

Modular Verification of Concurrent Programs via Sequential Model Checking

Dan Rasin¹, Orna Grumberg¹, and Sharon Shoham²

¹ Technion – Israel Institute of Technology

² Tel-Aviv University

Abstract. This work utilizes the plethora of work on verification of sequential programs for the purpose of verifying concurrent programs. We reduce the verification of a concurrent program to a series of verification tasks of sequential programs. Our approach is modular in the sense that each sequential verification task roughly corresponds to the verification of a single thread, with some additional information about the environment in which it operates. Information regarding the environment is gathered during the run of the algorithm, by need. While our approach is general, it specializes on concurrent programs where the threads are structured hierarchically. The idea is to exploit the hierarchy in order to minimize the amount of information that needs to be transferred between threads. To that end, we verify one of the threads, considered “main”, as a sequential program. Its verification process initiates queries to its “environment” (which may contain multiple threads). Those queries are answered by sequential verification, if the environment consists of a single thread, or, otherwise, by applying the same hierarchical algorithm on the environment. Our technique is fully automatic, and allows us to use any off-the-shelf sequential model checker. We implemented our technique in a tool called CoMuS and evaluated it against established tools for concurrent verification. Our experiments show that it works particularly well on hierarchically structured programs.

1 Introduction

Verification of concurrent programs is known to be extremely hard. On top of the challenges inherent in verifying sequential programs, it adds the need to consider a high (typically unbounded) number of thread interleavings. An appealing direction is to exploit the modular structure of such programs in verification. Usually, however, a property of the whole system cannot be partitioned into a set of properties that are local to the individual threads. Thus, some knowledge about the interaction of a thread with its environment is required.

In this work we develop a new approach, which utilizes the plethora of work on verification of sequential programs for the purpose of *modularly* verifying the safety of concurrent programs. Our technique automatically reduces the verification of a concurrent program to a series of verification tasks of sequential programs. This allows us to benefit from any past, as well as future, progress in sequential verification.

Our approach is modular in the sense that each sequential verification task roughly corresponds to the verification of a single thread, with some additional information

about the environment in which it operates. This information is *automatically* and *lazily* discovered during the run of the algorithm, when needed.

While our approach is general, it specializes on concurrent programs where the threads are structured hierarchically as it takes a *hierarchical view* of the program. Namely, for the purpose of verification, one of the threads, t_M , is considered “main”, and all other threads are considered its “environment”. The idea is to exploit the hierarchy in order to minimize the amount of information that needs to be transferred between the verification tasks of different threads.

We first analyze t_M using sequential verification, where, for soundness, all interferences from the environment are abstracted (over-approximated) by a function `env_move`, which is called by t_M whenever a context switch should be considered. Initially, `env_move` havoc all shared variables; it is gradually refined during the run of the algorithm. When the sequential model checker discovers a violation of safety in t_M , it also returns a path leading to the violation. The path may include calls to `env_move`, in which case the violation may be spurious (due to the over-approximation). Therefore, the algorithm initiates *queries* to the environment of t_M whose goal is to check whether *certain* interferences, as observed on the violating path, are feasible. Whenever an interference turns out to be infeasible, the `env_move` function is refined to exclude it. Eventually, `env_move` becomes precise enough to enable full verification of the desired property on the augmented t_M . Alternatively, it can reveal a real counterexample in t_M .

The queries are checked on the environment (that may consist of multiple threads) in the same modular manner. Thus we obtain a *hierarchical modular verification*. Along the algorithm, each thread learns about the next threads in the hierarchy, and is provided with assumptions from former threads in the hierarchy to guide its learning. When the program has a hierarchical structure that is aligned with the verification process, this makes the assumptions simpler and speeds up verification.

Our technique is fully automatic and performs *unbounded verification*, i.e., it can both find bugs and prove safety in concurrent programs even with unbounded executions (e.g., due to loops), as long as the number of threads is fixed. It works on the level of program code and generates standard sequential programs in its intermediate checks. This allows us to use any off-the-shelf sequential model-checker. In particular, we can handle concurrent programs with an infinite state-space, provided that the sequential model checker supports such programs (as is the case in our implementation).

We implemented our technique in a prototype called **Concurrent to Multiple Sequential** (CoMuS) and evaluated it against established tools for unbounded verification of concurrent C programs. We use SeaHorn [18] to model check sequential programs.

Our experiments show that the approach works particularly well on programs in which the threads are arranged as a chain, t_1, t_2, \dots, t_k , where thread t_i depends only on its immediate successor t_{i+1} in the chain. This induces a natural hierarchical structure in which t_1 is the main thread with environment t_2, \dots, t_k ; thread t_2 is the main thread in the environment, and so on. This structure often occurs in concurrent implementations of dynamic programming algorithms.

To summarize, the main contributions of our work are as follows:

- We present a new modular verification approach that reduces the verification of a concurrent program to a series of verification tasks of sequential programs. Any off-the-shelf model checker for sequential programs can thus be used.
- Our approach takes a hierarchical view of the program, where each thread learns about the next threads in the hierarchy, and is provided with assumptions from former threads to guide its learning.
- The needed information on a thread’s environment is gathered in the code, automatically and lazily, during the run of the algorithm.
- We implemented our approach and showed that as the number of threads grows, it outperforms existing tools on programs that have a hierarchical structure, such as concurrent implementations of dynamic programming algorithms.

1.1 Related Work

The idea of code transformation to a sequential program appeared in [32,33,24]. However, these works translate the concurrent program to a single nondeterministic sequential program. In contrast, our technique exploits the modular structure of the program.

In the rest of this section, we address unbounded modular techniques for proving safety properties of concurrent programs. Other techniques use bounded model checking, where the bound can address different parameters, such as the number of context switches [19,33], write operations [32] or loop iterations [29,1,35].

The work most closely related to ours is [16,17]. Their technique uses predicate abstraction of both states and environment transitions (similar to our `env_move`), as part of an automatic modular verification framework. The technique also iteratively refines this abstraction by checking possible witnesses of errors. However, they treat all threads symmetrically, whereas our approach exploits a hierarchical view of the program. In addition, [16,17] explore abstract single threads using reachability trees, which are inherent to their technique. We, on the other hand, represent threads (augmented with some environment information) as stand-alone C programs. Thus, we can use any off-the-shelf model checker to address the “sequential part” of the verification problem.

The works in [9,20,12] suggest to apply rely-guarantee reasoning for concurrent (or asynchronous) programs, while the different sections of the program can be verified sequentially. However, their technique requires human effort to specify the rely-guarantee conditions, whereas our approach is completely automatic.

[10] suggests a modular algorithm with rely-guarantee reasoning and automatic condition inference. [21] formalizes the algorithm in the framework of abstract interpretation. However, their algorithm requires finite state systems, and its inferred conditions only refer to changes in global variables. Hence, they fail to prove properties where local variables are necessary for the proof. In our approach, reasoning about local variables is allowed, when we learn that they are necessary for verification. Such variables are then turned into global variables, but their behavior is abstracted, preserving modularity. [6] also tackles the incompetence of modular proofs by exposing local variables as global, according to counterexamples. However, their approach uses BDDs and suits finite state systems. Similar to [16], they treat threads symmetrically. Our approach is applicable to infinite state systems and uses a guided search to derive cross-thread information.

Our queries resemble queries in learning-based compositional verification [5,25], which are also answered by a model checker. Our hierarchical recursive approach resembles the n-way decomposition handled in [25]. However, these works represent programs, assumptions and specification as LTSs, and although extended to deal with shared memory in [31] these algorithms are suitable for finite state systems.

Several works such as [11,14,34,28], tackle the interleaving explosion problem by performing a thread interleaving reduction. [34] combines partial order reduction [13] with the impact algorithm [22], whereas [28] identifies reducible blocks for compositional verification. These approaches are complementary to ours, as our first step is performing an interleaving reduction (to identify cut-points for `env_move` calls).

2 Preliminaries

Sequential Programs. A *sequential program* P is defined by a control flow graph whose nodes are a set of program locations L (also called labels), and whose edges E are a subset of $L \times L$. The program has an initial label, denoted $l^{init} \in L$. Each node l is associated with a command $c \in \text{cmds}$, denoted $\text{cmd}(l)$, which can be an assignment or an if command, as well as `havoc`, `assume` and `assert` (explained below). Intuitively, we think of standard C programs (that may contain loops as well), which can be trivially compiled to such control flow graphs. The program may also include non-recursive functions, which will be handled by inlining.

The program is defined over a set of variables V . Conditions in the program are quantifier-free first-order-logic formulas over V . A special variable `pc` $\notin V$, ranging over L , indicates the program location. A *state* s of P is a pair (l, σ) where $l \in L$ is the value of `pc` and σ is a valuation of V . Variables may have unbounded domains, resulting in a potentially infinite state-space. We also assume the existence of a special error state, denoted $\epsilon = (l_\epsilon, \perp)$. We denote by $l(s)$ and $\sigma(s)$ the first and second components (resp.) of a state $s = (l, \sigma)$. Given an initialization formula ϕ_{init} over V , the set of *initial states* consists of all states (l^{init}, σ) where $\sigma \models \phi_{init}$.

For $s = (l, \sigma)$, let $\text{cmd}(s) = \text{cmd}(l)$. We denote $\text{next}(s) = \{s' \mid s' \text{ can be obtained from } s \text{ using } \text{cmd}(s)\}$. This set is defined according to the command. In particular, $s' \in \text{next}(s)$ implies that $(l(s), l(s')) \in E$. The definition of $\text{next}(s)$ for assignments and if commands is standard. A `v=havoc()` command assigns a non-deterministic value to the variable v . An `assume(b)` command is used to disregard any computation in which the condition b does not hold. Formally, if $s = (l, \sigma)$ and $\text{cmd}(s) = \text{assume}(b)$, then $\sigma \models b \Rightarrow \text{next}(s) = \{(l', \sigma)\}$ where $l' \neq l_\epsilon$ is the unique label such that $(l, l') \in E$, and $\sigma \not\models b \Rightarrow \text{next}(s) = \emptyset$. An `assert(b)` command is defined similarly, except that it moves to the error state if b is violated.

A *computation* ρ of P is a sequence $\rho = s_0 \rightarrow s_1 \rightarrow \dots \rightarrow s_n$ for some $n \geq 0$ s.t. for every two adjacent states s_i, s_{i+1} : $s_{i+1} \in \text{next}(s_i)$. ρ is an *initial computation* in P if it starts from an initial state. ρ is a *reachable computation* in P if there exists an initial computation ρ' for which ρ is the suffix. The *path* of a computation $(l_0, \sigma_0) \rightarrow \dots \rightarrow (l_n, \sigma_n)$ is the sequence of program locations l_0, \dots, l_n .

Preconditions and Postconditions. Given a condition q over V and an edge $e = (l, l')$, a *precondition* of q w.r.t. e , denoted $\text{pre}(e, q)$, is any condition p such that for every

state s , if $\sigma(s) \models p$ and $l(s) = l$ then there exists $s' \in \text{next}(s)$ s.t. $\sigma(s') \models q$ and $l(s') = l'$ ³. A precondition extends to a path $\pi = l_0, \dots, l_n$ in the natural way. The *weakest precondition* of q w.r.t. e (resp., π) is implied by any other precondition, and can be computed in the standard way [8]. We denote it $wp(e, q)$ (resp., $wp(\pi, q)$).

A postcondition of p w.r.t. $e = (l, l')$, denoted $\text{post}(e, p)$, is any condition q such that if $\sigma(s) \models p$, $l(s) = l$ then for every $s' \in \text{next}(s)$, if $l(s') = l'$ then $\sigma(s') \models q$. Postconditions can also be extended to paths $\pi = l_0, \dots, l_n$. We use $\text{post}(\pi, p)$ to denote a postcondition of condition p w.r.t. path π .

Concurrent Programs A concurrent program P consists of multiple threads t_1, \dots, t_m , where each thread t_i has the same syntax as a sequential program over a set of variables V_i and a program location variable pc_i . The threads communicate through shared variables, meaning that generally V_i, V_j are not disjoint for $i \neq j$. A variable is *written* by t_i if it appears on the left hand side of any assignment in t_i . A variable v is *shared* between two threads t_i, t_j if $v \in V_i \cap V_j$. A variable $v \in V_i$ is a *local* variable of t_i if $v \notin V_j$ for every $j \neq i$. Let $V = \bigcup_{i=1}^m V_i$. A *state* of P is a pair (\bar{l}, σ) , where σ is a valuation of V and $\bar{l} = (l_1, \dots, l_m)$ where l_i is the value of pc_i . We also assume one common error state ϵ . Given an initialization formula ϕ_{init} over V , the set of *initial states* consists of all states $(\bar{l}^{\text{init}}, \sigma)$ where $\sigma \models \phi_{\text{init}}$ and l_i^{init} is the initial label of t_i .

The execution of a concurrent program is interleaving, meaning that exactly one thread performs a command at each step, and the next thread to execute is chosen non-deterministically. We consider a sequentially consistent semantics in which the effect of a single command on the memory is immediate. For $s = (\bar{l}, \sigma)$, let $\text{cmd}(s, t_i)$ denote the command of thread t_i at label l_i . We denote $\text{next}(s, t_i) = \{s' \mid s' \text{ can be obtained from } s \text{ after } t_i \text{ performs } \text{cmd}(s, t_i)\}$. A *computation* ρ of the concurrent program P is a sequence $s_0 \xrightarrow{t^1} s_1 \xrightarrow{t^2} \dots \xrightarrow{t^n} s_n$ s.t. for every two adjacent states s_i, s_{i+1} : $s_{i+1} \in \text{next}(s_i, t^{i+1})$. We say that ρ is a computation of thread t in P if $t^j = t$ for every $1 \leq j \leq n$. We define *initial* and *reachable* computations as in the sequential case, but w.r.t. computations of the concurrent program.

We support synchronization operations by modeling them with atomic control commands. For example, `Lock(lock)` is modeled by atomic execution of `assume(lock = false); lock = true`. Since our technique models context switches by explicit calls to `env_move`, we are able to prevent context switches between these commands.

Safety. A computation of a (sequential or concurrent) program is *violating* if it ends in the error state ϵ . The computation is *safe* otherwise. A (sequential or concurrent) program is *safe* if it has no initial violating computations. In the case of a sequential program, we refer to the path of a violating computation as a *violating path*.

A *Sequential Model Checker* is a tool which receives a sequential program as input, and checks whether the program is safe. If it is, it returns “SAFE”. Otherwise, it returns a *counterexample* in the form of a violating path.

³ Note that our definition of a precondition does not require all the successors to satisfy q .

3 Our Methodology

In this section we describe our methodology for verifying safety properties of concurrent programs, given via assertions. The main idea is to use a sequential model checker in order to verify the concurrent program. Our approach handles any (fixed) number of threads. However, for simplicity, we describe our approach for a concurrent program with two threads. The extension to any number of threads can be found in [30].

In the sequel, we fix a concurrent program P with two threads. We refer to one as *the main thread* (t_M) and to the other as *the environment thread* (t_E), with variables V_M and V_E and program location variables p_{c_M} and p_{c_E} , respectively. V_M and V_E might intersect. Let $V = V_M \cup V_E$. Given a state $s = (\bar{l}, \sigma)$, we denote by $l_M(s)$ and $l_E(s)$ the values of p_{c_M} and p_{c_E} , respectively. For simplicity, we assume that safety of P is specified by assertions in t_M (this is not a real restriction of our method).

Our algorithm generates and maintains a sequential program for each thread. Let P_M and P_E be the two sequential programs, with variables $\widehat{V}_M \supseteq V_M$ and $\widehat{V}_E \supseteq V_E$. Each sequential program might include variables of the other thread as well, together with additional auxiliary variables not in V . Our approach is asymmetric, meaning that P_M and P_E have different roles in the algorithm. P_M is based on the code of t_M , and uses a designated function, `env_move`, to abstract computations of t_E . P_E is based on the code of t_E , and is constructed in order to answer specific queries for information required by P_M , specified via assumptions and assertions. The algorithm iteratively applies model checking to each of these programs separately. In each iteration, the code of P_M is gradually modified, as the algorithm learns new information about the environment, and the code of P_E is adapted to answer the query of interest.

In Sec. 4, we first describe the way our algorithm operates on P_M . During the analysis of P_M , information about the environment is retrieved using *environment queries*: Intuitively, an environment query receives two conditions, α and β , and checks whether there exists a reachable computation of t_E in P from α to β . The idea is to perform specific guided queries in t_E , to search for computations that might “help” t_M to reach a violation. If such a computation exists, the environment query returns a formula ψ , which ensures that all states satisfying it can reach β using t_E only. We also require that α and ψ overlap. In order to ensure the reachability of β , the formula ψ might need to address local variables of t_E , as well as p_{c_E} . These variables will then be added to P_M , and may be used for the input of future environment queries. If no such computation of the environment exists, the environment query returns $\psi = \text{FALSE}$. Sec. 5 describes how our algorithm answers environment queries. The formal definition follows.

Definition 1 (Environment Query) *An environment query $\text{Reach}_E(\alpha, \beta)$ receives conditions α and β over $V \cup \{p_{c_E}\}$, and returns a formula ψ over $V \cup \{p_{c_E}\}$ such that:*

1. *If there exists a computation of t_E in P that is (1) reachable in P , (2) starts from a state s s.t. $s \models \alpha$ and (3) ends in a state s' s.t. $s' \models \beta$, then $\psi \wedge \alpha \not\equiv \text{FALSE}$.*
2. *If $\psi \not\equiv \text{FALSE}$ then $\alpha \wedge \psi \not\equiv \text{FALSE}$ and for every state s s.t. $s \models \psi$, there exists a computation (not necessarily reachable) of t_E in P from s to some s' s.t. $s' \models \beta$.*

Multiple threads. The key ingredients used by our technique are (i) an `env_move` function that is used in P_M to overapproximate finite computations (of any length) of t_E

Algorithm 1 Algorithm MainThreadCheck

```

1: procedure MAINTHREADCHECK( $t_M, t_E, \phi_{init}$ )
2:    $P_M = \text{add\_env\_move}$  calls in  $t_M$  and initialize  $\text{env\_move}()$ 
3:   while a violating path exists in  $P_M$  do // using sequential MC
4:     Let  $\pi = l_0, \dots, l_{n+1}$  be a path violating  $\text{assert}(b)$ .
5:     if there are no  $\text{env\_moves}$  in  $\pi$  then return “Real Violation”
6:     let  $l_k$  be the label of the last  $\text{env\_move}$  call in  $\pi$ 
7:     let  $\pi_{start} = l_0, \dots, l_k$  and  $\pi_{end} = l_{k+1}, \dots, l_n$ 
8:      $\beta = \text{wp}(\pi_{end}, \neg b)$  // see (1) in Sec. 4.3
9:      $\alpha = \text{post}(\pi_{start}, \phi_{init})$  // see (2) in Sec. 4.3
10:    Let  $\psi = \text{Reach}_E(\alpha, \beta)$  // environment query for  $t_E$  (see Sec. 5)
11:    if  $\psi$  is FALSE then
12:      Let  $(\alpha', \beta') = \text{Gen}_E(\alpha, \beta)$  // see (4) in Sec. 4.3.
13:       $P_M = \text{RefineEnvMove}(P_M, \alpha', \beta')$  // see (4) in Sec. 4.3
14:    else // see (5) in Sec. 4.3
15:      Add  $\text{assert}(\neg\psi)$  in  $P_M$  at new label  $l'$  right before  $l_k$ 
16:    return “Program is Safe”.

```

(see Sec. 4), and (ii) a `try_start` function that is used in P_E to overapproximate initial computations of P in order to let P_E simulate non-initial computations of t_E that follow them (see Sec. 5). When P has more than two threads, the environment of t_M consists of multiple threads, hence environment queries are evaluated by a recursive application of the same approach. Since the computations we consider in the environment are not necessarily initial, the main thread of the environment should now include both the `env_move` function and the `try_start` function. For more details see [30].

4 Analyzing the Main thread

In this section we describe our algorithm for analyzing the main thread of P for the purpose of proving P safe or unsafe (Algorithm 1). Algorithm 1 maintains a sequential program, P_M , over $\widehat{V}_M \supseteq V_M$, which represents the composition of t_M with an abstraction of t_E . The algorithm changes the code of P_M iteratively, by adding new assumptions and assertions, as it learns new information about the environment.

The abstraction of t_E is achieved by introducing a new function, `env_move`. Context switches from t_M to t_E are modeled explicitly by calls to `env_move`. The body of `env_move` changes during the run of Algorithm 1. However, it always has the property that it over-approximates the set of finite (possibly of length zero) computations of t_E in P that are reachable in P . This is formalized as follows:

Definition 2 (Overapproximation) For a state s_m of P_M (over \widehat{V}_M) s.t. $l(s_m)$ is the beginning or the end of `env_move`, we say that s_m matches a state s of P (over V) if (1) s_m and s agree on $\widehat{V}_M \cap V$, i.e. $\sigma(s_m)|_V = \sigma(s)|_{\widehat{V}_M}$, where $\sigma|_U$ is the projection of σ to the variables appearing in U , and (2) if $p_{c_E} \in \widehat{V}_M$, then $\sigma(s_m)(p_{c_E}) = l_E(s)$.

We say that `env_move` overapproximates the computations of t_E in P if for every reachable computation $\rho = s \xrightarrow{t_E} \dots \xrightarrow{t_E} s'$ of t_E in P (possibly of length 0), and for

```

1  bool claim0 = false, claim1 = false;
2  bool cs1 = false, cs0 = false;
3  int turn;
4  void t0() {
5      while (true) {
6          claim0 = true;
7          turn = 1;
8          while (claim1 && turn != 0) { }
9          cs0 = true;
10         // CRITICAL_SECTION
11         assert(!cs1);
12         cs0 = false;
13         claim0 = false;    }}
14 void t1() {
15     while (true) {
16         claim1 = true;
17         turn = 0;
18         while (claim0 && turn != 1) { }
19         cs1 = true;
20         // CRITICAL_SECTION
21         cs1 = false;
22         claim1 = false;    }}

```

Fig. 1. Peterson’s mutual exclusion algorithm for two threads t_0 and t_1 .

every state s_m s.t. $l(s_m)$ is the beginning of env_move and s_m matches s , there exists a computation $\rho_m = s_m \rightarrow \dots \rightarrow s'_m$ of P_M s.t. (1) ρ_m is a complete execution of env_move , i.e., $l(s'_m)$ is the end of env_move and for every other state s''_m in ρ_m , $l(s''_m)$ is a label within env_move , and (2) s'_m matches s' .

The code of P_M always consists of the original code of t_M , the body of the env_move function (which contains assumptions about the environment), calls to env_move that are added at initialization, and new assertions that are added during the algorithm. \widehat{V}_M always consists of V_M , possibly p_{CE} , some variables of V_E (that are gradually added by need), and some additional auxiliary variables needed for the algorithm (see (4) in Sec. 4.3).

4.1 Initialization

Algorithm 1 starts by constructing the initial version of P_M , based on the code of t_M . To do so, it adds explicit calls to env_move at every location where a context switch needs to be considered in t_M . The latter set of locations is determined by an *interleaving reduction analysis*, which identifies a set of locations, called *cut-points*, such that the original program is safe if and only if all the computations in which context-switches occur only at cut-points are safe.

In addition, the algorithm constructs the initial env_move function which havocs every shared variable of t_E and t_M that is written by t_E . This function will gradually be refined to represent the environment in a more precise way.

Example 3 We use Peterson’s algorithm [26] for mutual exclusion, presented in Fig. 1, as a running example. The algorithm contains a busy-wait loop in both threads, where a thread leaves that loop and enters its critical section only after the turn variable indicates that it is its turn to enter, or the other thread gave up on its claim to enter the critical section. In order to specify the safety property (mutual exclusion), we use additional variables cs_0, cs_1 which indicate that t_0 and t_1 (resp.) are in their critical sections. The safety property is that $\neg \text{cs}_0 \vee \neg \text{cs}_1$ always holds. It is specified by the `assert(!cs1)` command in t_0 between lines 9 and 12, where cs_0 is true

Assume that t_0 was chosen as the main thread and t_1 as the environment thread. We generate a sequential program P_0 , based on the code of t_0 : we add env_moves at

every cut point, as determined by our interleaving reduction mechanism. The initial `env_move` only havocs all variables of P_0 that are written by t_1 , i.e., `claim1`, `turn`, `cs1` (see Fig. 4).

4.2 Iteration of the MainThreadCheck Algorithm

Each iteration of Algorithm 1 starts by applying a sequential model checker to check whether there exists a violating path (that may involve calls to `env_move`) in P_M (line 3). If not, we conclude that the *concurrent* program is safe (line 16), as the `env_move` function over-approximates the computations of the environment. If an assertion violation is detected in P_M , the model checker returns a counterexample in the form of a violating path. If there are no `env_move` calls in the path (line 5), it means that the path represents a genuine violation obtained by a computation of the original main thread, and hence the program is unsafe.

Otherwise, the violation relies on environment moves, and as such it might be spurious. We therefore analyze this counterexample as described in Sec. 4.3. The purpose of the analysis is to check whether t_E indeed enables the environment transitions used along the path. If so, we find “promises of error” for the violated assertion at earlier stages along the path and add them as new assertions in P_M . Intuitively speaking, a “promise of error” is a property ensuring that t_E can make a sequence of steps that will allow t_M to violate its assertion. Such a property may depend on both threads, and hence it is defined over $V \cup \{\text{pce}\}$ (pce_M is given implicitly by the location of the assertion in P_M). Formally, we have the following definition:

Definition 4 Let ψ, ψ' be formulas over $V \cup \{\text{pce}\}$ and let l, l' be labels of t_M . We say that (l, ψ) is a promise of (l', ψ') if for every state s of P s.t. $l_M(s) = l$ and $s \models \psi$ there exists a computation in P starting from s to a state s' s.t. $l_M(s') = l'$ and $s' \models \psi'$.

If (l, ψ) is a promise of $(l', \neg b)$ and l' has an `assert (b)` command, then we say that (l, ψ) is a promise of error.

Note that the definition is transitive. Specifically, if (l, ψ) is a promise of (l', ψ') and (l', ψ') is a promise of error, then (l, ψ) is also a promise of error.

Outcome. Each iteration of Algorithm 1 ends with one of these three scenarios:

1. The algorithm terminates having found a genuine counterexample for P (line 5).
2. The obtained counterexample is found to be spurious since an execution of `env_move` along the path is proved to be infeasible. The counterexample is eliminated by refining the `env_move` function (line 13, also see item (4) in the next section).
3. Spuriousness of the counterexample remains undetermined, but a new promise of error is generated before the last `env_move` call in the violating path. We augment P_M with a new assertion, representing this promise of error (line 15).

4.3 Analyzing a potentially spurious violating path

Let $\pi = l_0, \dots, l_{n+1}$ be a violating path of P_M , returned by the sequential model checker in an iteration of Algorithm 1, which is potentially spurious in P , i.e., contains




Fig. 2. (a) If $Reach_E(\alpha, \beta) = FALSE$, we search for more general α' and β' which restrict the environment transition; (b) If $Reach_E(\alpha, \beta) = \psi \neq FALSE$, then we know that ψ leads to β and that $\psi \wedge \alpha \neq FALSE$.

```

1 void P0() {
2   assert(!(cs1 || claim1);
3   env_move();
4   while (true) {
5     claim0 = true;
6     assert(!(cs1 || claim1);
7     env_move();
8     turn = 1;
9     assert(!(cs1 ||
10      (claim1 && turn != 0));
11   }

```

Fig. 3. The sequential program P_0 after a few iterations of Algorithm 1.

```

1 void env_move() {
2   bool claim1_copy = claim1;
3   int turn_copy = turn;
4   bool cs1_copy = cs1;
5   claim1 = havoc_bool();
6   turn = havoc_int();
7   cs1 = havoc_bool();
8   if (true) { assume(!cs1 || claim1); }

```

Fig. 4. The `env_move` function of P_0 : initially (without highlighted lines); and after one refinement (with highlighted lines).

at least one `env_move` call. Since π is violating, $l_{n+1} = l_e$ and $cmd(l_n) = \text{assert}(b)$ for some condition b . Let l_k , for some $0 \leq k \leq (n-1)$, be the location of the last `env_move` in π . We perform the following steps, illustrated by Fig. 2:

(1) *Computing condition after the environment step:* We compute (backwards) the weakest precondition of $\neg b$ w.r.t. the path $\pi_{end} = l_{k+1}, \dots, l_n$ to obtain $\beta = wp(\pi_{end}, \neg b)$ (line 8). Recall that $\neg b$ is necessarily reachable from β along π_{end} in P_M .

(2) *Computing condition before the environment step:* We compute (forward) a post-condition $\alpha = post(\pi_{start}, \phi_{init})$ starting from ϕ_{init} for the path $\pi_{start} = l_0, \dots, l_k$ (line 9). To ensure progress, we make sure that if π_{start} ends with a suffix of `asserts` then $\alpha \Rightarrow c$ for every `assert(c)` command that appears in this suffix (e.g., by conjoining α with c). Recall that α necessarily holds after executing π_{start} in P_M from ϕ_{init} .

(3) *Environment query:* We compute $\psi = Reach_E(\alpha, \beta)$ (line 10).

Example 5 Fig. 3 presents a prefix of P_M after a few iterations of the algorithm, before the first refinement of `env_move` (i.e., P_M still uses the initial `env_move` function). The previous iterations found new promises of error, and augmented P_M with new assertions. Consider the initial conditions from Fig. 1, i.e., $\phi_{init} \triangleq [claim0 =$

$claim1 = cs1 = cs0 = false]$. Assume that our sequential model checker found the violation given by the next path: 2, 3, 4, 5, 6, 7, 8, 9, .

To check whether the last `env_move` call in line 7 represents a real computation of t_1 , we compute the weakest precondition of the condition $\neg b \triangleq cs1 \wedge (\neg claim1 \vee turn = 0)$, taken from the violated assertion in line 9, w.r.t. the path $\pi_{end} = 8, 9$. The result is $\beta = wp(\pi_{end}, \neg b) = (cs1 \wedge \neg claim1)$. The computation of $\alpha = post(\pi_{start}, \phi)$ for the path $\pi_{start} = 2, 3, 4, 5, 6$ yields $\alpha = (\neg cs0 \wedge claim0 \wedge (\neg cs1 \vee claim1))$. We then generate an environment query $Reach_E(\alpha, \beta)$.

(4) *Refining the `env_move` function:* If $\psi = FALSE$ (line 11) it means that there is no reachable computation of t_E in P from a state s s.t. $s \models \alpha$ to a state s' s.t. $s' \models \beta$. We apply a generalization procedure $Gen_E(\alpha, \beta)$ that returns α', β' s.t. $\alpha \Rightarrow \alpha', \beta \Rightarrow \beta'$ and still $Reach_E(\alpha', \beta') = FALSE$ (line 12). To do so, Gen_E iteratively replaces α and/or β with α', β' s.t. $\alpha \Rightarrow \alpha', \beta \Rightarrow \beta'$ and rechecks $Reach_E(\alpha', \beta')$. For example, if α contains a subformula of the form $\delta_1 \wedge \delta_2$ that appears positively, we attempt to replace it by δ_1 or δ_2 to obtain α' .⁴ We then refine `env_move` to eliminate the environment transition from α' to β' (line 13). Fig. 2(a) illustrates this step.

The refinement is done by introducing in `env_move`, after the variables are havocked, the command `(if ($\alpha'(W_old)$) assume($\neg\beta'$))`, where W_old are the values of the variables before they are havocked in `env_move` (these values are copied by `env_move` to allow evaluating α' on the values of the variables before `env_move` is called). The command blocks all computations of `env_move` from α' to β' . Since such computations were proven by the environment query to be infeasible in t_E , we are ensured that `env_move` remains an overapproximation of the computations of t_E .

Example 6 The call to $Reach_E(\alpha, \beta)$ in Example 5 results in $\psi = FALSE$. Hence, we apply generalization. We obtain two formulas $\alpha' = TRUE, \beta' = \beta$ which indeed satisfy $\alpha \Rightarrow \alpha', \beta \Rightarrow \beta'$ and $Reach_E(\alpha', \beta') = FALSE$. This means that when t_E is called with $\alpha' = TRUE$, then no computation of t_E reaches a state satisfying $\beta' = cs1 \wedge \neg claim1$. Fig. 4 presents the `env_move` function before and after the refinement step based on (α', β') takes place. The refinement step adds the highlighted line to the initial `env_move` function. This line has the constraint `if (true) assume(!cs1 || claim1)`, derived from the observation above.

(5) *Adding assertions:* If $\psi \neq FALSE$, then for every state satisfying ψ there is a computation of t_E in P to a state satisfying β . Since $\beta = wp(\pi_{end}, \neg b)$, it is guaranteed that this computation can be extended (in t_M) along the path π_{end} , which does not use any environment moves, to reach a state s' that violates the assertion `assert(b)`. This is illustrated in Fig. 2(b). We therefore conclude that if ψ is satisfied before the `env_move` at label l_k , a genuine violation can be reached, making (\hat{l}_k, ψ) a promise of error, where \hat{l}_k denotes the label in t_M that corresponds to l_k (the label reached after executing the `env_move` called at label l_k). Therefore, we add a new assertion `assert($\neg\psi$)` right before l_k (line 15). In addition, if ψ includes a variable v that is not in \widehat{V}_M (e.g., `pcE`), then v is added to \widehat{V}_M , its declaration (and initialization, if exists) is added to P_M , and `env_move` is extended to `havoc v` as well (if it is written by t_E).

⁴ More information about the generalization appears in the optimizations section in [30].

5 Answering Environment Queries

Recall that an environment query $Reach_E(\alpha, \beta)$ checks whether there exists a reachable computation ρ of t_E in P from a state $s \models \alpha$ to a state $s' \models \beta$. This computation may involve any finite number of steps of t_E , executed without interference of t_M .

If $\alpha \wedge \beta \neq FALSE$, we simply return β , which represents a computation of length zero. Otherwise, we wish to apply a sequential model checker on t_E in order to reveal such computations, or conclude there are none. However, the computation ρ may not be initial, while our sequential model checker can only search for violating paths starting from an initial state. Hence we construct a modified sequential program P_E , based on the code of t_E , which also represents (over-approximates) non-initial, but reachable, computations ρ of t_E in P . For that, we add in P_E calls to a new function, `try_start`, which models the runs of t_M until the start of ρ . The calls to `try_start` are added in all cut-points computed by an interleaving reduction (similar to the one applied to t_M).

The try_start function. The `try_start` function is responsible for non-deterministically setting the start point of ρ , where context switches to t_M are no longer allowed. This is done by setting a new `start` variable to true (provided that its value is not yet true). We refer to the latter call as the *activation* `try_start`. As long as `start` is false (i.e., prior to the activation call), `try_start` havoc the variables written by t_M . When `start` is set to true, we add an `assume` (α) command after the havoc commands as this is the state chosen to start the computation. To handle the case where `pc_E` appears in α , `try_start` receives the original location (in t_E) in which it is called as a parameter, and updates the explicit `pc_E` variable. Whenever `start` is already true, `try_start` immediately exits, ensuring that ρ indeed only uses transitions of t_E .

In P_E , we also add assertions of the form `assert (!start || $\neg\beta$)` after every call to `try_start`. Hence, a violating path, if found, reaches `start` \wedge β , i.e., it captures a computation in which α was satisfied (when `start` was set to true), and reached β .

Returning Result. If a violating path is not found, we return $Reach_E(\alpha, \beta) = FALSE$. If a violating path m_0, \dots, m_{n+1} is found, let m_k be the label of the activation `try_start` for some $0 \leq k \leq (n-1)$. Let π_E be the projection of m_{k+1}, \dots, m_{n-1} to t_E . We compute the weakest precondition of β w.r.t. the path π_E and obtain $\psi = wp(\pi_E, \beta)$. The computed ψ satisfies the desired requirement: For every state s of P s.t. $s \models wp(\pi_E, \beta)$, there exists a computation ρ of t_E starting from s which follows the path π_E and reaches a state s' satisfying β . Note that ρ might not be reachable, as in the prefix we used an abstraction of t_M . That means that $Reach_E(\alpha, \beta)$ is not “exact” and may return $\psi \neq FALSE$ when there is no reachable computations as required. However, it satisfies the requirements of Def. 1, which is sufficient for soundness and progress. The intuition is that checking the reachability of ψ is done by the main thread.

For an example demonstrating how an environment query is answered see [30].

6 Soundness and Progress

Our algorithm for verifying the concurrent program P terminates when either (i) all the assertions in P_M are proven safe (i.e., neither the original error nor all the new promises

of error can be reached in P_M), in which case Algorithm 1 returns “Program is Safe”. (ii) a violation of some assertion in P_M , which indicates either the original error or a promise of error, is reached without any `env_move` calls, in which case Algorithm 1 returns “Real Violation”. The following theorem summarizes its soundness⁵.

Theorem 1. *If Algorithm 1 returns “Safe” then the concurrent program P has no violating computation; If it returns “Real violation” then P has a violating computation.*

The proof of the first claim shows that our algorithm maintains the overapproximation property of `env_move` (see Def. 2), from which the claim follows immediately. In the proof of the second claim, we show that the properties of an environment query (see Def. 1) and of promises of errors (Def. 4) are satisfied.

While termination is not guaranteed for programs over infinite domains, the algorithm is ensured to make *progress* in the following sense. Each iteration either refines `env_move` (step (4) in Sec. 4.3), making it more precise w.r.t. the real environment, or generates new promises of errors at earlier stages along the violating path (step (5) in Sec. 4.3). In the former case, the set of pairs of states (s, s') represented by the start and end states of computations of `env_move` is strictly decreasing – this set overapproximates the set of pairs of states (s, s') for which t_E has a reachable computation from state s to state s' (see Def. 2). In the latter case, the set of states known to lead to a real violation of safety is strictly increasing. In both cases, the other set remains unchanged.

When the domain of all variables is finite, these two sets are bounded, hence the algorithm is guaranteed to terminate.

7 Experimental Results and Conclusion

Setup. We implemented our algorithm in a prototype tool called CoMuS. The implementation is written in Python 3.5, uses `pyparser` [2] for parsing and transforming C programs, uses `SeaHorn` [18] for sequential model checking, and uses `Z3` [7] to check logical implications for some optimizations. A description of the optimizations can be found in [30]. CoMuS currently supports only a subset of the syntax of C (see Sec. 2). It does not perform alias analysis and hence has limited pointers support. It also does not support dynamic thread creations, although we support any fixed number of threads.

We compare CoMuS with `Threader` [27], `VVT` [15] and `UL-CSeq` [23], the last two being the top scoring model checkers on the concurrency benchmark among *sound unbounded* tools in `SVCOMP'16` and `SVCOMP'17` (resp.). On the concurrency benchmark, `VVT` was 4th overall in `SVCOMP'16`, and `UL-CSeq` was 8th overall in `SVCOMP'17`⁶. `Threader` performs modular verification, abstracts each thread separately and uses an interference abstraction for each pair of threads. `UL-CSeq` performs a reduction to a single non-deterministic sequential program. We used it in its default mode,

⁵ Full proofs appear in <https://tinyurl.com/comusfull>.

⁶ The same benchmark was used for unbounded sound tools and tools which perform unbounded reductions. Bounded tools are typically ranked higher. Our method is unbounded and is able to provide proofs, hence we find the selected tools more suitable for comparison.

with CPAChecker [3] as a backend. VVT combines bounded model checking for bug finding with an IC3 [4] based method for full verification.

We ran the experiments on a x86-64 Linux machine, running Ubuntu 16.04 (Xenial) using Intel(R) Xeon(R) CPU E5-2680 v3 @ 2.50GHz with 8GB of RAM.

Experiments. We evaluated the tools using three experiments. One compares the four tools on concurrent programs with a clear hierarchy. The second compares syntactically similar programs with and without hierarchical structure to evaluate the effect of the structure on the verification time. The last one looked at general concurrent programs.

Hierarchically structured programs. For the first experiment, we used three concurrent dynamic-programming algorithms: Sum-Matrix, Pascal-Triangle and Longest-Increasing-Subsequence. The Sum-Matrix programs receive a matrix A as input. For every pair of indexes (i, j) , it computes the sum of all elements $A[k, l]$, where $k \geq i$ and $l \geq j$. In their concurrent version, each thread is responsible for the computation of a single row. The Pascal-Triangle programs compute all the binomial coefficients up to a given bound. Each thread computes one row of the triangle, where each element in the row depends on two results of the previous row. The Longest-Increasing-Subsequence programs receive an array, and compute for each index i , the length of the longest increasing subsequence that ends at index i . Each thread is responsible for computing the result for a given index of the array, depending on the result of all prefixes. Both these and the matrix programs are infinite state, as the elements of the array (resp. the matrix) are unbounded inputs.

These algorithms have a natural definition for any finite number of threads. Typically, the verification becomes harder as the number of threads increases. For evaluation, we used programs with an increasing number of threads, and check the influence of the number on the different tools. For each instance, we use both a safe and an unsafe version. Both versions differ from each other either only by a change of specification, or by a slight modification that introduces a bug.

The chosen programs have two meaningful characteristics: (i) They exhibit non-trivial concurrency. This means that each thread performs a series of computations, and it can advance when the data for each computation is ready, without waiting for the threads it depends on to complete. Consider the Sum-Matrix problem as an example. Assume thread t_i needs to compute the result at some location (i, j) , and that each row is computed backwards (from the last cell to the first). The computation exploits the results of thread t_{i+1} . Thread t_i needs to wait for thread t_{i+1} to compute the result for location $(i + 1, j)$. However, t_i does not wait for t_{i+1} to terminate, as it can compute the cell (i, j) , while t_{i+1} continues to compute $(i + 1, j - 1)$. (ii) Their data flow graph has a clear chain structure. That is, the threads can be ordered in a chain hierarchy, and each thread only requires information computed by its immediate successor.

Fig. 5 summarizes the results for these programs. The timeout was set to 3600 seconds. The code of the programs is available at tinyurl.com/comusatva18. We include in the table also our running example, the Peterson algorithm.

The results demonstrate a clear advantage for CoMuS for verification (i.e., for safe programs) as the number of threads increases. This can be attributed to the chain structure that lets CoMuS minimize the amount of information transferred between threads.

class	threads	Safe					Unsafe				
		Threader	VVT	ULCSeq	Time	Seahorn Calls	Threader	VVT	ULCSeq	Time	Seahorn Calls
mat	2	0.18	23.30	13.61	3.53	15	0.20	0.15	8.30	1.53	5
mat	3	0.41	455.13	1194.49	36.85	83	0.42	0.28	Unknown	9.01	22
mat	4	4.82	T/O	T/O	68.46	99	4.58	0.52	Unknown	11.75	25
mat	5	243.81	T/O	T/O	270.69	196	201.63	0.74	T/O	116.59	112
mat	6	T/O	T/O	T/O	3,491.10	772	T/O	1.05	T/O	141.39	118
pas	2	0.67	2.44	6.28	0.65	2	0.24	0.14	3.59	1.45	5
pas	3	7.48	95.11	477.46	15.44	43	5.34	0.70	9.70	18.03	37
pas	4	421.46	689.85	T/O	139.54	142	611.25	3.99	1,073.55	19.33	25
pas	5	T/O	T/O	T/O	T/O	-	T/O	20.95	T/O	110.31	110
long	2	0.43	0.17	13.90	0.67	2	0.22	0.14	3.35	0.50	1
long	3	24.94	6.13	201.23	10.64	29	19.85	0.30	52.93	9.41	25
long	4	T/O	624.00	T/O	96.07	106	T/O	2.56	T/O	110.47	109
long	5	T/O	3,299.94	T/O	552.18	222	T/O	123.09	T/O	619.02	238
Pet	2	10.16	13.38	106.18	17.56	45					

Fig. 5. Run times [secs] for all four tools for verifying concurrent dynamic programs algorithms.

For falsification, CoMuS is outperformed by VVT’s bounded method. However, it still performs significantly better than the two other tools when the number of threads grows.

Hierarchical vs. non-hierarchical programs. The programs used for this evaluation are variants of the “fib_bench” examples of the SV-COMP concurrency benchmark. We compare programs in which the data flow graph has a ring topology, vs. programs in which it has a chain topology. For the ring case, consider a program with threads t_0, \dots, t_{n-1} and variables v_0, \dots, v_{n-1} . Each thread t_i runs in a loop, and iteratively performs $v_i += v_{(i+1 \pmod n)}$. The checked property is that v_0 does not surpass an upper bound. The chain case is identical except that for the last thread, t_{n-1} , we break the chain and perform $v_{n-1} += 1$ instead of $v_{n-1} += v_0$. Fig. 6 presents the results of this comparison. All the programs in the table are safe and with two loop iterations. The timeout was set to 1200 seconds.

For the ring case, all tools fail to verify programs with ≥ 4 threads. Threader presents similar results for both ring and chain topologies. VVT benefits from the less dependent chain topology, but still timeouts on > 3 threads. CoMuS, on the other hand, is designed to exploit hierarchy, and benefits significantly from the chain topology, where it verifies all instances. ULCSeq is excluded from the table as it times-out on the “fib_bench” examples (both in our experiments and in the SV-COMP results).

The reason for CoMuS’s different runtime on the chain and ring variants is that for programs that have no clear hierarchy (as in the ring programs), the conditions passed to the environment queries must include information relevant to the caller thread; a manual inspection shows that they typically become more complex. As similar phenomenon happens if the verification order used by CoMuS is not aligned with the hierarchy of the program. For example, switching the verification order of the last two threads in the long_th3_safe example, increases the verification time from 10 to 25 seconds.

General concurrent programs. We also evaluated the tools on a partial subset of the SV-COMP concurrency benchmark, whose code is supported by CoMuS. Typically, on these runs CoMuS was outperformed by the other tools. We conclude that even

threads	Ring				Chain			
	CoMus				CoMus			
	Threader	VVT	Time	Seahorn Calls	Threader	VVT	Time	Seahorn Calls
2	17.70	48.93	8.52	33	11.70	18.47	3.61	16
3	T/O	T/O	126.81	294	T/O	54.51	13.35	42
4	T/O	T/O	T/O	-	T/O	T/O	40.36	88
5	T/O	T/O	T/O	-	T/O	T/O	99.09	151
6	T/O	T/O	T/O	-	T/O	T/O	275.30	247
7	T/O	T/O	T/O	-	T/O	T/O	617.32	368

Fig. 6. Run times [secs] for fib_bench programs with ring topology vs. chain topology.

though our method can be applied to programs without a clear hierarchical structure, it is particularly beneficial for programs in which the hierarchy is inherent.

Conclusion. In this work we develop an automatic, modular and hierarchical method for proving or disproving safety of concurrent programs by exploiting model checking for sequential programs. The method can handle infinite-state programs. It is sound and unbounded. We implemented our approach in a prototype tool called CoMuS, which compares favorably with top scoring model checkers on a particular class of problems, as previously characterized. In the future we intend to exploit internal information gathered by the sequential model checker (e.g., SeaHorn) to further speedup our results. We would also like to examine how to apply our approach to other hierarchies (e.g., trees).

Acknowledgement. This publication is part of a project that has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No [759102-SVIS]). The research was partially supported by Len Blavatnik and the Blavatnik Family foundation, the Blavatnik Interdisciplinary Cyber Research Center, Tel Aviv University, and the United States-Israel Binational Science Foundation (BSF) grants No. 2016260 and 2012259.

References

1. J. Alglave, D. Kroening, and M. Tautschnig. Partial orders for efficient Bounded Model Checking of concurrent software. In *Computer Aided Verification (CAV)*, 2013.
2. E. Bendersky. <https://github.com/eliben/pycparser>.
3. D. Beyer and M. E. Keremoglu. CPAchecker: A tool for configurable software verification. In *Computer Aided Verification (CAV)*, pages 184–190, 2011.
4. A. R. Bradley. SAT-based model checking without unrolling. In *VMCAI*, 2011.
5. J. M. Cobleigh, D. Giannakopoulou, and C. S. Pasareanu. Learning assumptions for compositional verification. In *TACAS*, pages 331–346, 2003.
6. A. Cohen and K. S. Namjoshi. Local proofs for global safety properties. In *CAV*, 2007.
7. L. De Moura and N. Bjørner. Z3: An efficient SMT solver. *TACAS*, 2008.
8. E. W. Dijkstra. *A Discipline of Programming*. Prentice-Hall, 1976.
9. C. Flanagan, S. N. Freund, and S. Qadeer. Thread-modular verification for shared-memory programs. In *ESOP*, 2002.
10. C. Flanagan and S. Qadeer. Thread-modular model checking. In *SPIN*, 2003.

11. C. Flanagan and S. Qadeer. Transactions for software model checking. *Electr. Notes Theor. Comput. Sci.*, 89(3):518–539, 2003.
12. I. Gavran, F. Niksic, A. Kanade, R. Majumdar, and V. Vafeiadis. Rely/guarantee reasoning for asynchronous programs. In *CONCUR*, 2015.
13. P. Godefroid. *Partial-Order Methods for the Verification of Concurrent Systems - An Approach to the State-Explosion Problem*, volume 1032 of *LNCS*. Springer, 1996.
14. G. Gueta, C. Flanagan, E. Yahav, and M. Sagiv. Cartesian partial-order reduction. *SPIN'07*.
15. H. Günther, A. Laarman, and G. Weissenbacher. Vienna verification tool: IC3 for parallel software - (competition contribution). In *TACAS*, 2016.
16. A. Gupta, C. Popeea, and A. Rybalchenko. Predicate abstraction and refinement for verifying multi-threaded programs. In *POPL*, 2011.
17. A. Gupta, C. Popeea, and A. Rybalchenko. Threader: A constraint-based verifier for multi-threaded programs. In *Computer Aided Verification (CAV)*, pages 412–417, 2011.
18. A. Gurfinkel, T. Kahsai, A. Komuravelli, and J. A. Navas. The seahorn verification framework. In *Computer Aided Verification (CAV)*, pages 343–361, 2015.
19. A. Lal and T. W. Reps. Reducing concurrent analysis under a context bound to sequential analysis. In *Computer Aided Verification (CAV)*, pages 37–51, 2008.
20. K. R. M. Leino and P. Müller. A basis for verifying multi-threaded programs. In *ESOP'09*.
21. A. Malkis, A. Podelski, and A. Rybalchenko. Thread-modular verification is cartesian abstract interpretation. In *Theoretical Aspects of Computing - ICTAC*, pages 183–197, 2006.
22. K. L. McMillan. Lazy abstraction with interpolants. In *CAV*, 2006.
23. T. L. Nguyen, B. Fischer, S. La Torre, and G. Parlato. Unbounded lazy-CSeq: A lazy sequentialization tool for C programs with unbounded context switches - (competition contribution). In *TACAS*, pages 461–463, 2015.
24. T. L. Nguyen, B. Fischer, S. La Torre, and G. Parlato. Lazy sequentialization for the safety verification of unbounded concurrent programs. In *ATVA*, 2016.
25. C. S. Pasareanu, D. Giannakopoulou, M. G. Bobaru, J. M. Cobleigh, and H. Barringer. Learning to divide and conquer: applying the L* algorithm to automate assume-guarantee reasoning. *Formal Methods in System Design*, 32(3):175–205, 2008.
26. G. L. Peterson. Myths about the mutual exclusion problem. *Inf. Process. Lett.*, 12(3):115–116, 1981.
27. C. Popeea and A. Rybalchenko. Threader: A verifier for multi-threaded programs - (competition contribution). In *TACAS*, pages 633–636, 2013.
28. C. Popeea, A. Rybalchenko, and A. Wilhelm. Reduction for compositional verification of multi-threaded programs. In *FMCAD*, pages 187–194, 2014.
29. I. Rabinovitz and O. Grumberg. Bounded model checking of concurrent programs. *CAV'05*.
30. D. Rasin. Modular verification of concurrent programs via sequential model checking. M.Sc. thesis, Technion — Israel Institute of Technology, 2018.
31. N. Sinha and E. M. Clarke. SAT-based compositional verification using lazy learning. In *CAV*, pages 39–54, 2007.
32. E. Tomasco, O. Inverso, B. Fischer, S. La Torre, and G. Parlato. Verifying concurrent programs by memory unwinding. In *TACAS*, 2015.
33. E. Tomasco, T. L. Nguyen, O. Inverso, B. Fischer, S. La Torre, and G. Parlato. Lazy sequentialization for TSO and PSO via shared memory abstractions. In *FMCAD*, 2016.
34. B. Wachter, D. Kroening, and J. Ouaknine. Verifying multi-threaded software with impact. In *Formal Methods in Computer-Aided Design, FMCAD*, pages 210–217, 2013.
35. M. Zheng, J. G. Edenhofner, Z. Luo, M. J. Gerrard, M. S. Rogers, M. B. Dwyer, and S. F. Siegel. CIVL: applying a general concurrency verification framework to c/threads programs (competition contribution). In *TACAS*, pages 908–911, 2016.