are pairwise disjoint.
 - $\mathcal{L}(t) \cap \mathcal{R}(t) = \emptyset$ Suppose $p \in \mathcal{L}(t)$ then $ls(t) = \langle l, false \rangle$. Therefore $p \notin \mathcal{R}(t)$. Suppose $p \in \mathcal{R}(t)$ then $ls(t) = \langle l, true \rangle$. Therefore $p \notin \mathcal{L}(t)$.
 - $\mathcal{L}(t) \cap \mathcal{N}(t) = \emptyset$ By the definition of $\mathcal{N}(t)$.
 - $\mathcal{R}(t) \cap \mathcal{N}(t) = \emptyset$ As above.
2. For all t , $(\mathcal{L}(t)/\rightarrow_t \setminus \mathcal{R}(t))$ is false. Suppose $(g, ls(t)) \rightarrow_t (g', ls'(t))$ by some transition x where $(g, ls(t)) \in \mathcal{L}(t)$ and $(g', ls'(t)) \in \mathcal{R}(t)$. Then $ls(t) = \langle l, false \rangle$ and $ls'(t) = \langle l', true \rangle$. For the phase to change from *false* to *true* we must execute $x \in RM \setminus LM$. However a state at which x is enabled cannot exist in $\mathcal{L}(t)$, thus we have a contradiction.
3. For all $u \neq t$, $(\rightarrow_t \setminus \mathcal{R}(t))$ right-commutes with \rightarrow_u . For all states $(g_1, ls_1), (g_2, ls_2)$ such that $(g_1, ls_1(t)) \rightarrow_t (g_2, ls_2(t))$ and $(g_2, ls_2) \in \mathcal{R}(t)$, we have $((g_1, ls_1), (g_2, ls_2)) \in (\rightarrow_t \setminus \mathcal{R}(t))$. By the definition of $\mathcal{R}(t)$, we have $RM(t, g_1, ls_1(t), g_2, ls_2(t))$. Now for all states (g_3, ls_3) such that $(g_2, ls_2(u)) \rightarrow_u (g_3, ls_3(u))$, by Assumption 1 of *RM* and *LM* we have that $RM(t, g_1, ls_1(t), g_2, ls_2(t))$ and $T_G(u, g_2, ls_2(u), g_3, ls_3(u))$. Therefore there exists g_4 such that $T_G(u, g_1, ls_1(u), g_4, ls_2(u))$ and $RM(t, g_4, ls_2(t), g_3, ls_3(t))$. What remains to be shown is that the phase of thread u does not change by commuting the actions of t and u . Only thread u may access $ls_i(u)$ for any i . By Lemma 1 we have that that a transition in *RM* or *LM* that is enabled at (g_2, ls_2) will also be enabled at (g_1, ls_1) . Therefore if $RM(t, g_1, ls_1(t), g_2, ls_2(t))$ and $T^\#(u, g_2, ls_2(u) = \langle l_3, p_u \rangle, i, g_3, ls_3(u) = \langle l_4, p'_u \rangle)$, there exists g_4 such that $T^\#(u, g_1, ls_1(u) = \langle l_3, p_u \rangle, i, g_4, ls_2(u) = \langle l_4, p'_u \rangle)$ and $RM(t, g_4, ls_2(t), g_3, ls_3(t))$ Therefore $(\rightarrow_t \setminus \mathcal{R}(t)) \circ \rightarrow_u \subseteq \rightarrow_u \circ (\rightarrow_t \setminus \mathcal{R}(t))$ which is our definition of right-commutes.
4. For all $u \neq t$, $(\mathcal{L}(t)/\rightarrow_t)$ left-commutes with \rightarrow_u . For all states $(g_2, ls_2), (g_3, ls_3)$ such that $(g_2, ls_2) \in \mathcal{L}(t)$ and $(g_2, ls_2(t)) \rightarrow_t (g_3, ls_3(t))$ then $((g_2, ls_2), (g_3, ls_3)) \in (\mathcal{L}(t)/\rightarrow_t)$. By the definition of $\mathcal{L}(t)$ we have $LM(t, g_2, ls_2(t), g_3, ls_3(t))$. Now for all states (g_1, ls_1) such that $(g_1, ls_1(u)) \rightarrow_u (g_2, ls_2(u))$ from Assumption 2 of *RM* and *LM* we have that if $T_G(u, g_1, ls_1(u), g_2, ls_2(u))$ and $LM(t, g_2, ls_2(t), g_3, ls_3(t))$ then there exists g_4 such that $LM(t, g_1, ls_1(t), g_4, ls_2(t))$ and $T_G(u, g_4, ls_2(u), g_3, ls_3(u))$. Again it remains to be shown that the phase of thread u does not change by commuting the actions of t and u . Only thread u may access $ls(u)$. By Lemma 2 we have that a transition in *RM* or *LM* that is enabled at (g_1, ls_1) will also be enabled at (g_4, ls_2) . Therefore if $T^\#(u, g_1, ls_1(u) = \langle l_3, p_u \rangle, i, g_2, ls_2(u) = \langle l_4, p'_u \rangle)$ and $LM(t, g_2, ls_2(t), g_3, ls_3(t))$ then there exists g_4 such that $LM(t, g_1, ls_1(t), g_4, ls_2(t))$ and $T^\#(u, g_4, ls_2(u) = \langle l_3, p_u \rangle, i, g_3, ls_3(u) =$

$\langle l_4, p'_u \rangle$). Therefore $\rightarrow_u \circ (\mathcal{L}(t)/\rightarrow_t) \subseteq (\mathcal{L}(t)/\rightarrow_t) \circ \rightarrow_u$ which is our definition of left-commutes.

5. For all $u \neq t$, if $(g, \ell s) \rightarrow_t (g', \ell s')$, then $(g, \ell s) \in \mathcal{R}(u) \Leftrightarrow (g', \ell s') \in \mathcal{R}(u)$, $(g, \ell s) \in \mathcal{L}(u) \Leftrightarrow (g', \ell s') \in \mathcal{L}(u)$, and $(g, \ell s) \in \mathcal{W}(u) \Leftrightarrow (g', \ell s') \in \mathcal{W}(u)$. We will consider each case individually. The following observation will be useful in each of the cases: Since thread t cannot modify the local store or phase of thread u , we have $\ell s(u) = \langle l, p \rangle \Leftrightarrow \ell s'(u) = \langle l, p \rangle$ for all l, p . Suppose $\ell s(u) = \langle l, p \rangle$. Then $\ell s'(u) = \langle l, p \rangle$.

- $(g, \ell s) \in \mathcal{R}(u) \Leftrightarrow (g', \ell s') \in \mathcal{R}(u)$. This case is easily proved from the definition of $\mathcal{R}(u)$ and the fact that $\ell s(u) = \langle l, p \rangle = \ell s'(u)$.
- $(g, \ell s) \in \mathcal{W}(u) \Leftrightarrow (g', \ell s') \in \mathcal{W}(u)$. This case is easily proved from the definition of $\mathcal{W}(u)$ and the fact that $\ell s(u) = \langle l, p \rangle = \ell s'(u)$.
- $(g, \ell s) \in \mathcal{L}(u) \Leftrightarrow (g', \ell s') \in \mathcal{L}(u)$. This case is proved from the definition of $\mathcal{L}(u)$, the fact that $\ell s(u) = \langle l, p \rangle = \ell s'(u)$, and from the following observation about LM derived from Assumptions 3 and 4: $(\exists g'', l''. LM(u, g, l, g'', l'')) \Leftrightarrow (\exists g'', l''. LM(u, g', l, g'', l''))$.

We substitute our concrete transition relation for the abstract relation of Theorem 1 and the state partition sets as defined. We use the consequents of Theorem 1 to get that if a state $(g', \ell s')$ is reachable from $(g_0, \ell s_0)$ where $(g', \ell s') \in \mathcal{W}(t)$, there is another state $(g'', \ell s'')$ that is reachable from $(g_0, \ell s_0)$ by a transaction sequence and $(g'', \ell s'') \in \mathcal{W}(t)$.

$$\begin{array}{c}
 \text{(INIT)} \\
 \hline
 \Sigma(g_0, \ell s_0) \\
 \\
 \text{(STEP)} \\
 \frac{\forall u \neq t. (g, \ell s) \in \mathcal{N}(u) \quad \Sigma(g, \ell s) \quad T^\#(t, g, \ell s(t), i, g', \ell')}{\Sigma(g', \ell s[t := \ell'])}
 \end{array}$$

Fig. 5. Model checking with unsound reduction.

Using the values of $\mathcal{N}(t)$ for all $t \in Tid$, we model check the multithreaded program by computing the least fixpoint of the set of rules in Figure 5. This model checking algorithm schedules a thread only when no other thread is executing inside a transaction. This algorithm is potentially unsound for the following reason. If a transaction in thread t commits but never finishes, the shared variables modified by this transaction become visible to other threads. However, the algorithm does not explore transitions of other threads from any state after the transaction commits. Section 4 presents a more sophisticated algorithm which ensures that all threads are explored from some state in the post-commit phase of every transaction.

4 Commit Point Completion

```
record TraversalInfo {
  state : State
  tid : Tid
  numTids : Tid
  choice : int
  LM : Boolean
  RM : Boolean
  Xend : Boolean
  CPC : Boolean
}
```

Fig. 6. TraversalInfo declaration.

This section presents the CPC algorithm and its soundness proof, which are the core new technical contributions of this paper. The algorithm uses Depth First Search (DFS). Each state in the DFS stack is encapsulated using a `TraversalInfo` record shown in Figure 6. In addition to the state, the `TraversalInfo` records: (1) `tid`, the id of the thread used to reach the state, (2) `numTids`, the number of threads active in the state, (3) `choice`, the current index among the nondeterministic choices executable by thread `tid` in this state, (4) `LM`, a boolean which is set to true iff the action used to reach this state is a left mover, (5) `RM`, a boolean which is set to true iff the action used to reach this state is a right mover, (6) `Xend`, a boolean which is set to true iff the algorithm decides to schedule other threads at this state, and (7) `CPC`, a boolean which is relevant for only states with `phase` equal to false, and is set to true by the algorithm if there exists a path of transitions of the thread generating the state to a state where all threads are scheduled.

The helper functions for the CPC algorithm perform the following actions. `Enabled` determines whether the current thread has a transition enabled at a given state. `Execute` applies the transition relation $T^\#$ to the current state. `Update` schedules the next thread to run.

Figure 8 gives two variants of the CPC algorithm. We will discuss first the variant with line L19 commented out, then motivate and discuss the addition of the optimization.

The statement at L4 peeks at the `TraversalInfo` `q` on top of the stack and explores all successors of the state using actions from thread `q.tid`. If the phase of `q` is false, then for each such successor `q'`, if the action used to generate `q'` is not a left-mover, then we update `q.Xend` to true at label L7. The invariant associated with the CPC flag is the following: If `q` is about to be popped from the stack and `q.CPC` is true and $Phase(q.tid, q.state)$ is false then there exists a path to a state where `Xend` is true. Thus, at label L8 we set `q.CPC` to true if `q.Xend` is true. The `Xend` and `CPC` fields are also updated when a `TraversalInfo` is

```

Boolean Enabled(TraversalInfo q) {
  let (g, ℓs) = q.state in
  return (∃ g', ℓ'. T#(q.tid, g, ℓs(q.tid), q.choice, g', ℓ'))
}

TraversalInfo Execute(TraversalInfo q) {
  let (g, ℓs) = q.state in
  let (g', ℓ') = T#(q.tid, g, ℓs(q.tid), q.choice) in
  State succ = (g', ℓs[q.tid := ℓ'])
  q.choice = q.choice + 1;
  return { state = succ,
          tid = q.tid,
          numTids = 1,
          choice = 1,
          CPC = false,
          Xend = false,
          LM = LM(q.tid, q.state, succ)
          RM = RM(q.tid, q.state, succ) }
}

TraversalInfo Update(TraversalInfo q) {
  Tid nextTid = ite((q.tid == |Tid|), 1, q.tid + 1)
  return { state = q.state,
          tid = nextTid,
          numTids = q.numTids + 1,
          choice = 1,
          CPC = q.CPC,
          Xend = q.Xend,
          LM = q.LM,
          RM = q.RM }
}

```

Fig. 7. Helper procedures for the CPC algorithm.

```

Hashtable table;
Stack stack;
TraversalInfo q, q', q'', pred;

stack = new Stack
table = new Hashtable

L0: q' = { state = (go, ls0),
        tid = 1,
        numTids = 1,
        choice = 1,
        CPC = true,
        Xend = true,
        LM = false,
        RM = false }

L1: table.Add(q'.state, q')
L2: stack.Push(q')

L3: while (stack.Count > 0)
L4:   q = stack.Peek()
L5:   if (Enabled(q))
L6:     q' = Execute(q)
L7:     q.Xend = q.Xend || (¬Phase(q.tid, q.state) && ¬q'.LM)
L8:     q.CPC = q.CPC || q.Xend
L9:     if (IsMember(table, q'.state))
L10:      q'' = Lookup(table, q'.state)
L11:      q.CPC = q.CPC || q''.CPC
L12:     else /* undiscovered state */
L13:      table.Add(q'.state, q')
L14:      stack.Push(q')
L15:     end if
L16:     q.choice = q.choice + 1
L17:   else
L18:     q.Xend = q.Xend || (¬Phase(q.tid, q.state) && ¬q.CPC
L19:                        (* && ¬q.RM *))
L20:     q.CPC = q.CPC || q.Xend

L21:   stack.Pop()
L22:   pred = stack.Peek()
L23:   pred.CPC = pred.CPC || q.CPC

L24:   if (q.Xend && q.numTids < |Tid|)
L25:     q' = Update(q)
L26:     stack.Push(q')
L27:   end if
L28: end if
L29: end while

```

Fig. 8. CPC algorithm for sound reduction.

popped from the stack. In particular, at label L18, when q is about to be popped from the stack, if its phase is false and $q.CPC$ is false, then we set $q.Xend$ to true and force scheduling of all threads at q . If $q.Xend$ is true, then at label L24 we ensure that all threads are scheduled from q . Figure 7 contains helper procedures for the CPC algorithm.

We want to show that if a `TraversalInfo` record q exists on the stack such that $q.CPC$ is true then there is a sequence of left-mover transitions to a state represented in some `TraversalInfo` record q' such that $q'.Xend$ is true. To show this we will argue by induction on the order in which `TraversalInfo` records are popped from the stack. However we will require an auxiliary lemma about the stack to complete the induction. The auxiliary lemma assumes that Lemma 4 holds for states that have been popped from the stack and is invoked within Lemma 4 in two cases: (1) the base case where nothing has been popped from the stack and (2) the inductive case where we assume that Lemma 4 holds for all `TraversalInfo` records that have been popped from the stack. We proceed with the auxiliary lemma.

Lemma 3. *Suppose Lemma 4 holds for all `TraversalInfo` records q not on the stack such that $q.CPC$ is true and $Phase(q.tid, q.state)$ is false. Suppose there exists a stack of `TraversalInfo` records such that for all `TraversalInfo` records q on the stack such that $q.CPC$ is true and $Phase(q.tid, q.state)$ is false, there exists a sequence of left-mover transitions of thread $q.tid$ such that every state in the sequence is represented in a `TraversalInfo` record q' and $q'.CPC$ is true and $Phase(q'.tid, q'.state)$ is false and the final `TraversalInfo` q' , where $q'.state$ is the final state in the sequence, $q'.Xend$ is true.*

Proof. We will proceed by induction on the size of the stack. For the base case, let the stack be size 1. Then q is the initial state and $q.CPC$ is true and $q.Xend$ is true as set on line L0.

For the inductive case, suppose the stack depth is i . If $q.Xend$ is true then we are done. If $q.Xend$ is false then $q.CPC$ could only have been set at line L11 or L23. Suppose $q.CPC$ is set by a successor that is not represented on the stack at line L11 or L23. If the transition were not a left-mover then $q.Xend$ would be true (we have assumed it is false so there would be a contradiction). If $Phase(q'.tid, q'.state)$ were true then the action executed could not have been a left-mover. We have shown that there is a left-mover transition to a state that is not represented on the stack such that the CPC flag of the associated `TraversalInfo` is true and the phase is false. Since the successor state is not represented on the stack and we have assumed the property holds for such states we are done. Otherwise the successor state that was revisited must also be represented on the stack. Since q is about to be popped from the stack, all other states on the stack must be represented at some depth $l < i$. If the transition executed by thread $q.tid$ were not a left-mover then $q.Xend$ would be true and we would have a contradiction. Therefore a left-mover transition exists from $q.state$ to a state represented in a `TraversalInfo` q' such that q' is lower in the stack than q . If $Phase(q'.tid, q'.state)$ were true then it would be different than

$Phase(q.tid, q.state)$ and we would have a contradiction (as we have shown that $q'.LM$ was true and in order for the phase to change from false to true, $q'.LM$ must be false). We have assumed that $q.CPC$ is true and set at line L10. This could only result from $q'.CPC$ also being true. Therefore we have assumed that the sequence exists for all states on the stack with such properties and shown that the last state on the stack extends the sequence.

Lemma 4 (CPC). *Let q be a `TraversalInfo` record such that $q.CPC$ is true. Then there exists a sequence of left-mover transitions to another state held in a `TraversalInfo` record q' such that $q'.Xend$ is true and for every intermediate `TraversalInfo` record q'' where $q''.state$ is in the sequence, $q''.CPC$ is true.*

Proof. If $q.Xend$ is true then we are done. Suppose $q.Xend$ is false. $q.CPC$ is set to true in exactly four places:

1. L8 when $q.Xend$ is true. In this case, let $q = q'$.
2. L11 when a successor is discovered to be in the hash table and the successor's CPC field is true.
3. L20 when $q.Xend$ is true. In this case, let $q = q'$.
4. L23 when a successor's CPC bit is true and the successor is about to be popped from the stack.

Suppose q is about to be popped from the stack at line L21. We will proceed by induction on the order that `TraversalInfo` records are popped from the stack. For the base case let q be the first traversal info record to be popped from the stack. We know that there are no states that are represented in the hash table that are not also on the stack. If $q.CPC$ is true then it was set at (1) L8 or L20 as a result of $q.Xend$ also being true or (2) L11 when a `TraversalInfo` q'' was found to be represented in the hash table such that $q''.CPC$ is true and $q''.state$ is a successor of $q.state$. If case (1) holds then we are done. Suppose case (2) holds. Let $\widehat{q''}$ be the successor generated by the algorithm and q'' the `TraversalInfo` record in the hash table. We know that q'' is on the stack at a depth less than q . We also know that $\widehat{q''}.LM$ must be true otherwise $q.Xend$ would be true. Therefore a left-mover transition exists from $q.state$ to $q''.state$. From this we conclude that $Phase(q''.tid, q''.state)$ is false. From Lemma 3 a sequence of left-mover transitions exists to a state represented in a `TraversalInfo` record q' such that $q'.Xend$ is true. This concludes the base case.

For the inductive case let q be the `TraversalInfo` record be record $n + 1$ to be popped from the stack. We assume that the Lemma holds for all n of those `TraversalInfo` records already popped from the stack. Again if $q.CPC$ is true it was set at (1) L8 or L20 as a result of $q.Xend$ also being true or (2) L23 when a successor's CPC bit is true and the successor is about to be popped from the stack or (3) L11 when a `TraversalInfo` q'' was found to be represented in the hash table such that $q''.CPC$ is true and $q''.state$ is a successor of $q.state$. If case (1) holds then we are done. Suppose case (2) holds. If the action executed was not a left-mover then $q.Xend$ would be true. As the successor has been popped before q we assume the induction hypothesis holds. Suppose case (3) holds. Again there

are two cases. Suppose the successor is represented in a `TraversalInfo` record that is not on the stack. As in (2) we see that a left-mover action was executed and we can apply the induction hypothesis. Suppose the successor is represented in a `TraversalInfo` record that is currently held in the stack. Let $\widehat{q''}$ be the successor generated by the algorithm and q'' the `TraversalInfo` record in the hash table. We know that $\widehat{q''}.LM$ must be true otherwise $q.Xend$ would be true. Therefore a left-mover transition exists from $q.state$ to $q''.state$. From this we conclude that $Phase(q''.tid, q''.state)$ is false. So a successor exists and is represented in `TraversalInfo` q'' at a depth lower than q such that $q''.CPC$ is true, $q''.LM$ is true, and $Phase(q''.tid, q''.state)$ is false. By Lemma 3 a sequence of left-mover transitions exists to a state represented in a `TraversalInfo` record such that the `Xend` field is true. We also note that if any state in the sequence is not represented in the stack then the property holds by the application of the inductive hypothesis. This concludes the inductive case.

4.1 L19

Without line L19, the algorithm of figure 8 produces maximal length sequences. The optimization of line L19 schedules other threads at the last commit point on a branch resulting in a shorter but still sound sequence. A commit point is associated with the execution of every action $a \notin RM$. The default commit point is the state where the phase changes. Commit points are only examined when in the post commit phase of a transaction. Thus every transaction has at least one commit point and may have more than one. Along any branch in the post-commit phase of a transaction it is necessary that there exists a path from the last commit point to a state where other threads are scheduled to retain soundness.

Motivation for the addition of line L19 is not immediately obvious. It is the case however that without this optimization, the proposed algorithm performs a selective cycle detection. Consider the following case analysis on the post-commit phase of a transaction.

1. The transaction completes along every branch. Then the transaction begins and ends in a state in $\mathcal{N}(t)$ for thread t . Cycle detection is not needed in this case.
2. The transaction does not branch, contains no actions $a \in LM \setminus RM$ and does not terminate. In this case cycle detection performs as well as the CPC algorithm, except for the existence of `assume(false)` and similar statements.
3. The transaction does not branch, contains one or more actions $a \in LM \setminus RM$ and does not terminate. This is similar to the previous case.
4. The transaction contains $n > 1$ branches, contains no actions $a \in LM \setminus RM$ and does not terminate along any branch. Here there are n cycles in the post-commit phase of the transaction. If other threads are scheduled as required by L18 then this is equivalent to cycle detection and introduces n interleaving points. By adding in L19 the algorithm schedules threads above all cycles in the post-commit phase, thereby introducing only one interleaving point.

5. The transaction contains $n > 1$ branches, contains no actions $a \in LM \setminus RM$ and terminates along one or more branches. Without line L19 the algorithm will introduce an interleaving point for every branch that does not terminate before the terminating branch is discovered **and** will introduce an interleaving point for every branch that revisits a state that is on the stack where the CPC flag is not set. Thus without line L19 the algorithm performs cycle detection and ignores a subset of the cycles. However with L19 exactly one interleaving point is introduced.
6. The transaction contains $n > 1$ branches, contains $m > 1$ actions $a_m \in LM \setminus RM$ and does not terminate along any branch. Here each branch that contains a_m will introduce an interleaving point. However if for all m , a_m is on the same branch then exactly one interleaving point will be introduced.
7. The transaction contains $n > 1$ branches, contains $m > 1$ actions $a_m \in LM \setminus RM$ and terminates along $l \leq n$ of the branches. In this case each branch that contains a_m that does not terminate will introduce exactly one interleaving point. Thus in the worst case $n - l$ interleaving points are introduced, which is the same as without L19. However if $m < l$ then at most $m + 1$ interleaving points are introduced in the worst case. In the best case there could be as few as one interleaving point.

4.2 Soundness proof

Theorem 6. *Let q be a `TraversalInfo` constructed during the execution of the CPC algorithm such that $q.RM = \text{false}$. Then at line L21 there exists a sequence of left-mover transitions of thread $q.tid$ from $q.state$ to (g', ls') and all threads are explored from (g', ls') .*

Proof. $Phase(q.tid, q.state)$ must be false as we have assumed that $q.RM$ is false. If $q.CPC$ is true at line L18 then it remains true as on line L20 the value is updated to `true` || `q.Xend`. Suppose $q.CPC$ is false at line L18. Therefore at line L18 `q.Xend` is set to true. At L20, $q.CPC$ is set to true. Therefore at L21, $q.CPC$ is true. As $q.CPC$ is true and $Phase(q.tid, q.state)$ is false we can apply Lemma 4 directly and conclude that there exists a sequence of left-mover transitions of thread $q.tid$ from $q.state$ to (g', ls') and all threads are explored from (g', ls') .

Lemma 5. *Suppose we have a sequence of n committed transactions from p to q such that at q thread $t(n)$ goes wrong. Then there exists another sequence of committed transactions from p to q' such that at q' thread $t(n)$ goes wrong and each committed transaction $p'_i \xrightarrow{t(i)}^+ p'_{i+1}$ such that $p'_{i+1} \in \mathcal{L}(t(i))$ has no postfix of both-mover transitions.*

Proof. Let t be an indexing function from transaction indices to thread ids. Let the sequence of transactions from p to q be σ . For any sequence of committed transactions there is a fixed number n transactions where $k \leq n$ of those transactions end in a state in $\mathcal{L}(t(i))$ for some thread $t(i)$. We will proceed by creating

a series of transaction sequences $\alpha_1.\beta_1, \alpha_2.\beta_2, \dots, \alpha_k.\beta_k$ where each transaction sequence $\alpha_j.\beta_j$, $0 \leq j \leq k$, ends in a state q' where thread $t(n)$ is wrong. Let $\alpha_0.\beta_0 = \sigma$ where α_0 is the maximal prefix of σ such that for each thread u that executes the transaction of thread u ends in a state in $\mathcal{N}(u)$ and β_0 the remainder of the sequence. Now, for each j , $0 \leq j \leq k$, β_j begins with a transition $p_i \xrightarrow{+}_{t(i)} p_{i+1}$ such that $p_{i+1} \in \mathcal{L}(t(i))$. Let β'_j be the remainder of the transition sequence β_j after p_{i+1} . We know that thread $t(i)$ does not go wrong, nor does it move again in the remainder of the transaction sequence. We also know that at some point in the transaction thread $t(i)$ executed a transition in $LM \setminus RM$ causing the transaction to commit. From the commit transition to the end of the sequence all transitions of thread $t(i)$ have been in LM although some may have also been in RM . Let p'_{i+1} be the first state such that there is a subsequence from $p'_{i+1} \xrightarrow{*}_{t(i)} p_{i+1}$ and every transition in the subsequence is a both-mover. Since these transitions are all both-movers they can be postponed beyond the end of the sequence without any effect on other threads (i.e., if we postpone their execution, or schedule another thread at p'_{i+1} and don't schedule thread $t(i)$ again, thread $t(n)$ is still wrong at the end of the sequence). Let $\beta'_a.\beta'_b = \beta'_j$ where β'_a is the maximal prefix of β'_j where every transaction of some thread u ends in a state in $\mathcal{N}(u)$. Let $\alpha_{j+1} = \alpha_j.p_i \xrightarrow{+}_{t(i)} p'_{i+1}.\beta'_a$ and $\beta_{j+1} = \beta'_b$. When $j = k + 1$, there are no more transactions that end in a state in $\mathcal{L}(u)$ for some thread u so the final sequence is $\alpha_{k+1}.\beta_{k+1}$ and we are done.

Theorem 7 concludes that if there is a state in the multithreaded program where a thread goes wrong that is reachable from the initial state the CPC algorithm will find a state that is reachable from the initial state where that thread goes wrong.

Theorem 7. *If there is an execution of the multithreaded program from (g_0, ls_0) to (g, ls) and a thread t such that $ls(t) = \text{wrong}$, then there is another state (g', ls') where the CPC algorithm visits (g', ls') and $ls'(t) = \text{wrong}$.*

Proof. By Lemma 5 we know that there exists a sequence of committed transactions from (g_0, ls_0) to (g', ls') such that $ls'(t) = \text{wrong}$. We will assume that every state in the sequence is unique. Let t be a function from transaction indices to thread ids. Also, let k be the number of transactions in the sequence and i be the number of the currently explored transaction and n be the number of transitions in the sequence.

We will transform the sequence from Lemma 5 to a sequence that is explored by the CPC algorithm by an iterative process such the final sequence is explored by the CPC algorithm and in the final state $ls(t) = \text{wrong}$.

We will maintain the invariants (1) that for all $0 \leq m \leq n$, α_m is explored by the CPC algorithm and (2) in the final state, $ls(t) = \text{wrong}$. The initial state is explored by the CPC algorithm as shown on line L0. Let α_0 be the sequence containing the initial state and β_0 be the sequence containing the remainder of the sequence from Lemma 5.

For state $m + 1$ in the sequence suppose the CPC algorithm explores α_m such that the m states are represented in `TraversalInfo` records on the stack

and β_m is the remainder of the sequence. Let $t(i)$ be the thread that executed to generate (g, ℓ_s) by executing the m^{th} transition and \mathbf{q} the `TraversalInfo` record such that $\mathbf{q}.\text{state} = (g, \ell_s)$ and $\mathbf{q}.\text{tid} = t(i)$. We will do a case analysis on the state (g, ℓ_s) .

- $(g, \ell_s) \in \mathcal{N}(t)$. If $(g, \ell_s) \in \mathcal{N}(t)$ then it could be the case that $\ell_s(t) = \text{wrong}$ and we are done. Otherwise, $\text{Phase}(\mathbf{q}.\text{tid}, \mathbf{q}.\text{state})$ was false and $\mathbf{q}.\text{LM}$ was also false. In which case at line L7 $\mathbf{q}.\text{Xend}$ is set to true. For all threads $t(j)$ in a state in $\mathcal{N}(t(j))$ thread $t(j)$ is scheduled from $\mathbf{q}.\text{state}$ at line L25. Therefore thread $t(i+1)$ will also be scheduled at $\mathbf{q}.\text{state}$. Suppose thread $t(i+1)$ is the next thread scheduled by the algorithm. Since the transition relation is such that there exists at most one g', ℓ_s' such that $T^\#(t(i+1), g, \ell_s(t(i+1)), j, g', \ell_s')$ the state represented in `TraversalInfo` \mathbf{q}' by executing thread $t(i+1)$ at line L6 must be the next state in the transaction sequence. As the transition executed was from the same transition relation used to obtain the transaction sequence, the transition executed by the CPC algorithm would have the same mover status and therefore be in the same partition. Let β'_m be the remainder of the sequence where $(g, \ell_s) \rightarrow_{t(i+1)} (g', \ell_s'). \beta'_m = \beta_m$. Suppose (g', ℓ_s') is represented in the hash table at line L9 then the state must have been visited before. There are two cases. Suppose (g', ℓ_s') is represented in the stack. Then there exists two states in the sequence that are equivalent. Suppose (g', ℓ_s') is not represented in the stack. Since it has been popped from the stack all successors of that state generated by the CPC algorithm have already been visited. Let α_{m+1} be the execution sequence used when (g', ℓ_s') was visited for the first time. Then let $\beta_{m+1} = \beta'_m$. Otherwise the state is new and would be pushed onto the stack at line L14. In this case we will let $\alpha_{m+1} = \alpha_m.(g, \ell_s) \rightarrow_{t(i+1)} (g', \ell_s')$ and $\beta_{m+1} = \beta'_m$.
- $(g, \ell_s) \in \mathcal{L}(t)$ and thread $t(i+1)$ is the next thread scheduled in β_m . If $(g, \ell_s) \in \mathcal{L}(t)$ then $\text{Phase}(\mathbf{q}.\text{tid}, \mathbf{q}.\text{state})$ is false. Since this is the last state in the transition sequence for this transaction we also know that $\mathbf{q}.\text{RM}$ must be false from Lemma 5. We have assumed that T_G is deterministic so there was at most one transition enabled for $\mathbf{q}.\text{tid}$ from the predecessor of $\mathbf{q}.\text{state}$. If $\mathbf{q}.\text{LM}$ was false there would be a contradiction as $(g, \ell_s) \notin \mathcal{N}(\mathbf{q}.\text{tid})$. If $\mathbf{q}.\text{RM}$ were true then there would be two transitions enabled for $\mathbf{q}.\text{tid}$ from the predecessor of $\mathbf{q}.\text{state}$ or the CPC algorithm would have to apply a different transition relation. Therefore $\mathbf{q}.\text{RM}$ is indeed false. If there are any transitions enabled in thread $t(i)$ then the CPC algorithm will continue to execute transitions of thread $t(i)$ until all successors that are reachable by transitions of thread $t(i)$ have been explored. Suppose the CPC algorithm is about to execute line L18 (i.e., all successors of reachable by thread $t(i)$ have been explored). By Theorem 6 there exists a sequence of left-mover transitions to another state contained in `TraversalInfo` \mathbf{q}' where all threads are scheduled. Call this sequence α'_m . Suppose thread $t(i+1)$ is the next thread to be scheduled. By Lemma 2 the transitions of thread $t(i+1)$ are enabled at $\mathbf{q}'.\text{state}$. Moreover every transition that was enabled before executing the sequence α'_m remains enabled at $\mathbf{q}'.\text{state}$. As the transition relation is deter-

ministic the first transition in β_m is the enabled transition for thread $t(i+1)$ at $\mathbf{q'}.state$. We execute the transition relation for thread $t(i+1)$ to get a new state (g', ls') . Let β'_m be the remainder of β_m after the transition from $\mathbf{q'}.state \rightarrow_{t(i+1)} (g', ls')$. Suppose (g', ls') is represented in the hash table at line L9 then the state must have been visited before. There are two cases. Suppose (g', ls') is represented in the stack. Then there exists two states in the sequence that are equivalent. Suppose (g', ls') is not represented in the stack. Since it has been popped from the stack all successors of that state generated by the CPC algorithm have already been visited. Let α_{m+1} be the execution sequence used when (g', ls') was visited for the first time. Then let $\beta_{m+1} = \beta'_m$. Otherwise the state is new and would be pushed onto the stack at line L14. In this case we will let $\alpha_{m+1} = \alpha_m \cdot \alpha'_m \cdot \mathbf{q'}.state \rightarrow_{t(i+1)} (g', ls')$ and $\beta_{m+1} = \beta'_m$.

- $(g, ls) \in \mathcal{L}(t)$ and $t(i)$ is the next thread scheduled in β_m . If $(g, ls) \in \mathcal{L}(t)$ then $Phase(\mathbf{q}.tid, \mathbf{q}.state)$ is false. We know this state is not the last in the transaction. Since the transition relation is such that there exists at most one g', l' such that $T^\#(t(i), g, ls(t(i)), i, g', l')$ the state represented in **TraversalInfo** $\mathbf{q'}$ and $\mathbf{q'}.state$ must be the next state in the transaction sequence. Suppose (g', ls') is represented in the hash table at line L9 then the state must have been visited before. There are two cases. Suppose (g', ls') is represented in the stack. Then there exists two states in the sequence that are equivalent. Suppose (g', ls') is not represented in the stack. Since it has been popped from the stack all successors of that state generated by the CPC algorithm have already been visited. Let α_{m+1} be the execution sequence used when (g', ls') was visited for the first time. Then let $\beta_{m+1} = \beta'_m$. Otherwise the state is new and would be pushed onto the stack at line L14. In this case we will let $\alpha_{m+1} = \alpha_m \cdot (g, ls) \rightarrow_{\mathbf{q}.tid} (g', ls')$ and $\beta_{m+1} = \beta'_m$.
- $(g, ls) \in \mathcal{R}(t)$. If $(g, ls) \in \mathcal{R}(t)$ then $Phase(\mathbf{q}.tid, \mathbf{q}.state)$ must be true. Since all transactions in the sequence are committed, we know this state is not the last state in the transaction. Since the transition relation is such that there exists at most one g', l' such that $T^\#(t(i), g, ls(t(i)), i, g', l')$ the state represented in **TraversalInfo** $\mathbf{q'}$ and $\mathbf{q'}.state$ must be the next state in the transaction sequence. Suppose (g', ls') is represented in the hash table at line L9 then the state must have been visited before. There are two cases. Suppose (g', ls') is represented in the stack. Then there exists two states in the sequence that are equivalent. Suppose (g', ls') is not represented in the stack. Since it has been popped from the stack all successors of that state generated by the CPC algorithm have already been visited. Let α_{m+1} be the execution sequence used when (g', ls') was visited for the first time. Then let $\beta_{m+1} = \beta'_m$. Otherwise the state is new and would be pushed onto the stack at line L14. In this case we will let $\alpha_{m+1} = \alpha_m \cdot (g, ls) \rightarrow_{\mathbf{q}.tid} (g', ls')$ and $\beta_{m+1} = \beta'_m$.

When $m = n$ α_m is explored by the CPC algorithm and β_m is empty. The final state of α_n differs only as a result of left-mover transitions of threads $t(i) \neq t(k)$.

Example	Loc	Unsound Reduction	CPC	Cycle Detection
AuctionHouse	798	108	108	108
FlowTest	485	4656	4656	4656
Shipping	1844	206	206	222
Conc	392	512	512	2063
Peterson	793	1080	1213	3427
Bluetooth	2768	48109	52092	116559
TransactionManager	6927	1220517	1264894	1268571
AlternatingBit	130	1180	1180	1349
Philosophers	76	87399	87399	428896
Bakery	104	10221	14935	14254

Table 1. Number of states visited by Unsound Reduction, CPC and Cycle Detection algorithms.

Theorem 5 which satisfies the antecedents of Theorem 1 guarantees that the final state (g', ls') is such that $ls'(t) = \textit{wrong}$.

5 Experimental results

We implemented the CPC algorithm in Zing, which is a software model checker being developed in Microsoft Research. Table 1 gives the number of states explored by Zing on various example programs using three variants of the reduction algorithm. The column labeled “Loc” gives the number of lines of code in the Zing program. The column labeled “Unsound Reduction” gives the number of states explored by a reduction algorithm which does not solve the ignoring problem. This gives a lower bound on the number of states that need to be explored by any sound algorithm. The column labeled “CPC” gives the number of states explored by the CPC algorithm. The column labeled “Cycle Detection” gives the number of states explored by a sound algorithm which forcibly ends a transaction whenever a cycle is encountered in the post-commit part of the transaction. The number of states explored is a measure of the running time of the algorithm. The smaller the number of states explored by a sound algorithm, the faster the tool is.

The programs are classified into four groups. The first 3 programs, **AuctionHouse**, **FlowTest** and **Shipping** programs were produced by translating to Zing from a process co-ordination language called BPEL. They represent workflows for business processes, and have mostly acyclic state spaces. In these examples, the number of states explored by the all three algorithm are almost identical.

The next 3 programs **Conc**, **Peterson** and **Bluetooth** were produced by automatic abstraction refinement from concurrent C programs. We have adapted the SLAM toolkit [12] to concurrent programs by using Zing as a back-end model checker instead of Bebop. These examples all have loops that terminate non-deterministically in the abstraction. Thus, the cycle detection algorithm forces

interleaving of all threads in these loops whereas the CPC algorithm avoids interleaving all threads in the loops without losing soundness. The CPC algorithm really shines in comparison with the Cycle Detection algorithm on these examples.

The `TransactionManager` program was obtained from a product group in Microsoft. It was automatically translated to Zing from C#, after a few manual abstractions and manually closing the environment. It is one of the larger Zing examples we currently have. Since the manual abstraction did not result in non-deterministically terminating loops, the CPC algorithm performs only marginally better than the Cycle Detection algorithm.

The final 3 programs, `AlternatingBit`, `Philosophers` and `Bakery` are standard toy examples used by the formal verification community. In the first two examples, CPC performs better than Cycle Detection. In the `Bakery` example we find that the Cycle Detection algorithm performs slightly better than the CPC algorithm. This is possible, since the total number of states is counted over all transactions, and the CPC algorithm gives optimality only within a single transaction. Heuristically, this should translate to smaller number of states explored over all the transactions, but this example shows that this is not always the case.

Overall, the results clearly demonstrate that CPC is a good algorithm for making reduction sound, without forcing the interleaving of other threads in all loops. It generally explores fewer states than Cycle Detection, and out-performs Cycle Detection in examples with nondeterministic loops. Such examples arise commonly from automatic abstraction refinement.

6 Related work

Partial order reduction has numerous variants. The most commonly used ones are stubborn sets of Valmari [2], ample sets [4, 1], and sleep sets [5]. Most of these approaches handle the ignoring problem by using some variant of cycle detection. In another paper, Valmari proposes detecting Strongly Connected Components (SCCs) to solve the ignoring problem [13]. This algorithm from [13] involves detecting terminal strongly connected components, and forces scheduling of other threads from at least one state in each of the terminal strongly connected components (see Algorithm 1.28, Section 5 in [13]). In contrast, the CPC algorithm does not directly compute any strongly connected components. Also the CPC algorithm terminates transactions at fewer points than Valmari's algorithm. Consider the example from Figure 9. In this example, a transaction commits at the state after executing line L0, followed by a non-deterministic branch at line L2. Each of the branches produce terminal SCCs in the state space. Valmari's algorithm appears to force scheduling T2 at each of these terminal SCCs, whereas the CPC algorithm forces scheduling T2 only once, at the commit-point (label L1).

Transaction based reduction was originally developed by Lipton [7]. Work by Stoller and Cohen [10] uses a locking discipline to aggregate transitions into

```

int g = 0;

void T1() {
L0:   g = 1;
L1:   skip;
L2:   if (*) {
L2:       while(true){
L3:           skip;
L4:       }
      else {
L5:           while(true){
L6:               skip;
L7:           }
      }
L8:   return;
}

void T2() {
M0:   assert(g == 0);
M1:   return;
}

P = { T1() } || { T2() }

```

Fig. 9. Distinction between CPC algorithm and SCC-based algorithms

a sequence of transitions that may be viewed atomically. Flanagan and Qadeer augment this approach with right movers to get further reduction [9]. This idea is combined with procedure summarization by Qadeer, Rajamani, and Rehof in [8]. As mentioned earlier, all of these papers address the ignoring problem only indirectly by disallowing certain types of infinite executions, such as those consisting of only internal hidden actions, within each thread (see Condition **C** from Section 4.2 in [9] which forbids the transaction from having infinite executions after committing, but without completing, and well-formedness assumption **Wf-infinite-invis** from Section 4 in [10]). It is not clear how these assumptions are enforced. Two of the above papers [9, 8] do not have any accompanying implementation, and it is unclear how the ignoring problem is solved in the implementation associated with [10]. Our guess is that they use some form of cycle detection.

The Verisoft [6] implementation does not use the detection of cycles or strongly connected components, rather a timeout is used to detect an infinite execution that is local to a particular process. Other cycles are broken by limiting the search depth or using a driver that generates a finite number of external events. Dwyer et al [14] use the notion of a locking discipline is used to increase the number of transitions that can form an ample set for a process. The algorithms presented use the standard cycle detection technique to insure soundness.

7 Conclusion

Partial order reduction methods with ample sets usually use Cycle Detection to solve the ignoring problem. In the context of transaction based reduction, we propose a new technique called Commit Point Completion (CPC) to solve the ignoring problem. We have proved that this algorithm is correct, and have implemented it in the Zing model checker. Our experimental results demonstrate that with transaction based reduction, the CPC algorithm performs better than Cycle Detection. Though the CPC algorithm was presented using the terminology of Lipton's transactions, we believe that the idea is applicable to other variants of partial order reduction as well. Exploration of this idea is left to future work. The ignoring problem also arises when we attempt to build summaries for multithreaded programs[8]. Though not mentioned here, our implementation of summaries in Zing also uses the core idea of the CPC algorithm to ensure soundness.

References

1. Clarke, E., Grumberg, O., Peled, D.: Model Checking. MIT Press (1999)
2. Valmari, A.: A stubborn attack on state explosion. In: CAV 91: Computer Aided Verification, Springer-Verlag (1991) 156–165
3. Holzmann, G., Peled, D.: An improvement in formal verification. In: FORTE 94: Formal Description Techniques, Chapman & Hall (1994) 197–211
4. Peled, D.: Partial order reduction: Model-checking using representatives. In: MFCS 96: Mathematical Foundations of Computer Science, Springer-Verlag (1996) 93–112
5. Godefroid, P.: Partial-Order Methods for the Verification of Concurrent Systems: An Approach to the State-Explosion Problem. LNCS 1032. Springer-Verlag (1996)
6. Godefroid, P.: Model checking for programming languages using Verisoft. In: POPL 97: Principles of Programming Languages. (1997) 174–186
7. Lipton, R.J.: Reduction: A method of proving properties of parallel programs. In: Communications of the ACM. Volume 18:12. (1975) 717–721
8. Qadeer, S., Rajamani, S.K., Rehof, J.: Summarizing procedures in concurrent programs. In: Principles of Programming Languages, ACM (2004) 245–255
9. Flanagan, C., Qadeer, S.: Transactions for software model checking. In: SoftMC 03: Software Model Checking Workshop. (2003)
10. Stoller, S.D., Cohen, E.: Optimistic synchronization-based state-space reduction. In: TACAS 03. LNCS 2619, Springer-Verlag (2003) 489–504
11. Flanagan, C., Qadeer, S.: Types for atomicity. In: TLDI 03: Types in Language Design and Implementation, ACM (2003) 1–12
12. Ball, T., Rajamani, S.K.: The SLAM project: Debugging system software via static analysis. In: POPL 02: Principles of Programming Languages, ACM (2002) 1–3
13. Valmari, A.: Stubborn sets for reduced state space generation. In: Advances in Petrinets. LNCS 483, Springer-Verlag (1990)
14. Dwyer, M.B., Hatcliff, J., Robby, Ranganath, V.P.: Exploiting object escape and locking information in partial-order reductions for concurrent object-oriented programs. Formal Methods in System Design **25** (2004) 199–240