

Thread-Modular Model Checking

Cormac Flanagan¹ and Shaz Qadeer²

¹ Systems Research Center, HP Labs, 1501 Page Mill Road, Palo Alto, CA 94304

² Microsoft Research, One Microsoft Way, Redmond, WA 98052

Abstract. We present thread-modular model checking, a novel technique for verifying correctness properties of *loosely-coupled* multithreaded software systems. Thread-modular model checking verifies each thread separately using an automatically inferred environment assumption that abstracts the possible steps of other threads. Separate verification of each thread yields significant space and time savings. Suppose there are n threads, each with a local store of size L , where the threads communicate via a shared global store of size G . If each thread is finite-state (without a stack), the naive model checking algorithm requires $O(G.L^n)$ space, whereas thread-modular model checking requires only $O(n.G.(G + L))$ space. If each thread has a stack, the general model checking problem is undecidable, but thread-modular model checking terminates in polynomial time.

1 Introduction

Designing correct multithreaded software is difficult due to subtle interactions among threads operating concurrently on shared data. Errors in such systems are easy to introduce but difficult to diagnose and fix. Model checking [CE81, QS81] is a promising technique for verifying correctness properties of multithreaded software systems. However, due to the large state spaces of such systems, they are difficult to model check. In this paper, we present a novel technique called thread-modular model checking to alleviate the problem of exploring large state spaces of multithreaded software.

We consider multithreaded software systems with a finite number of threads where the shared global store and the local store of each thread are finite. However, each thread also has an unbounded stack which allows us to model procedure calls and recursion. We focus on the verification of safety properties such as assertions and global invariants. Verification of such safety properties can be reduced to the problem of checking whether an error state is reachable from the system's initial state. This problem is undecidable [HU79, Ram00] in general. Thread-modular model checking is a conservative (sound and incomplete) algorithm for this problem that is powerful enough to verify a variety of multithreaded software systems occurring in practice.

Thread-modular reasoning for shared-memory programs was first introduced by Jones [Jon83]. The basic idea behind this technique is to verify each thread separately using an environment assumption to model interleaved steps of the

other threads. The environment assumption of each thread is a binary relation over the set of global stores, and includes all global store updates that may be performed by other threads.

In earlier work, we extended the proof rule of Jones and implemented it in the Calvin checker [FFQ02,FQS02] for multithreaded Java programs. Our experience using Calvin indicates that the threads in most software systems are *loosely-coupled*, *i.e.*, there is little correlation among the local states of the various threads, and thread-modular reasoning is sufficiently powerful to verify these systems. However, a significant cost of using Calvin is that the programmer is required to provide the appropriate environment assumption. The thread-modular model checking technique in this paper avoids this cost by automatically inferring these environment assumptions.

Thread-modular model checking infers the environment assumption for each thread by first inferring a *guarantee* for each thread, which models all global store updates performed by that thread. The environment assumption of a thread is then the disjunction of the guarantees of all the other threads. The guarantee of each thread is initially the empty relation, and is iteratively extended during the model checking process. Each thread is verified using the standard algorithm for model checking a sequential pushdown system except that at each control point of the thread, the global state is allowed to mutate according to the guarantees of the other threads. In addition, whenever a thread modifies the global store, that transition on the global states is added to that thread's guarantee. The iteration continues until the reachable state space and guarantee of each thread converges. The complexity of this procedure is $O(n.G^3.L^3.F)$, where n is the number of threads, F is the number of stack symbols, G is the size of the global store, and L is the size of local store per thread.

Even if the threads do not have a stack and are consequently finite-state, thread-modular model checking offers significant savings over standard model checking. The naive model checking algorithm explicitly models the program counters of all threads. Therefore, it explores all interleavings of the various threads and its complexity is exponential in the number of threads. However, thread-modular model checking verifies each thread separately and its complexity $O(n.G^2.L.(n + L))$ is significantly better than that of the naive algorithm.

1.1 Example

To illustrate the benefits of thread-modular model checking, we consider its application to a simple multithreaded program. The multithreaded program *Simple*(n) has n threads which are concurrently executing the procedure `p`. Each thread is identified by unique integer value from the set $Tid = \{1, \dots, n\}$. These threads manipulate a shared integer variable `x` initialized to 1. The variable `x` is protected by a mutex `m`, which is either the (non-zero) identifier of the thread holding the lock, or else 0, if the lock is not held by any thread. Thus, the type $Mutex = \{0\} \cup Tid$. The mutex `m` is manipulated by two operations, `acquire` and `release`. The operation `acquire` blocks until `m = 0` and then atomically sets `m` to `tid`, the identifier of the current thread. The operation `release` sets `m`

back to 0. For each thread, there is an implicit local variable called `pc`, which is the program counter of the thread. The variable `pc` takes values from the set $\text{Loc} = \{1, \dots, 6\}$ of control locations. We denote the program counter of thread tid by $\text{pc}[tid]$.

A simple multithreaded program

```
int x := 1;
```

```
void p() {
1:  acquire;
2:  x := 0;
3:  x := x + 1;
4:  assert x > 0;
5:  release;
6: }
```

$$\text{Simple}(n) = \underbrace{\text{p}() \mid \dots \mid \text{p}()}_n$$

We would like to verify three correctness properties for the program $\text{Simple}(n)$. A correctness property is given by a set of error states; the program satisfies the correctness property if no error state is reachable.

1. There are no races on the data variable `x`. The error set is

$$\exists i, j \in \text{Tid}. i \neq j \wedge \text{pc}[i] \in \{2, 3, 4\} \wedge \text{pc}[j] \in \{2, 3, 4\} .$$

2. The assertion at control location 4 does not fail for any thread. The error set is

$$\exists i \in \text{Tid}. \text{pc}[i] = 4 \wedge x \leq 0 .$$

3. Every reachable state satisfies the invariant $m = 0 \Rightarrow x = 1$. The error set is

$$m = 0 \wedge x \neq 1 .$$

Thread-modular model checking can verify these correctness properties. Our algorithm computes the guarantee

$$\mathcal{G} \subseteq \text{Tid} \times (\text{Mutex} \times \text{int}) \times (\text{Mutex} \times \text{int})$$

where $\text{Mutex} \times \text{int}$ is the set of all global stores, and the thread-local reachable set

$$\mathcal{R} \subseteq \text{Tid} \times (\text{Mutex} \times \text{int}) \times \text{Loc} .$$

The set \mathcal{G} has the property that if the thread with identifier tid ever takes a step in which the pair (\mathbf{m}, \mathbf{x}) of global variables is modified from (m_1, x_1) to (m_2, x_2) , then $(tid, (m_1, x_1), (m_2, x_2)) \in \mathcal{G}$. The set \mathcal{R} has the property that if there is a reachable state in which the pair (\mathbf{m}, \mathbf{x}) has the value (m, v) and the program counter of thread with identifier tid has the value pc , then $(tid, (m, x), pc) \in \mathcal{R}$.

These sets are given by the following predicates:

$$\begin{aligned}
& \vee \mathbf{m} = 0 \wedge \mathbf{m}' = \mathit{tid} \wedge \mathbf{x} = \mathbf{x}' = 1 \\
\mathcal{G} \stackrel{\text{def}}{=} & \vee \mathbf{m} = \mathit{tid} \wedge \mathbf{m}' = 0 \wedge \mathbf{x} = \mathbf{x}' = 1 \\
& \vee \mathbf{m} = \mathbf{m}' = \mathit{tid} \wedge \mathbf{x} = 0 \wedge \mathbf{x}' = 1 \\
& \vee \mathbf{m} = \mathbf{m}' = \mathit{tid} \wedge \mathbf{x} = 1 \wedge \mathbf{x}' \in \{0, 1\} \\
& \vee \text{pc}[\mathit{tid}] \in \{1, 6\} \wedge \mathbf{m} = 0 \wedge \mathbf{x} = 1 \\
\mathcal{R} \stackrel{\text{def}}{=} & \vee \text{pc}[\mathit{tid}] \in \{1, 6\} \wedge \mathbf{m} \in \mathit{Tid} \setminus \{\mathit{tid}\} \wedge \mathbf{x} \in \{0, 1\} \\
& \vee \text{pc}[\mathit{tid}] \in \{2, 4, 5\} \wedge \mathbf{m} = \mathit{tid} \wedge \mathbf{x} = 1 \\
& \vee \text{pc}[\mathit{tid}] = 3 \wedge \mathbf{m} = \mathit{tid} \wedge \mathbf{x} = 0
\end{aligned}$$

The environment assumption of the thread tid can be computed from the guarantee as follows:

$$\mathcal{E}(\mathit{tid}) \stackrel{\text{def}}{=} \exists t \in \mathit{Tid} : t \neq \mathit{tid} \wedge \mathcal{G}[\mathit{tid} := t]$$

An examination of \mathcal{R} proves that $\mathit{Simple}(n)$ satisfies its three correctness properties:

1. The thread with identifier tid accesses \mathbf{x} only when $\text{pc}[\mathit{tid}] \in \{2, 3, 4\}$. Every member of \mathcal{R} satisfies the property that if $\text{pc}[\mathit{tid}] \in \{2, 3, 4\}$ then $\mathbf{m} = \mathit{tid}$. Therefore, it is impossible for two different threads to be at a control location in $\{2, 3, 4\}$ simultaneously. Consequently, there is no race on the variable \mathbf{x} .
2. Every member of \mathcal{R} satisfies the property that $\mathbf{x} = 1$ when $\text{pc} = 4$. Therefore, the assertion at control location 4 holds.
3. Every member of \mathcal{R} satisfies the condition $\mathbf{m} = 0 \Rightarrow \mathbf{x} = 1$, which is therefore an invariant of $\mathit{Simple}(n)$.

To verify the program $\mathit{Simple}(n)$, the thread-modular model checking algorithm analyzes each thread separately. When analyzing thread tid , each global state stored by the algorithm contains values for \mathbf{m} , \mathbf{x} , and the program counter of thread tid . The algorithm explores $O(n)$ states and transitions for each thread. Since there are n threads, the number of explored states and transitions is $O(n^2)$.

On the other hand, each state stored by a naive model checking algorithm will provide values for \mathbf{m} , \mathbf{x} , and the program counters of all the threads. Consequently, the number of states and transitions explored are $O(2^n)$. Thus, for this example, the thread-modular model checking algorithm provides exponential savings in the time and space required for state-space enumeration.

1.2 Limitations

Thread-modular model checking is a sound but incomplete verification algorithm. In the program $\mathit{Simple}(n)$, our model of the mutex \mathbf{m} is crucial for the success of thread-modular model checking. Suppose we model the mutex \mathbf{m} as a boolean instead, such that \mathbf{m} is *true* when locked and *false* when unlocked. The operation `acquire` blocks until $\mathbf{m} = \mathit{false}$ and then atomically sets \mathbf{m} to *true*. The

operation `release` sets m back to *false*. Then, thread-modular model checking computes the following sets for \mathcal{G} and \mathcal{R} :

$$\begin{aligned} \mathcal{G} &\stackrel{\text{def}}{=} \begin{aligned} &\vee m = \textit{false} \wedge m' = \textit{true} \wedge \mathbf{x} = \mathbf{x}' \in \{\textit{false}, \textit{true}\} \\ &\vee m = \{\textit{false}, \textit{true}\} \wedge m' = \textit{false} \wedge \mathbf{x} = \mathbf{x}' = \{\textit{false}, \textit{true}\} \\ &\vee m = m' \wedge \mathbf{x} \in \{0, 1\} \wedge \mathbf{x}' \in \{0, 1\} \end{aligned} \\ \mathcal{R} &\stackrel{\text{def}}{=} \mathbf{x} \in \{0, 1\} \end{aligned}$$

With the \mathcal{R} computed above, we cannot prove any of the three correctness properties from Section 1.1. A reason for the incompleteness of thread-modular model checking is that the computed thread guarantee is represented only in terms of updates to the shared variables. The guarantee does not keep track of either the sequence in which actions updating shared variables are performed or the number of times each such action is executed. Often, this information is necessary for verifying the program.

Although thread-modular reasoning is incomplete in general, we can verify any correctness property with it by providing auxiliary information in the program. This information is provided by introducing auxiliary shared variables to the program and adding appropriate updates to them. In *Simple(n)*, we provided this information by storing in the mutex m the identifier of the current holder of the mutex. A different way to achieve the same effect is to introduce for each thread tid , a boolean *shared* variable $\text{apc}[tid]$ representing an abstraction of the program counter $\text{pc}[tid]$. The variable $\text{apc}[tid]$ is initialized to *false*, an `acquire` operation by thread tid sets $\text{apc}[tid]$ to *true*, and a `release` operation by thread tid sets $\text{apc}[tid]$ to *false*.

1.3 Related work

We refer the reader to our earlier papers [FFQ02,FQS02] for a discussion of the related work on verification of multithreaded software by compositional reasoning and model checking.

Cobleigh et al. [CGP03] share our motivation of reducing the annotation cost of compositional reasoning. They use a counterexample-guided learning algorithm to infer environment assumptions, an approach that is very different from ours. Our algorithm is based entirely on model checking; the correctness properties of the program are verified and appropriate environment assumptions are inferred solely by state-space enumeration.

Bouajjani et al. [BET03] present a generic approach to the static analysis of concurrent programs. Unlike our work on shared-memory programs, they focus on synchronous message-passing programs. They present several abstractions for such programs, including a commutative abstraction that ignores message order. Our thread-modular model checking algorithm focuses on concurrent updates to the shared heap, but has a similar flavor to this commutative abstraction, in that it ignores the order of heap updates performed by a thread.

2 Concurrent finite-state systems

A *concurrent finite state system* consists of a number of concurrently executing threads. The threads communicate through a global store, which is shared by all threads. In addition, each thread has its own local store containing data not manipulated by other threads, such as the program counter of the thread. Each thread also has an associated thread identifier. A state of the system consists of a global store g and a mapping ls from thread identifiers to local stores. We use the notation $ls[t := l]$ to denote a mapping that is identical to ls except that it maps thread identifier t to local store l .

Domains

$$\begin{array}{l}
 t, e \in \quad Tid = \{1, \dots, n\} \\
 g \in \quad GlobalStore \\
 l \in \quad LocalStore \\
 ls \in \quad LocalStores = Tid \rightarrow LocalStore \\
 \Sigma \in \quad State = GlobalStore \times LocalStores
 \end{array}$$

We model the behavior of the individual threads as the transition relation T :

$$T \subseteq Tid \times (GlobalStore \times LocalStore) \times (GlobalStore \times LocalStore)$$

The relation $T(t, g, l, g', l')$ holds if the thread t can take a step from a state with global store g and where thread t has local store l , yielding a new state with global and local stores g' and l' , respectively.

We assume that program execution starts in an initial state $\Sigma_0 = (g_0, ls_0)$ consisting of an initial global store g_0 and a mapping ls_0 that provides the initial local store for each thread. The correctness condition for the program in our system is provided by an *error set* $E \subseteq GlobalStore \times LocalStores$. A state (g, ls) is *erroneous* if $E(g, ls)$ is true. Our goal is to determine if, when started from the initial state Σ_0 , the system can reach an erroneous state.

2.1 Standard model checking

Since the set of possible states is finite, we can use standard model checking to determine if any erroneous state is reachable from the initial state. In particular, the least solution $R \subseteq State$ to the following inference rules describes the set of reachable states.

Standard model checking

$$\begin{array}{c}
 \text{(BASIC INIT)} \qquad \qquad \qquad \text{(BASIC STEP)} \\
 \frac{}{R(g_0, ls_0)} \qquad \qquad \frac{R(g, ls) \quad T(t, g, ls(t), g', l')}{R(g', ls[t := l'])}
 \end{array}$$

Although we provide a declarative definition of R here, it is easily computed using a worklist-based algorithm. Having computed R , it is straightforward to

determine if any erroneous state is reachable, *i.e.*, if there exist t , g , and ls such that $R(g, ls) \wedge E(t, g, ls)$.

Unfortunately, the computational cost of this algorithm becomes excessive in the presence of multiple threads. Let $n = |Tid|$ be the number of threads and let $G = |GlobalStore|$ and $L = |LocalStore|$ be the sizes of the global and local stores, respectively. Then the size of R and the space complexity of this algorithm is $O(G.L^n)$. Furthermore, for each entry in R there may be $n.G.L$ applications of (BASIC STEP). Hence the time complexity of this algorithm is $O(n.G^2.L^{n+1})$. A more accurate time complexity can be obtained by accounting for the bounded nondeterminism of the transition relation of each thread. Let d be the bound on the number of (g', l') pairs for any thread t , global store g , and local store l such that $T(t, g, l, g', l')$ holds. Then, for each entry in R , there are at most $n.d$ applications of (BASIC STEP) and the time complexity is $O(n.d.G.L^n)$.

2.2 Thread-modular model checking

The complexity of standard model checking is exponential in the number of threads, since it explicitly correlates the local states (and program counters) of all the various threads. However, since the threads in most software systems are predominantly loosely-coupled, this correlation is largely redundant. Thread-modular model checking provides a means to avoid this redundancy.

Under thread-modular model checking, each thread is checked separately, using the guarantees that abstract the behavior of interleaved steps of other threads. The algorithm works by computing two relations: \mathcal{R} , which specifies the reachable states of each thread, and \mathcal{G} , which is the guarantee of each thread. Thus, the guarantee is inferred automatically during the model checking process.

$$\begin{aligned}\mathcal{R} &\subseteq Tid \times GlobalStore \times LocalStore \\ \mathcal{G} &\subseteq Tid \times GlobalStore \times GlobalStore\end{aligned}$$

The relation $\mathcal{R}(t, g, l)$ holds if the system can reach a state with global store g and where the thread t has local store l . Similarly, $\mathcal{G}(t, g, g')$ holds if a step by thread t can go from a reachable state with global store g to a state with global store g' . While model checking a thread with identifier different from t , we know that whenever the global store is g and $\mathcal{G}(t, g, g')$ holds, an interleaved step of thread t can change the global store to g' . The relations \mathcal{R} and \mathcal{G} are defined as the least solution to the following rules.

Thread-modular model checking		
(AG INIT)	(AG ENV)	(AG STEP)
$\frac{}{\mathcal{R}(t, g_0, ls(t))}$	$\frac{\mathcal{R}(t, g, l) \quad \mathcal{G}(e, g, g') \quad t \neq e}{\mathcal{R}(t, g', l)}$	$\frac{\mathcal{R}(t, g, l) \quad T(t, g, l, g', l')}{\mathcal{R}(t, g', l') \quad \mathcal{G}(t, g, g')}$

The set of reachable states determined using thread-modular reasoning is a conservative approximation of the set of actual reachable states, as illustrated by the following lemma.

Lemma 1. *For all global stores g and local store maps ls , if $R(g, ls)$ then for all thread identifiers t , $\mathcal{R}(t, g, ls(t))$.*

Our algorithm reports an error if there is an erroneous state (g, ls) such that $\mathcal{R}(t, g, ls(t))$ for all $t \in Tid$. If a software error causes an erroneous state to be reachable, *i.e.*,

$$\exists g, ls. (E(g, ls) \wedge R(g, ls)) ,$$

then the thread-modular algorithm will catch that error, *i.e.*,

$$\exists g, ls. E(g, ls) \wedge \forall t. \mathcal{R}(t, g, ls(t)) .$$

Thread-modular model checking can be performed using a worklist-based algorithm, whose complexity is much less than that of standard model checking. The space complexity is $O(n.G.(G + L))$. There may be $n^2.G^2.L$ applications of (AG ENV) and $n.G^2.L^2$ applications of (AG STEP). Hence the time complexity of this algorithm is $O(n.G^2.L.(n + L))$. Again, we improve the bound to $O(n.G.L.(n.G + d))$ using the bound d on the nondeterminism of the transition relations of the thread.

3 Concurrent pushdown systems

The thread-modular approach described so far works well for checking multi-threaded finite state software systems. However, its applicability to realistic systems is somewhat limited, because such systems are typically constructed using procedures and procedure calls, and hence rely on the presence of an unbounded stack for each thread. In this section, we extend our thread-modular approach to handle such systems.

We assume that, in addition to a local store, each thread now also has its own private stack, which is sequence of frames. We leave the exact structure of each frame unspecified, but it might contain, for example, the return address for a procedure call. A state of the concurrent pushdown system consists of a global store, a collection of local stores, one for each thread, and a collection of stacks, one for each thread.

Domains

$$\begin{array}{l} f \in \text{Frame} \\ s \in \text{Stack} = \text{Frame}^* \\ ss \in \text{Stacks} = \text{Tid} \rightarrow \text{Stack} \\ \Sigma \in \text{State} = \text{GlobalStore} \times \text{LocalStores} \times \text{Stacks} \end{array}$$

We model the behavior of the individual threads using three relations:

$$\begin{array}{l} T \subseteq \text{Tid} \times (\text{GlobalStore} \times \text{LocalStore}) \times (\text{GlobalStore} \times \text{LocalStore}) \\ T^+ \subseteq \text{Tid} \times (\text{GlobalStore} \times \text{LocalStore}) \times (\text{LocalStore} \times \text{Frame}) \\ T^- \subseteq \text{Tid} \times (\text{GlobalStore} \times \text{LocalStore} \times \text{Frame}) \times \text{LocalStore} \end{array}$$

The relation T models thread steps that do not manipulate the stack. The relation $T(t, g, l, g', l')$ holds if the thread t can take a step from a state with global and local stores g and l , respectively, yielding (possibly modified) stores g' and l' , and where the stack is not accessed or updated during this step. The relation $T^+(t, g, l, l', f)$ models steps of thread t that push a frame onto the stack. The global and local stores are initially g and l , the global store is unmodified during this step, the local store is updated to l' , and the frame f is pushed onto the stack. Similarly, the relation $T^-(t, g, l, f, l')$ models steps of thread t that pop a frame from the stack. The global and local stores are initially g and l and the frame f is initially on top of the stack. After the step, the global store is unmodified, the local store is updated to l' , and the frame f has been popped from the stack.

The correctness condition is still specified by an error set $E \subseteq \text{GlobalStore} \times \text{LocalStores}$. Note that although the error set depends on the local stores of the threads, it does not depend on their stacks. A state (g, ls, ss) is erroneous if $(g, ls) \in E$.

We assume that all stacks are empty in the initial state, and let ss_0 map each thread identifier to the empty stack. The set of reachable states is then defined by the least relation $R \subseteq \text{State}$ satisfying the following rules.

Basic PDA model checking

(BASIC PDA INIT)	(BASIC PDA STEP)
$\frac{}{R(g_0, ls_0, ss_0)}$	$\frac{R(g, ls, ss) \quad T(t, g, ls(t), g', l')}{R(g', ls[t := l'], ss)}$
(BASIC PDA PUSH)	(BASIC PDA POP)
$\frac{R(g, ls, ss) \quad T^+(t, g, ls(t), l', f)}{R(g', ls[t := l'], ss[t := ss(t).f])}$	$\frac{R(g, ls, ss) \quad ss(t) = s.f \quad T^-(t, g, ls(t), f, l')}{R(g, ls[t := l'], ss[t := s])}$

Since the stack sizes are unbounded, the set of reachable states may also be unbounded. Consequently, any algorithm to compute R may diverge. In fact, the model checking problem for concurrent pushdown systems is undecidable, a result that can be proved by reduction from the undecidable problem of determining if the intersection of two context-free languages is empty [Ram00].

3.1 Thread-modular model checking

Although sound and complete model checking of concurrent pushdown systems is undecidable, thread-modular reasoning allows us to model check such systems in a conservative yet useful manner. Again, we model check each thread separately, using the guarantees to reason about the effect of interleaved steps of other threads. The algorithm works by computing the guarantee relation \mathcal{G} and the reachability relations \mathcal{P} and \mathcal{Q} .

$$\begin{aligned} \mathcal{G} &\subseteq \text{ThreadId} \times \text{GlobalStore} \times \text{GlobalStore} \\ \mathcal{P} &\subseteq \text{ThreadId} \times \text{GlobalStore} \times \text{LocalStore} \times \text{GlobalStore} \times \text{LocalStore} \\ \mathcal{Q} &\subseteq \text{ThreadId} \times \text{GlobalStore} \times \text{LocalStore} \times \text{Frame} \times \text{GlobalStore} \times \text{LocalStore} \end{aligned}$$

The guarantee $\mathcal{G}(t, g, g')$ holds if a step by thread t can go from a reachable state with global store g to a state with global store g' . The reachability relation $\mathcal{P}(t, g, l, g', l')$ holds if (1) the system can reach a state with global store g and where thread t has local store l , and (2) from any such state, the system can later reach a state with global store g' and where thread t has local store l' , and where the stack is identical to that in the first state. Similarly, the reachability relation $\mathcal{Q}(t, g, l, f, g', l')$ holds if (1) the system can reach a state with global store g and where thread t has local store l , and (2) from any such state, the system can later reach a state with global store g' and where thread t has local store l' , and where the stack is identical to that in the first state except that the frame f has been added to it. These relations are defined as the least solution to the following rules.

Thread-modular PDA model checking

(AG PDA INIT)	
$\overline{\mathcal{P}(t, g_0, ls_0(t), g_0, ls_0(t))}$	
(AG PDA ENV1)	(AG PDA ENV2)
$\frac{\mathcal{P}(t, g_1, l_1, g_2, l_2) \quad \mathcal{G}(e, g_2, g_3) \quad e \neq t}{\mathcal{P}(t, g_1, l_1, g_3, l_2)}$	$\frac{\mathcal{Q}(t, g_1, l_1, f, g_2, l_2) \quad \mathcal{G}(e, g_2, g_3) \quad e \neq t}{\mathcal{Q}(t, g_1, l_1, f, g_3, l_2)}$
(AG PDA STEP1)	(AG PDA PUSH)
$\frac{\mathcal{P}(t, g_1, l_1, g_2, l_2) \quad T(t, g_2, l_2, g_3, l_3)}{\mathcal{P}(t, g_1, l_1, g_3, l_3) \quad \mathcal{G}(t, g_2, g_3)}$	$\frac{\mathcal{P}(t, g_1, l_1, g_2, l_2) \quad T^+(t, g_2, l_2, l_3, f)}{\mathcal{Q}(t, g_1, l_1, f, g_2, l_3) \quad \mathcal{P}(t, g_2, l_3, g_2, l_3)}$
(AG PDA STEP2)	(AG PDA POP)
$\frac{\mathcal{Q}(t, g_1, l_1, f, g_2, l_2) \quad \mathcal{P}(t, g_2, l_2, g_3, l_3)}{\mathcal{Q}(t, g_1, l_1, f, g_3, l_3)}$	$\frac{\mathcal{Q}(t, g_1, l_1, f, g_2, l_2) \quad T^-(t, g_2, l_2, f, l_3)}{\mathcal{P}(t, g_1, l_1, g_2, l_3)}$

The set of reachable states determined using thread-modular reasoning is a conservative approximation of the set of actual reachable states, as illustrated by the following lemma.

Lemma 2. *For all global stores g and local store maps ls and stack maps ss , if $R(g, ls, ss)$ then for all thread identifiers t , there exists some g', l' such that $\mathcal{P}(t, g', l', g, ls(t))$.*

Our algorithm reports an error if there is $(g, ls) \in E$ such that for all $t \in Tid$, there is a global store g' and a local store l' with $\mathcal{P}(t, g', l', g, ls(t))$. If a software error causes an erroneous state to be reachable, *i.e.*,

$$\exists g, ls, ss. (E(g, ls) \wedge R(g, ls, ss))$$

then the thread-modular algorithm will catch that error, *i.e.*,

$$\exists g, ls. E(g, ls) \wedge \forall t. \exists g', l'. \mathcal{P}(t, g', l', g, ls(t)) .$$

Let $F = |Frame|$. Then, the space complexity of this algorithm is $O(n.G^2.L^2.F)$. The time complexity of this algorithm is $O(n^2.G^3.L^3.F)$ since each inference rule can be applied at most $n^2.G^3.L^3.F$ times.

4 Discussion

We have presented a new technique called thread-modular model checking for verifying multithreaded software systems. Although incomplete for general systems, this technique is particularly effective for loosely-coupled multithreaded software where the various threads synchronize using primitives such as mutexes, readers-writer locks, etc. If the synchronization primitives are modeled with appropriate auxiliary information, these systems can be verified one thread at a time.

Realistic software systems often have dynamic thread creation that may lead to unbounded number of threads. This aspect of multithreaded software is currently not handled by our algorithm. However, the set of thread identifiers, even if infinite, is a scalarset type [ID96]. Consequently, these systems are amenable to symmetry reduction which we plan to exploit in future work.

The thread-modular model checking algorithm constructs a particular abstraction of multithreaded software using environment assumptions. However, the abstraction might be too coarse to verify the relevant correctness property. If the algorithm reports an error, we would like an efficient procedure to check whether the violation is real or introduced due to the abstraction process. In the second case, we would like to automatically refine the environment assumptions by possibly explicating some aspect of the program counters of the other threads in the environment. After the refinement, the model checking algorithm can be repeated. Thus, the thread-modular model checking algorithm may be converted to a semi-algorithm that is sound and also complete on termination.

References

- [BET03] A. Bouajjani, J. Esparza, and T. Touili. A generic approach to the static analysis of concurrent programs with procedures. In *POPL 03: Principles of Programming Languages*, 2003. to appear.
- [CE81] E.M. Clarke and E.A. Emerson. Design and synthesis of synchronization skeletons using branching-time temporal logic. In *Workshop on Logic of Programs*, Lecture Notes in Computer Science 131, pages 52–71. Springer-Verlag, 1981.
- [CGP03] J.M. Cobleigh, D. Giannakopoulou, and C.S. Păsăreanu. Learning assumptions for compositional verification. In *TACAS 03: Tools and Algorithms for the Construction and Analysis of Systems*, 2003. to appear.
- [FFQ02] C. Flanagan, S.N. Freund, and S. Qadeer. Thread-modular verification for shared-memory programs. In *ESOP 02: European Symposium on Programming*, Lecture Notes in Computer Science 2305, pages 262–277, 2002.
- [FQS02] C. Flanagan, S. Qadeer, and S. Seshia. A modular checker for multithreaded programs. In *CAV 02: Computer Aided Verification*, Lecture Notes in Computer Science 2404, pages 180–194, 2002.

- [HU79] J.E. Hopcroft and J.D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley Publishing Company, 1979.
- [ID96] C.N. Ip and D.L. Dill. Better verification through symmetry. *Formal Methods in System Design*, 9(1–2):41–75, 1996.
- [Jon83] C. B. Jones. Tentative steps toward a development method for interfering programs. *ACM Transactions on Programming Languages and Systems*, 5(4):596–619, 1983.
- [QS81] J. Queille and J. Sifakis. Specification and verification of concurrent systems in CESAR. In *Fifth International Symposium on Programming*, Lecture Notes in Computer Science 137, pages 337–351. Springer-Verlag, 1981.
- [Ram00] G. Ramalingam. Context-sensitive synchronization-sensitive analysis is undecidable. *ACM Transactions on Programming Languages and Systems*, 22(2):416–430, 2000.