

Abstracting and Counting Synchronizing Processes

Zeinab Ganjei, Ahmed Rezine*, Petru Eles, and Zebo Peng

Linköping University, Sweden

Abstract. We address the problem of automatically establishing synchronization dependent correctness (e.g. due to using barriers or ensuring absence of deadlocks) of programs generating an arbitrary number of concurrent processes and manipulating variables ranging over an infinite domain. This is beyond the capabilities of current automatic verification techniques. For this purpose, we define an original logic that mixes variables referring to the number of processes satisfying certain properties and variables directly manipulated by the concurrent processes. We then combine existing works on counter, predicate, and constrained monotonic abstraction and build an original nested counter example based refinement scheme for establishing correctness (expressed as non reachability of configurations satisfying formulas in our logic). We have implemented a tool (PACMAN, for predicated constrained monotonic abstraction) and used it to perform parameterized verification for several programs whose correctness crucially depends on precisely capturing the number of processes synchronizing using shared variables.

Key words: parameterized verification, counting logic, barrier synchronization, deadlock freedom, multithreaded programs, counter abstraction, predicate abstraction, constrained monotonic abstraction

1 Introduction

We address the problem of automatic and parameterized verification for concurrent multithreaded programs. We focus on synchronization related correctness as in the usage of barriers or integer shared variables for counting the number of processes at different stages of the computation. Such synchronizations orchestrate the different phases of the executions of possibly arbitrary many processes spawned during runs of multithreaded programs. Correctness is stated in terms of a new counting logic that we introduce. The counting logic makes it possible to express statements about program variables and variables counting the number of processes satisfying some properties on the program variables. Such statements can capture both individual properties, such as assertion violations, and global properties such as deadlocks or relations between the numbers of processes satisfying certain properties.

* In part supported by the 12.04 CENIIT project.

Synchronization among concurrent processes is central to the correctness of many shared memory based concurrent programs. This is particularly true in certain applications such as scientific computing where a number of processes, parameterized by the size of the problem or the number of cores, is spawned in order to perform heavy computations in phases. For this reason, when not implemented individually using shared variables, constructs such as (dynamic) barriers are made available in mainstream libraries and programming languages such as Pthreads, `java.util.concurrent` or OpenMP.

Automatically taking into account the different phases by which arbitrary many processes can pass is beyond the capabilities of current automatic verification techniques. Indeed, and as an example, handling programs with barriers of arbitrary sizes is a non trivial task even in the case where all processes only manipulate boolean variables. To enforce the correct behaviour of a barrier, a verification procedure needs to capture relations between the number of processes satisfying certain properties, for instance that there are no processes that are not waiting for the barrier before any process can cross it. This amounts to testing that the number of processes at certain locations is zero. Checking violations of program assertions is then tantamount to checking state reachability of a counter machine where counters track the number of processes satisfying predicates such as being at some program location. No sound verification techniques can therefore be complete for such systems.

Our approach to get around this problem builds on the following observation. In case there are no tests for the number of processes satisfying certain properties (e.g. being in specific programs locations for barriers), symmetric boolean concurrent programs can be exactly encoded as counter machines without tests, i.e., essentially vector addition systems (VASS). For such systems, state reachability can be decided using a backwards exploration that only manipulates sets that are upward closed with respect to the component wise ordering [2, 7]. The approach is exact because of monotonicity of the induced transition system (more processes can fire more transitions since there are no tests on the numbers of processes). Termination is guaranteed by well quasi ordering of the component wise ordering on the natural numbers. The induced transition system is no more monotonic in the presence of tests on the number of processes. The idea in monotonic abstraction [10] is to modify the semantics of the entailed tests (e.g., tests for zero for barriers), such that processes not satisfying the tests are removed (e.g., tests for zero are replaced by resets). This results in a monotonic over-approximation of the original transition system and spurious traces are possible. This is also true for verification approaches that generate (broadcast) concurrent boolean programs as abstractions of concurrent programs manipulating integer variables. Such boolean approximations are monotonic even when the original program (before abstraction) can encode tests on the number of processes and is therefore not monotonic. Indeed, having more processes while respecting important relations between their numbers and certain variables in the original programs does not necessarily allow to fire more transitions (which is what abstracted programs do in such approaches).

Our approach consists in two nested counter example guided abstraction refinement loops. We summarize our contributions in the following points.

1. We propose an original *counting logic* that allows to express statements about program variables and about the number of processes satisfying certain predicates on the program variables.
2. We implement the outer loop by leveraging on existing symmetric predicate abstraction techniques. We encode resulting boolean programs in terms of a monotonic counter machine where reachability of configurations satisfying a *counting property* from our logic is captured as a state reachability problem.
3. We explain how to strengthen the counter machine using *counting invariants*, i.e. properties from our logic that hold on all runs. these can be automatically generated using classical thread modular analysis techniques.
4. We leverage on existing constrained monotonic abstraction techniques to implement the inner loop and to address the state reachability problem.
5. We have implemented both loops, together with automatic counting invariants generation, in a prototype (PACMAN) that allowed us to automatically establish or refute counting properties such as deadlock freedom and assertions. All programs we report on may generate arbitrary many processes.

Related work . Several works consider parameterized verification for concurrent programs. In [9] the authors use counter abstraction and truncate the counters to obtain a finite state system. Environment abstraction [4] combines predicate and counter abstraction. Both [9, 4] can require considerable interaction and human ingenuity to find the right predicates. The works in [8, 1] explore finite instances and automatically check for cutoff conditions. Except for checking larger instances, it is unclear how to refine entailed abstractions. The closest works to ours are [10, 3, 5]. We introduced (constrained) monotonic abstraction in [10, 3]. Monotonic abstraction was not combined with predicate abstraction, nor did it explicitly target counting properties or dynamic barrier based synchronization. In [5], the authors propose a predicate abstraction framework for concurrent multithreaded programs. As explained earlier such abstractions cannot exclude behaviours forbidden by synchronization mechanisms such as barriers. In our work, we build on [5] in order to handle shared and local integer variables. To the best of our knowledge, our work is the first automatic verification approach that specifically targets parameterized correctness of programs involving constructs where the number of processes are kept track of (e.g, using barriers).

Outline. We start by illustrating our approach using an example in Sec. 2 and introduce some preliminaries in Sec. 3. We then define concurrent programs and describe our counting logic in Sec. 4. Next, we explain the different phases of our nested loop in Sec. 5 and report on our experimental results in Sec. 6. We finally conclude in Sec. 7.

2 A Motivating Example

Consider the concurrent program described in Fig. 1. In this example, a *main* process spawns (transition t_1) an arbitrary number (*count*) of *proc* processes (at location $proc@lc_{ent}$). All processes share four integer variables (namely *max*, *prev*, *wait* and *count*) and a single boolean variable *proceed*. Initially, the variables *wait* and *count* are 0 while *proceed* is false. The other variables may assume any value belonging to their respective domains. Each *proc* process possesses a local integer variable *val* that can only be read or written by its owner. Each *proc* process assigns to *max* the value of its local variable *val* (may be any integer value) in case the later is larger than the former. Transitions t_6 and t_7 essentially implement a barrier in the sense that all *proc* processes must have reached $proc@lc_3$ in order for any of them to move to location $proc@lc_4$. After the barrier, the *max* value should be larger or equal to any previous local *val* value stored in the shared *prev* (i.e., $prev \leq max$ should hold). Violation of this assertion can be captured with the *countingpredicate* (introduced in Sec 4) $(proc@lc_4 \wedge \neg(prev \leq max))^{\#} \geq 1$ stating that the number of processes at location $proc@lc_4$ and witnessing that $prev > max$ is larger or equal than 1.

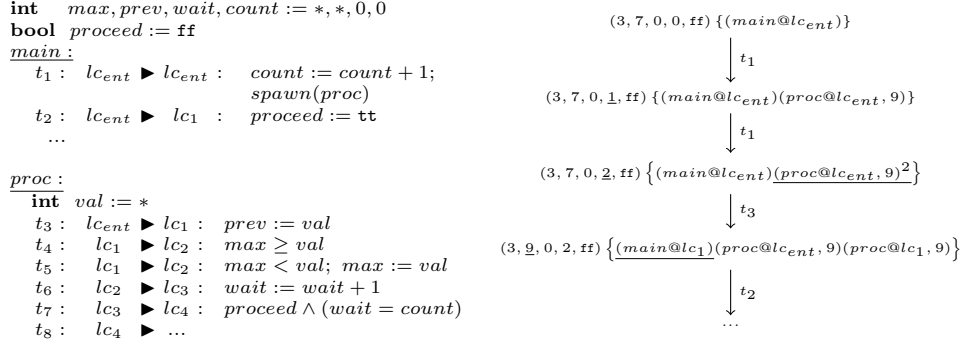


Fig. 1. The max example (left) and a possible run (right). The assertion $(proc@lc_4 \wedge \neg(prev \leq max))^{\#} \geq 1$ cannot be violated when starting with a single *main* process.

A possible run of the concurrent program is depicted to the right of Fig. 1. Observe that we write $(proc@lc_{ent}, 9)^2$ to mean that there are two *proc* processes at location lc_{ent} s.t. their local *val* are both equal to 9. In other words, we use counter abstraction since the processes are symmetric, i.e., processes running the same lines of code with equal local variables are interchangeable and we do not need to differentiate them. The initial configuration of this run is given in terms of the values of *max*, *prev*, *wait*, *count* and *proceed*, here $(3, 7, 0, 0, \mathbf{ff})$, and of the *main* process being at location lc_{ent} . The *main* process starts by incrementing the *count* variable and by spawning a *proc* process twice.

The assertion $(proc@lc_4 \wedge \neg(prev \leq max))^\# \geq 1$ is never violated under any run starting from a single main process. In order to establish this fact, any verification procedure needs to take into account the barrier in t_7 in addition to the two sources of infiniteness; namely, the infinite domain of the shared and local variables and the number of *procs* that may participate in the run. Until now, the closest works to ours deal with these two sources of infiniteness separately and cannot capture facts that relate them, namely, the values of the program variables and the number of generated processes. Any sound analysis that does not take into account that the *count* variable holds the number of spawned *proc* processes and that *wait* represents the number of *proc* processes at locations lc_3 or later will not be able to discard scenarios were a *proc* process executes $prev := val$ although one of them is at $proc@lc_4$. Such an analysis will therefore fail to show that $prev \leq max$ each time a process is at line $proc@lc_4$.

Our original nested CEGAR loop, called **Predicated Constrained Monotonic Abstraction** and depicted in Fig. 2, systematically leverages on simple facts that relate numbers of processes to the variables manipulated in the program. This allows us to verify or refute safety properties (e.g., assertions, deadlock freedom) depending on complex behaviors induced by constructs such as dynamic barriers. We illustrate our approach in the remaining on the max example of Fig. 1.

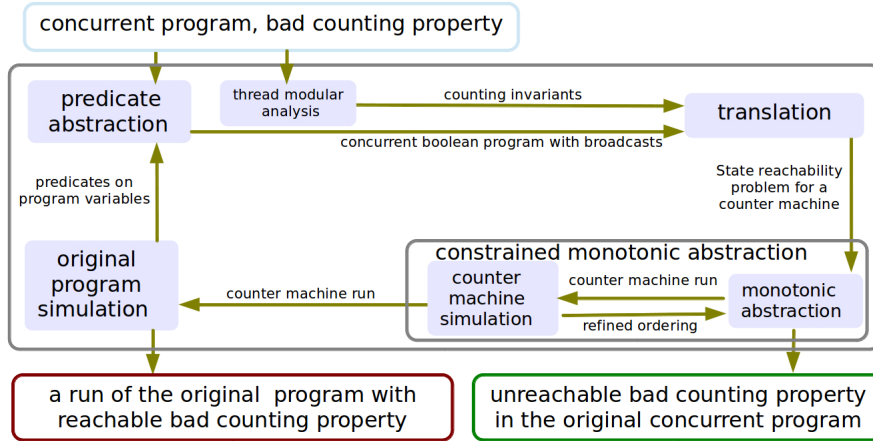


Fig. 2. Predicated Constrained Monotonic Abstraction

From concurrent programs to boolean concurrent programs. We build on the recent predicate abstraction techniques for concurrent programs. Such techniques would discard all variables and predicates and only keep the control flow together with the *spawn* and *join* statements. This leads to a number of counter example guided abstraction refinement steps (the outer CEGAR loop in Fig. 2) that require the addition of new predicates. Our implementation adds the

predicates $proceed$, $prev \leq val$, $prev \leq max$, $wait \leq count$, $count \leq wait$. It is worth noticing that all variables of the obtained concurrent program are finite state (in fact booleans). Hence, one would need a finite number of counters in order to faithfully capture the behavior of the abstracted program using counter abstraction. In addition, some of the transitions of the program, such as t_3 where a shared variable is updated, are abstracted using a broadcast transition.

From concurrent boolean programs to counter machines. Given a concurrent boolean program, we generate a counter machine that essentially boils down to a vector addition system with transfers. Each counter in the machine counts the number of processes at some location and with some local variables combination. One state in the counter machine represents reaching a configuration satisfying the counting property we want to verify. The other states correspond to the possible combinations of the global variables. The transfers represent broadcast transitions. Such a machine cannot relate the number of processes in certain locations (for instance the number of spawned processes $proc$ so far) to the predicates that are valid at certain states (for instance that $count = wait$). In order to remedy to this fact, we make use of *counting invariants* that relate program variables, $count$ and $wait$ in the following invariants, to the number of processes at certain locations.

$$count = (proc@lc_{ent})^\# + \sum_{i \geq 1} (proc@lc_i)^\# \quad wait = \sum_{i \geq 3} (proc@lc_i)^\#$$

We automatically generate such invariants using a simple thread modular analysis that tracks the number of processes at each location. Given such counting invariants, we constrain the counter machine and generate a more precise machine for which state reachability may now be undecidable.

Constrained monotonic abstraction. We monotonically abstract the resulting counter machine in order to answer the state reachability problem. Spurious traces are now possible. For instance, we remove processes violating the constraint imposed by the barrier in Fig.1. This example illustrates a situation where such approximations yield false positives. To see this, suppose two processes are spawned and $proceed$ is set to \mathbf{tt} . A first process gets to lc_3 and waits. The second process moves to lc_1 . Removing the second process opens the barrier for the first process. However, the assertion can now be violated because the removed process did not have time to update the variable max . Constrained monotonic abstraction eliminates spurious traces by refining the preorder used in monotonic abstraction. For the example of Fig.1, if the number of processes at lc_1 is zero, then closing upwards will not alter this fact. By doing so, the process that was removed in forward at lc_1 is not allowed to be there to start with, and the assertion is automatically established for any number of processes. The inner loop of our approach can automatically add more elaborate refinements such as comparing the number of processes at different locations. Exact traces of the counter machine are sent to the next step and unreachability of the control location establishes safety of the concurrent program.

Trace Simulation. Traces obtained in the counter machine are real traces as far as the concurrent boolean program is concerned. Those traces can be simulated on the original program to find new predicates (e.g., using Craig interpolation) and use them in the next iteration of the outer loop.

3 Preliminaries

We write \mathbb{N} to mean the set of natural numbers and \mathbb{Z} to mean the one of integer numbers. We write k to mean a constant in \mathbb{Z} and b to mean a boolean value in $\{\mathbf{tt}, \mathbf{ff}\}$. We use v, s, l, c, a to mean integer variables and $\tilde{v}, \tilde{s}, \tilde{l}$ to mean boolean variables. We let V, S, L, C and A (resp. \tilde{V}, \tilde{S} and \tilde{L}) denote sets of integer variables (resp. sets of boolean variables). We let \sim be an element in $\{<, \leq, =, \geq, >\}$. An expression e (resp. predicate π) belonging to the set $\mathbf{exprsOf}(V)$ (resp. $\mathbf{predsOf}(\tilde{V}, E)$) of arithmetic expressions (resp. boolean predicates) over integer variables V (resp. boolean variables \tilde{V} and arithmetic expressions E) is defined as follows.

$$\begin{aligned} e &::= k \mid v \mid (e + e) \mid (e - e) \mid k e & v \in V \\ \pi &::= b \mid \tilde{v} \mid (e \sim e) \mid \neg \pi \mid \pi \wedge \pi \mid \pi \vee \pi & \tilde{v} \in \tilde{V}, e \in E \end{aligned}$$

We write $\mathit{vars}(e)$ to mean all variables v appearing in e , and $\mathit{vars}(\pi)$ to mean all variables \tilde{v} and v appearing in π or in e in π . We also write $\mathit{atoms}(\pi)$ (the set of atomic predicates) to mean all boolean variables \tilde{v} and all comparisons $(e \sim e)$ appearing in π . We use the letters $\sigma, \eta, \theta, \nu$ (resp. $\tilde{\sigma}, \tilde{\eta}, \tilde{\nu}$) to mean mappings from sets of variables to \mathbb{Z} (resp. \mathbb{B}). Given n mappings $\nu_i : V_i \rightarrow \mathbb{Z}$ such that $V_i \cap V_j = \emptyset$ for each $i, j : 1 \leq i \neq j \leq n$, and an expression $e \in \mathbf{exprsOf}(V)$, we write $\mathit{val}_{\nu_1, \dots, \nu_n}(e)$ to mean the expression obtained by replacing each occurrence of a variable v appearing in some V_i by the corresponding $\nu_i(v)$. In a similar manner, we write $\mathit{val}_{\nu, \tilde{\nu}, \dots}(\pi)$ to mean the predicate obtained by replacing the occurrence of integer and boolean variables as stated by the mappings $\nu, \tilde{\nu}$, etc. Given a mapping $\nu : V \rightarrow \mathbb{Z}$ and a set $\mathit{subst} = \{v_i \leftarrow k_i \mid 1 \leq i \leq n\}$ where variables v_1, \dots, v_n are pairwise different, we write $\nu[\mathit{subst}]$ to mean the mapping ν' such that $\nu'(v_i) = k_i$ for each $1 \leq i \leq n$ and $\nu'(v) = \nu(v)$ otherwise. We abuse notation and write $\nu[\{v_i \leftarrow v'_i \mid 1 \leq i \leq n\}]$, for $\nu : V \rightarrow \mathbb{Z}$ where variables v_1, \dots, v_n are in V and pairwise different, to mean the mapping $\nu' : (V \setminus \{v_i \mid 1 \leq i \leq n\}) \cup \{v'_i \mid 1 \leq i \leq n\} \rightarrow \mathbb{Z}$ and such that $\nu'(v') = \nu(v)$ for each $v \in \{v_i \mid 1 \leq i \leq n\}$ and $\nu'(v) = \nu(v)$ otherwise. We define $\tilde{\nu}[\{v_i \leftarrow b_i \mid 1 \leq i \leq n\}]$ and $\tilde{\nu}[\{v_i \leftarrow \tilde{v}'_i \mid 1 \leq i \leq n\}]$ in a similar manner.

A multiset m over a set X is a mapping $X \rightarrow \mathbb{Z}$. We write $x \in m$ to mean $m(x) \geq 1$. The size $|m|$ of a multiset m is $\sum_{x \in X} m(x)$. We sometimes view a multiset m as a sequence $x_1, x_2, \dots, x_{|m|}$ where each element x appears $m(x)$ times. We write $x \oplus m$ to mean the multiset m' such that $m'(y)$ equals $m(y) + 1$ if $x = y$ and $m(y)$ otherwise.

4 Concurrent Programs and Counting Logic

To simplify the presentation, we assume a concurrent program (or program for short) to consist in a single non-recursive procedure manipulating integer variables. Arguments and return values are passed using shared variables. Programs where arbitrary many processes run a finite number of procedures can be encoded by having the processes choose a procedure at the beginning.

Syntax. A procedure in a program (S, L, T) is given in terms of a set T of transitions $(lc_1 \blacktriangleright lc'_1 : stmt_1), (lc_2 \blacktriangleright lc'_2 : stmt_2), \dots$ operating on two finite sets of integer variables, namely a set S of shared variables (denoted s_1, s_2, \dots) and a set L of local variables (denoted l_1, l_2, \dots). Each transition $(lc \blacktriangleright lc' : stmt)$ involves two locations lc and lc' and a statement $stmt$. We write Loc to mean the set of all locations appearing in T . We always distinguish two locations, namely an entry location lc_{ent} and an exit location lc_{ext} . In any transition, location lc_{ext} may appear as the source location only if the destination is also lc_{ext} (i.e., a sink location). Program syntax is given in terms of pairwise different v_1, \dots, v_n in $S \cup L$, e_1, \dots, e_n in $\text{exprsOf}(S \cup L)$ and π is in $\text{predsOf}(\text{exprsOf}(S \cup L))$.

$$\begin{aligned} \text{prog} ::= & (s := (k \mid *)^* \quad \underline{\text{proc}} : (l := (k \mid *)^* \quad (lc \blacktriangleright lc : stmt)^+ \\ \text{stmt} ::= & \text{spawn} \mid \text{join} \mid \pi \mid v_1, \dots, v_n := e_1, \dots, e_n \mid \text{stmt}; \text{stmt} \end{aligned}$$

Semantics. Initially, a single process starts executing the procedure with both local and shared variables initialized as stated in their definitions. Executions might involve an arbitrary number of spawned processes. The execution of any process (whether initial or spawned with the statement *spawn*) starts at the entry location lc_{ent} . Any process at an exit point lc_{ext} can be eliminated by a process executing a *join* statement. An assume π statement blocks if the predicate π over local and shared variables does not evaluate to true. Each transition is executed atomically without interruption from other processes.

More formally, a *configuration* is given in terms of a pair (σ, m) where the *shared state* $\sigma : S \rightarrow \mathbb{N}$ is a mapping that associates a natural number to each variable in S . An *initial shared state* (written σ_{init}) is a mapping that complies with the initial constraints for the shared variables. The multiset m contains *process configurations*, i.e., pairs (lc, η) where the location lc belongs to Loc and the *process state* $\eta : L \rightarrow \mathbb{N}$ maps each local variable to a natural number. We also write η_{init} to mean an *initial process state*. An *initial multiset* (written m_{init}) maps all (lc, η) to 0 except for a single (lc_{ent}, η_{init}) mapped to 1. We introduce a relation \xrightarrow{stmt}_P in order to define statements semantics (Fig. 3). We write $(\sigma, \eta, m) \xrightarrow{stmt}_P (\sigma', \eta', m')$, where σ, σ' are shared states, η, η' are process states, and m, m' are multisets of process configurations, in order to mean that a process at process state η when the shared state is σ and the other process configurations are represented by m , can execute the statement *stmt* and take the program to a configuration where the process is at state η' , the shared

state is σ' and the configurations of the other processes are captured by m' . For instance, a process can always execute a join if there is another process at location lc_{ext} (rule *join*). A process executing a multiple assignment atomically updates shared and local variables values according to the values taken by the expressions of the assignment before the execution (rule *assign*).

$$\begin{array}{c}
 \frac{(\sigma, \eta, m) \xrightarrow{stmt} (\sigma', \eta', m')}{(\sigma, (lc, \eta) \oplus m) \xrightarrow{(lc \blacktriangleright lc' : stmt)} (\sigma', (lc', \eta') \oplus m')} : trans \qquad \frac{val_{\sigma, \eta}(\pi)}{(\sigma, \eta, m) \xrightarrow{\pi} (\sigma, \eta, m)} : assume \\
 \\
 \frac{(\sigma, \eta, m) \xrightarrow{stmt} (\sigma', \eta', m') \quad (\sigma', \eta', m') \xrightarrow{stmt'} (\sigma'', \eta'', m'')}{(\sigma, \eta, m) \xrightarrow{stmt : stmt'} (\sigma'', \eta'', m'')} : seq \qquad \frac{m = ((lc_{ext}, \eta') \oplus m')}{(\sigma, \eta, m) \xrightarrow{join} (\sigma, \eta, m')} : join \\
 \\
 \frac{subst_A = \{v_i \leftarrow val_{\sigma', \eta'}(e_i) \mid v_i \in A\}}{(\sigma, \eta, m) \xrightarrow{v_1, \dots, v_n := e_1, \dots, e_n} (\sigma[subst_S], \eta[subst_L], m)} : assign \qquad \frac{m' = (lc_{ent}, \eta_{init}) \oplus m}{(\sigma, \eta, m) \xrightarrow{spawn} (\sigma, \eta, m')} : spawn
 \end{array}$$

Fig. 3. Semantics of concurrent programs.

We write $(\sigma, m) \longrightarrow_P (\sigma', m')$ if $(\sigma, m) \xrightarrow{t}_P (\sigma', m')$ for a transition t . A P run ρ is a sequence $(\sigma_0, m_0), t_1, \dots, t_n, (\sigma_n, m_n)$. The run is P feasible if $(\sigma_i, m_i) \xrightarrow{t_{i+1}}_P (\sigma_{i+1}, m_{i+1})$ for each $i : 0 \leq i < n$ and σ_0 and m_0 are initial. We say that a configuration (σ, m) is *reachable* if there is a feasible P run $(\sigma_0, m_0), t_1, \dots, t_n, (\sigma_n, m_n)$ s.t. $(\sigma, m) = (\sigma_n, m_n)$.

Counting Logic. We use $@Loc$ to mean the set $\{@lc \mid lc \in Loc\}$ of boolean variables. Intuitively, $@lc$ evaluates to \mathbf{tt} exactly when the process evaluating it is at location lc . We associate a *counting variable* $(\pi)^\#$ to each predicate π in $\mathbf{predsOf}(@Loc, \mathbf{exprsOf}(S \cup L))$. Intuitively, in a given program configuration, the variable $(\pi)^\#$ counts the number of processes for which the predicate π holds. We let Ω mean the set $\{(\pi)^\# \mid \pi \in \mathbf{predsOf}(@Loc, \mathbf{exprsOf}(S \cup L))\}$ of counting variables. A *counting predicate* is a predicate in $\mathbf{predsOf}(@Loc, \mathbf{exprsOf}(S \cup \Omega))$.

Elements in $\mathbf{exprsOf}(S \cup L)$ and $\mathbf{predsOf}(@Loc, \mathbf{exprsOf}(S \cup L))$ are evaluated wrt. a shared configuration σ and a process configuration (lc, η) . For instance, $val_{\sigma, (lc, \eta)}(v)$ is $\sigma(v)$ if $v \in S$ and $\eta(v)$ if $v \in L$ and $val_{\sigma, (lc, \eta)}(@lc') = (lc = lc')$. We abuse notation and write $val_{\sigma, m}(\omega)$ to mean the evaluation of the counting predicate ω wrt. a configuration (σ, m) . More precisely, $val_{\sigma, m}((\pi)^\#) = \sum_{(lc, \eta) \text{ s.t. } val_{\sigma, (lc, \eta)}(\pi)} m((lc, \eta))$ and $val_{\sigma, m}(v) = \sigma(v)$ for $v \in S$.

Our counting logic is quite expressive as we can capture assertion violations and deadlocks. For location lc , we let $enabled(lc)$ in $\mathbf{predsOf}(\mathbf{exprsOf}(S \cup L))$ define when a process can fire some transition from lc . We capture the violation of an $assert(\pi)$ at some location lc and the deadlock configurations using the following two counting predicates.

$$\omega_{assert} = (@lc \wedge \neg \pi)^\# \geq 1 \qquad \omega_{deadlock} = \bigwedge_{lc \in Loc} (@lc \wedge enabled(lc))^\# = 0$$

5 Relating layers of abstractions

We formally describe in the following the four steps involved in our predicated constrained monotonic abstraction approach (see Fig. 2).

5.1 Predicate abstraction

Given a program $P = (S, L, T)$ and a number of predicates Π on the variables $S \cup L$, we leverage on existing techniques (such as [5]) in order to generate an abstraction in the form of a boolean program $\tilde{P} = (\tilde{S}, \tilde{L}, \tilde{T})$ where all shared and local variables take boolean values. To achieve this, Π is partitioned into three sets Π_{shr} , Π_{loc} and Π_{mix} . Predicates in Π_{shr} only mention variables in S and those in Π_{loc} only mention variables in L . Predicates in Π_{mix} mention both shared and local variables of P . A one to one mapping associates a predicate $\text{origPredOf}(\tilde{v})$ in Π_{shr} (resp. $\Pi_{mix} \cup \Pi_{loc}$) to each \tilde{v} in \tilde{S} (resp. \tilde{L}).

In addition, there are as many transitions in T as in \tilde{T} . For each $(lc \blacktriangleright lc' : stmt)$ in T there is a corresponding $(lc \blacktriangleright lc' : \text{abstOf}(stmt))$ with the same source and destination locations lc, lc' , but with an abstracted statement $\text{abstOf}(stmt)$ that may operate on the variables $\tilde{S} \cup \tilde{L}$. For instance, the statement $(count := count + 1)$ in Fig. 1 is abstracted with the multiple assignment:

$$\left(\begin{array}{l} wait_leq_count, \\ count_leq_wait \end{array} \right) := \left(\begin{array}{l} choose(wait_leq_count, \mathbf{ff}), \\ choose(\neg wait_leq_count \wedge count_leq_wait, wait_leq_count) \end{array} \right) \quad (1)$$

The value of the variable $count_leq_wait$ after execution of the multiple assignment 1 is \mathbf{tt} if $\neg wait_leq_count \wedge count_leq_wait$ holds, \mathbf{ff} if $wait_leq_count$ holds, and is equal to a non deterministically chosen boolean value otherwise. In addition, abstracted statements can mention the local variables of passive processes, i.e., processes other than the one executing the transition. For this, we make use of the variables $\tilde{L}_p = \{\tilde{l}_p | \tilde{l} \text{ in } \tilde{L}\}$ where each \tilde{l}_p denotes the local variable \tilde{l} of passive processes. For instance, the statement $prev := val$ in Fig. 1 is abstracted with the multiple assignment 2. Here, the local variable $prev_leq_val$ of each process other than the one executing the statement (written $prev_leq_val_p$) is separately updated. This corresponds to a broadcast where the local variables of all passive processes need to be updated.

$$\left(\begin{array}{l} prev_leq_val, \\ prev_leq_max, \\ prev_leq_val_p \end{array} \right) := \left(\begin{array}{l} \mathbf{tt}, \\ choose \left(\begin{array}{l} \neg prev_leq_val \quad prev_leq_val \\ \wedge prev_leq_max, \quad \wedge \neg prev_leq_max \end{array} \right), \\ choose \left(\begin{array}{l} \neg prev_leq_val \quad prev_leq_val \\ \wedge prev_leq_val_p, \quad \wedge \neg prev_leq_val_p \end{array} \right) \end{array} \right) \quad (2)$$

Syntax and semantics of boolean programs. The syntax of boolean programs is described below. Variables $\tilde{v}_1, \dots, \tilde{v}_n$ are in $\tilde{S} \cup \tilde{L} \cup \tilde{L}_p$. Predicate π is in

$\text{predsOf}(\tilde{S} \cup \tilde{L})$, and predicates π_1, \dots, π_n are in $\text{predsOf}(\tilde{S} \cup \tilde{L} \cup \tilde{L}_p)$. We further require for the multiple assignment that if $\tilde{v}_i \in \tilde{S} \cup \tilde{L}$ then $\text{vars}(\pi_i) \subseteq \tilde{S} \cup \tilde{L}$.

$$\begin{aligned} \text{prog} &::= (\tilde{s} := (\mathbf{tt} \mid \mathbf{ff} \mid *)^* \text{proc} : (\tilde{l} := (\mathbf{tt} \mid \mathbf{ff} \mid *)^* (\text{lc} \blacktriangleright \text{lc} : \text{stmt})^+ \\ \text{stmt} &::= \text{spawn} \mid \text{join} \mid \pi \mid \tilde{v}_1, \dots, \tilde{v}_n := \pi_1, \dots, \pi_n \mid \text{stmt}; \text{stmt} \end{aligned}$$

Apart from the fact that all variables are now boolean, the main difference of Fig. 4 with Fig. 3 is the *assign* statement. For this, we write $(\tilde{\sigma}, \tilde{\eta}, \tilde{\eta}_p) \xrightarrow{\tilde{v}_1, \dots, \tilde{v}_n := \pi_1, \dots, \pi_n} (\tilde{\sigma}', \tilde{\eta}', \tilde{\eta}'_p)$ to mean that $\tilde{\eta}'_p$ is obtained in the following way. First, we change the domain of $\tilde{\eta}_p$ from \tilde{L} to \tilde{L}_p and obtain the mapping $\tilde{\eta}_{p,1} = \tilde{\eta}_p \left[\left\{ \tilde{l} \leftarrow \tilde{l}_p \mid \tilde{l} \in \tilde{L} \right\} \right]$, then we let $\tilde{\eta}_{p,2} = \tilde{\eta}_{p,1} \left[\left\{ \tilde{v}_i \leftarrow \text{val}_{\tilde{\sigma}, \tilde{\eta}, \tilde{\eta}_{p,1}}(\pi_i) \mid \tilde{v}_i \in \tilde{L}_p \text{ in lhs of the assignment} \right\} \right]$. Finally, we obtain $\tilde{\eta}'_p = \tilde{\eta}_{p,2} \left[\left\{ \tilde{l}_p \leftarrow \tilde{l} \mid \tilde{l} \in \tilde{L} \right\} \right]$. This step corresponds to a broadcast. We write $(\tilde{\sigma}, \tilde{m}) \xrightarrow{\tilde{t}}_{\tilde{P}} (\tilde{\sigma}', \tilde{m}')$ if $(\tilde{\sigma}, \tilde{m}) \xrightarrow{\tilde{t}}_{\tilde{P}} (\tilde{\sigma}', \tilde{m}')$ for some \tilde{t} . A \tilde{P} run \tilde{p} is a sequence $(\tilde{\sigma}_0, \tilde{m}_0), \tilde{t}_1, \dots, \tilde{t}_n, (\tilde{\sigma}_n, \tilde{m}_n)$. The run is *feasible* if $(\tilde{\sigma}_i, \tilde{m}_i) \xrightarrow{\tilde{t}_{i+1}}_{\tilde{P}} (\tilde{\sigma}_{i+1}, \tilde{m}_{i+1})$ for each $i : 0 \leq i < n$ and both $\tilde{\sigma}_0$ and \tilde{m}_0 are initial.

$$\begin{array}{c} \frac{(\tilde{\sigma}, \tilde{\eta}, \tilde{m}) \xrightarrow{\text{stmt}}_{\tilde{P}} (\tilde{\sigma}', \tilde{\eta}', \tilde{m}')}{(\tilde{\sigma}, (\text{lc}, \tilde{\eta}) \oplus \tilde{m}) \xrightarrow{(\text{lc} \blacktriangleright \text{lc}' : \text{stmt})}_{\tilde{P}} (\tilde{\sigma}', (\text{lc}', \tilde{\eta}') \oplus \tilde{m}')} : \text{trans}}{\quad} \quad \frac{\text{val}_{\tilde{\sigma}, \tilde{\eta}}(\pi)}{(\tilde{\sigma}, \tilde{\eta}, \tilde{m}) \xrightarrow{\pi}_{\tilde{P}} (\tilde{\sigma}, \tilde{\eta}, \tilde{m})} : \text{assume}} \\ \\ \frac{(\tilde{\sigma}, \tilde{\eta}, \tilde{m}) \xrightarrow{\text{stmt}}_{\tilde{P}} (\tilde{\sigma}', \tilde{\eta}', \tilde{m}') \text{ and } (\tilde{\sigma}', \tilde{\eta}', \tilde{m}') \xrightarrow{\text{stmt}'}}_{\tilde{P}} (\tilde{\sigma}'', \tilde{\eta}'', \tilde{m}'')}{(\tilde{\sigma}, \tilde{\eta}, \tilde{m}) \xrightarrow{\text{stmt}; \text{stmt}'}}_{\tilde{P}} (\tilde{\sigma}'', \tilde{\eta}'', \tilde{m}'') : \text{sequence}} \\ \\ \frac{\tilde{m}' = (\text{lc}_{\text{ent}}, \tilde{\eta}_{\text{init}}) \oplus \tilde{m}}{(\tilde{\sigma}, \tilde{\eta}, \tilde{m}) \xrightarrow{\text{spawn}}_{\tilde{P}} (\tilde{\sigma}, \tilde{\eta}, \tilde{m}')} : \text{spawn}} \quad \frac{\tilde{m} = ((\text{lc}_{\text{ext}}, \tilde{\eta}') \oplus \tilde{m}')}{(\tilde{\sigma}, \tilde{\eta}, \tilde{m}) \xrightarrow{\text{join}}_{\tilde{P}} (\tilde{\sigma}, \tilde{\eta}, \tilde{m}')} : \text{join}} \\ \\ \frac{\begin{array}{l} \tilde{\sigma}' = \tilde{\sigma} \left[\left\{ \tilde{v}_i \leftarrow \text{val}_{\tilde{\sigma}, \tilde{\eta}}(\pi_i) \mid \tilde{v}_i \in \tilde{S} \right\} \right] \\ \tilde{\eta}' = \tilde{\eta} \left[\left\{ \tilde{v}_i \leftarrow \text{val}_{\tilde{\sigma}, \tilde{\eta}}(\pi_i) \mid \tilde{v}_i \in \tilde{L} \right\} \right] \\ h : \tilde{m} \rightarrow \tilde{m}' \text{ a bijection with } h((\text{lc}_p, \tilde{\eta}_p)) = (\text{lc}_p, \tilde{\eta}'_p) \\ \text{for some } \tilde{\eta}' \text{ s.t. } (\tilde{\sigma}, \tilde{\eta}, \tilde{\eta}_p) \xrightarrow{\tilde{v}_1, \dots, \tilde{v}_n := \pi_1, \dots, \pi_n}_{\tilde{P}} (\tilde{\sigma}', \tilde{\eta}', \tilde{\eta}'_p) \end{array}}{(\tilde{\sigma}, \tilde{\eta}, \tilde{m}) \xrightarrow{\tilde{v}_1, \dots, \tilde{v}_n := \pi_1, \dots, \pi_n}_{\tilde{P}} (\tilde{\sigma}', \tilde{\eta}', \tilde{m}')} : \text{assign}} \end{array}$$

Fig. 4. Semantics of boolean concurrent programs.

Relation between P and \tilde{P} . Given a shared configuration $\tilde{\sigma}$, we write $\text{origPredOf}(\tilde{\sigma})$ to mean the predicate $\bigwedge_{\tilde{s} \in \tilde{S}} (\tilde{\sigma}(\tilde{s}) \Leftrightarrow \text{origPredOf}(\tilde{s}))$. In a similar manner, we write $\text{origPredOf}(\tilde{\eta})$ to mean the predicate $\bigwedge_{\tilde{l} \in \tilde{L}} (\tilde{\eta}(\tilde{l}) \Leftrightarrow \text{origPredOf}(\tilde{l}))$. Observe $\text{vars}(\text{origPredOf}(\tilde{\sigma})) \subseteq S$ and $\text{vars}(\text{origPredOf}(\tilde{\eta})) \subseteq S \cup L$. We abuse notation and write $\text{val}_{\tilde{\sigma}}(\tilde{\sigma})$ (resp. $\text{val}_{\tilde{\eta}}(\tilde{\eta})$) to mean that $\text{val}_{\tilde{\sigma}}(\text{origPredOf}(\tilde{\sigma}))$

(resp. $val_\eta(\text{origPredOf}(\tilde{\eta}))$) holds. We also write $val_{\tilde{\sigma}, \tilde{\eta}}(\pi)$, for a boolean combination π of predicates in Π , to mean the predicate obtained by replacing each π' in $\Pi_{mix} \cup \Pi_{loc}$ (resp. Π_{shr}) with $\tilde{\eta}(\tilde{v})$ (resp. $\tilde{\sigma}(\tilde{v})$) where $\text{origPredOf}(\tilde{v}) = \pi'$. We let $val_m(\tilde{m})$ mean there is a bijection $h : |m| \rightarrow |\tilde{m}|$ s.t. we can associate to each $(lc, \eta)_i$ in m an $(lc, \tilde{\eta})_{h(i)}$ in \tilde{m} such that $val_\eta(\tilde{\eta})$ for each $i : 1 \leq i \leq |m|$. To each \tilde{P} configuration $(\tilde{\sigma}, \tilde{m})$ corresponds a set $\gamma((\tilde{\sigma}, \tilde{m})) = \{(\sigma, m) | val_\sigma(\tilde{\sigma}) \wedge val_m(\tilde{m})\}$, and to each \tilde{P} configuration (σ, m) a singleton $\alpha((\sigma, m)) = \{(\tilde{\sigma}, \tilde{m}) | val_\sigma(\tilde{\sigma}) \wedge val_m(\tilde{m})\}$. We initialize the \tilde{P} variables s.t. for each σ_{init}, m_{init} there are $\tilde{\sigma}_{init}, \tilde{m}_{init}$ s.t. $\alpha((\sigma_{init}, m_{init})) = \{(\tilde{\sigma}_{init}, \tilde{m}_{init})\}$. To each \tilde{P} run $\tilde{\rho} = (\tilde{\sigma}_0, \tilde{m}_0), \tilde{t}_1, \dots, (\tilde{\sigma}_n, \tilde{m}_n)$ corresponds a set of P runs $\gamma(\tilde{\rho}) = \{(\sigma_0, m_0), t_1, \dots, (\sigma_n, m_n) | (\sigma_i, m_i) \in \gamma((\tilde{\sigma}, \tilde{m})), \tilde{t}_i = \text{abstOf}(t_i)\}$, and to each P run $\rho = (\sigma_0, m_0), t_1, \dots, t_n, (\sigma_n, m_n)$ corresponds a set of \tilde{P} runs $\alpha(\rho)$ defined as $\{\alpha((\sigma_0, m_0)), \tilde{t}_1, \dots, \tilde{t}_n, \alpha((\sigma_n, m_n)) | t_i = \text{abstOf}(t_i)\}$.

Definition 1 (predicate abstraction). *Let $P = (S, L, T)$ be a program and $\tilde{P} = (\tilde{S}, \tilde{L}, \tilde{T})$ be its abstraction wrt. Π as described in this Section. The abstraction is said to be effective and sound if \tilde{P} can be effectively computed and to each feasible P run ρ corresponds a non empty set $\alpha(\rho)$ of feasible \tilde{P} runs.*

5.2 Translation into an extended counter machine

Assume a program $P = (S, L, T)$, a set $\Pi \subseteq \text{predsOf}(\text{exprsOf}(S \cup L))$ of predicates and two counting predicates in $\text{predsOf}(@Loc, \text{exprsOf}(S \cup \Omega))$: an invariant ω_{inv} and a target ω_{trgt} . We write $\tilde{P} = (\tilde{S}, \tilde{L}, \tilde{T})$ to mean the boolean abstraction of P wrt. $\Pi \cup \text{atoms}(\omega_{trgt}) \cup \text{atoms}(\omega_{inv})$. Intuitively, this step results in the state reachability problem of an extended counter machine $M_{P, \Pi, \omega_{inv}, \omega_{trgt}}$ that captures reachability of \tilde{P} configurations (abstracting ω_{trgt} P configurations) with \tilde{P} runs that are strengthened wrt. the P counting invariant ω_{inv} .

An extended counter machine M is a tuple $(Q, C, \Delta, Q_{Init}, \Theta_{Init}, q_{trgt})$ where Q is a finite set of states, C is a finite set of counters (i.e., variables ranging over \mathbb{N}), Δ is a finite set of transitions, $Q_{Init} \subseteq Q$ is a set of initial states, Θ_{Init} is an initial set of counters valuations, (i.e., mappings from C to \mathbb{N}) and q_{trgt} is a state in Q . A transition δ in Δ is of the form $[q : op : q']$ where the operation op is either the identity operation nop , a guarded command $grd \Rightarrow cmd$, or a sequential composition of operations. We use a set A of auxiliary variables ranging over \mathbb{N} . These are meant to be existentially quantified when firing the transitions as explained in Fig. 5. A guard grd is a predicate in $\text{predsOf}(\text{exprsOf}(A \cup C))$ and a command cmd is a multiple assignment $c_1, \dots, c_n := e_1, \dots, e_n$ that involves e_1, \dots, e_n in $\text{exprsOf}(A \cup C)$ and pairwise different c_1, \dots, c_n .

A *machine configuration* is a pair (q, θ) where q is a state in Q and θ is a mapping $C \rightarrow \mathbb{N}$. Semantics are given in Fig. 5. A configuration (q, θ) is *initial* if $q \in Q_{Init}$ and $\theta \in \Theta_{Init}$. An M run ρ_M is a sequence $(q_0, \theta_0); \delta_1; \dots (q_n, \theta_n)$. It is *feasible* if (q_0, θ_0) is initial and $(q_i, \theta_i) \xrightarrow{\delta_{i+1}}_M (q_{i+1}, \theta_{i+1})$ for $i : 0 \leq i < n$. We

write \rightarrow_M to mean $\cup_{\delta \in \Delta} \xrightarrow{\delta}_M$. The machine state reachability problem is to decide whether there is an M feasible run $(q_0, \theta_0); \delta_1; \dots (q_n, \theta_n)$ s.t. $q_n = q_{tgt}$.

$$\begin{array}{c}
 \frac{\delta = [q : op : q'] \text{ and } \theta \xrightarrow{op}_M \theta'}{(q, \theta) \xrightarrow{\delta}_M (q', \theta')} : \text{transition} \\
 \\
 \frac{}{\theta \xrightarrow{nop}_M \theta} : \text{nop} \quad \frac{\theta \xrightarrow{op}_M \theta' \text{ and } \theta' \xrightarrow{op'}_M \theta''}{\theta \xrightarrow{op;op'}_M \theta''} : \text{seq} \\
 \\
 \frac{\exists A. \text{val}_\theta(\text{grd}) \wedge \forall i : 1 \leq i \leq n. \theta'(c_i) = \text{val}_\theta(e_i)}{\theta \xrightarrow{\text{grd} \Rightarrow \text{cmd}}_M \theta'} : \text{gcmd}
 \end{array}$$

Fig. 5. Semantics of an extended counter machine

Translation. We describe a machine $(Q, C, \Delta, Q_{Init}, \Theta_{Init}, q_{tgt})$ that captures the behaviour of the program \tilde{P} and encodes reaching abstractions of P configurations satisfying ω_{tgt} in terms of a machine state reachability problem. Each state in Q is either the target state q_{tgt} or is associated to a shared configuration $\tilde{\sigma}$. We write $q_{\tilde{\sigma}}$ to make the association explicit. There is a one to one mapping that associates a process configuration $(lc, \tilde{\eta})$ to each counter $c_{(lc, \tilde{\eta})}$ in C . Transitions are generated as described in Fig. 7. We associate a program configuration $(\tilde{\sigma}, \tilde{m})$ to each machine configuration $\text{enc}_M((\tilde{\sigma}, \tilde{m})) = (q_{\tilde{\sigma}}, \theta)$. Intuitively, states in Q encode values of the shared variables while each counter $c_{(lc, \tilde{\eta})}$ counts the number of processes at location lc and satisfying $\tilde{\eta}$. In other words $\tilde{m}((lc, \tilde{\eta})) = c_{(lc, \tilde{\eta})}$ for each $(lc, \tilde{\eta})$. The encoding of a \tilde{P} run $\rho_{\tilde{P}} = (\tilde{\sigma}_0, \tilde{m}_0); \tilde{t}_1; \dots (\tilde{\sigma}_n, \tilde{m}_n)$ is the set $\{(q_0, \theta_0); \delta_1 \dots; (q_n, \theta_n) \mid (q_i, \theta_i) = \text{enc}_M((\tilde{\sigma}_i, \tilde{m}_i)) \text{ and } \delta_i \in \Delta_{\tilde{t}_i}\}$. The machine encodes the behaviour of the boolean program as specified in Lem. 1.

Lemma 1 (monotonic translation). *Any \tilde{P} configuration $(\tilde{\sigma}, \tilde{m})$ such that $\omega_{tgt}[\text{origPredDf}(\tilde{s})/\tilde{\sigma}(\tilde{s})] \left[(\pi)^\# / \sum_{\{(lc, \tilde{\eta}) \mid \text{val}_{(lc, \tilde{\eta})}(\pi)\}} \tilde{m}((lc, \tilde{\eta})) \right]$ holds is \tilde{P} reachable iff q_{tgt} is M reachable.*

Let us write $\theta \preceq \theta'$ if $\theta(c) \leq \theta'(c)$ for each $c \in C$. Observe that all machine transitions of Fig. 7 are monotonic in that if $(q_1, \theta_1) \xrightarrow{\delta}_M (q_2, \theta_2)$ for some $\delta \in \Delta$, and if $\theta_1 \preceq \theta_3$, then there is a θ_4 such that $\theta_2 \preceq \theta_4$ and $(q_1, \theta_3) \xrightarrow{\delta}_M (q_2, \theta_4)$. Combined with the fact that \preceq is a well quasi ordering over $\mathbb{N}^{|C|}$, we get that:

Lemma 2 (monotonic decidability). *State reachability of any monotonic translation is decidable.*

However, monotonic translations correspond to coarse over-approximations that are incapable of dealing with statements of our counting logic (Sec. 4). Intuitively, bad configurations (such as those where a deadlock occurs) are no more

$$\frac{[q_{\tilde{\sigma}} : op : q_{\tilde{\sigma}'}] \in \Delta}{strong_{\omega_{inv}}(\tilde{\sigma}) = \exists S. \text{origPredOf}(\tilde{\sigma}) \wedge \omega_{inv} \left[(\pi)^{\#} / \sum_{\{(lc, \tilde{\eta}) \mid val_{(lc, \tilde{\eta})}(\pi)\}} c_{(lc, \tilde{\eta})} \right]} [q_{\tilde{\sigma}} : strong_{\omega_{inv}}(\tilde{\sigma}); op; strong_{\omega_{inv}}(\tilde{\sigma}') : q_{\tilde{\sigma}'}] \in \Delta' \quad \text{strengthen}$$

Fig. 6. Strengthening of counter transitions given a counting invariant ω_{inv} .

upward closed wrt. \preceq . This loss of precision makes the verification of such properties out of the reach of techniques solely based on monotonic translations. To regain some of the lost precision, we constrain the runs using counting invariants.

Lemma 3 (strengthened soundness). *Any feasible P run ρ_P has a \tilde{P} feasible run $\rho_{\tilde{P}}$ in $\alpha(\rho_P)$ with an M feasible run in $enc_M(\rho_{\tilde{P}})$, where M is any machine strengthened as described in Fig. 7 and Fig. 6.*

$$\begin{array}{c} \frac{(lc \blacktriangleright lc' : stmt) \text{ and } [(\tilde{\sigma}, \tilde{\eta}) : op : (\tilde{\sigma}', \tilde{\eta}')]_{stmt}}{(q_{\tilde{\sigma}} : c_{(lc, \tilde{\eta})} \geq 1; (c_{(lc, \tilde{\eta})})^{--}; op; (c_{(lc', \tilde{\eta}')})^{++}; q_{\tilde{\sigma}'} \in \Delta_{(lc \blacktriangleright lc' : stmt)}} : \text{transition} \\ \\ \frac{(lc \blacktriangleright lc' : stmt) \text{ and } [(\tilde{\sigma}, \tilde{\eta}) : op : (\tilde{\sigma}', \tilde{\eta}')]_{stmt}}{(q_{\tilde{\sigma}} : \omega_{trgt} [\tilde{s}/\tilde{\sigma}(\tilde{s})] [(\pi)^{\#} / \sum_{(lc, \tilde{\eta}) \models \pi} c_{(lc, \tilde{\eta})}] : q_{trgt}) \in \Delta_{trgt}} : \text{target} \\ \\ \frac{val_{\tilde{\sigma}, \tilde{\eta}}(\pi)}{[(\tilde{\sigma}, \tilde{\eta}) : nop : (\tilde{\sigma}, \tilde{\eta})]_{\pi}} : \text{assume} \frac{}{[(\tilde{\sigma}, \tilde{\eta}) : (c_{(lc_{cent}, \tilde{\eta}_{init})})^{++}; (\tilde{\sigma}, \tilde{\eta})]_{spawn}} : \text{spawn} \\ \\ \frac{[(\tilde{\sigma}, \tilde{\eta}) : op : (\tilde{\sigma}', \tilde{\eta}')]_{stmt} \text{ and } [(\tilde{\sigma}', \tilde{\eta}') : op' : (\tilde{\sigma}'', \tilde{\eta}'')]_{stmt'}}{[(\tilde{\sigma}, \tilde{\eta}) : op; op' : (\tilde{\sigma}'', \tilde{\eta}'')]_{stmt; stmt'}} : \text{sequence} \\ \\ \frac{}{[(\tilde{\sigma}, \tilde{\eta}) : c_{(lc_{ext}, \tilde{\eta}')} \geq 1; (c_{(lc_{ext}, \tilde{\eta}')})^{--}; (\tilde{\sigma}, \tilde{\eta})]_{join}} : \text{join} \\ \\ \tilde{\sigma}' = \tilde{\sigma} \left\{ \tilde{v}_i \leftarrow val_{\tilde{\sigma}, \tilde{\eta}}(\pi_i) \mid \tilde{v}_i \in \tilde{S} \right\} \quad \tilde{\eta}' = \tilde{\eta} \left\{ \tilde{v}_i \leftarrow val_{\tilde{\sigma}, \tilde{\eta}}(\pi_i) \mid \tilde{v}_i \in \tilde{L} \right\} \\ E = \left\{ a_{(lc, \tilde{\eta}_p), (lc, \tilde{\eta}'_p)} \mid (\tilde{\sigma}, \tilde{\eta}, \tilde{\eta}_p) \xrightarrow{\tilde{v}_1, \dots, \tilde{v}_n, := \pi_1, \dots, \pi_n} M (\tilde{\sigma}', \tilde{\eta}', \tilde{\eta}'_p), a_- \in A \right\} \\ Dom_E = \left\{ (lc, \tilde{\eta}_p) \mid a_{(lc, \tilde{\eta}_p), (lc, \tilde{\eta}'_p)} \in E \right\} \quad Img_E = \left\{ (lc, \tilde{\eta}'_p) \mid a_{(lc, \tilde{\eta}_p), (lc, \tilde{\eta}'_p)} \in E \right\} \\ \left[(\tilde{\sigma}, \tilde{\eta}) : \left(\bigcup \left\{ \begin{array}{l} c_{(d)} = \sum_{i \in Img_E} a_{d,i} \mid d \in Dom_E \\ c_{(i)} := \sum_{d \in Dom_E} a_{d,i} \mid i \in Img_E \end{array} \right\} : (\tilde{\sigma}', \tilde{\eta}') \right) \right]_{\tilde{v}_1, \dots, \tilde{v}_n, := \pi_1, \dots, \pi_n} : \text{assign} \end{array}$$

Fig. 7. Translation of the transitions of a boolean program $(\tilde{S}, \tilde{L}, \tilde{T})$, given a counting target ω_{trgt} , to the transitions $\Delta = \bigcup_{t \in \tilde{T}} \Delta_t \cup \Delta_{trgt}$ of a counter machine.

The resulting machine is not monotonic in general and we can encode the state reachability of a two counter machine.

Lemma 4 (strengthened undecidability). *State reachability is in general undecidable after strengthening.*

5.3 Constrained monotonic abstraction and preorder refinement

This step addresses the state reachability problem for an extended counter machine $M = (Q, C, \Delta, Q_{Init}, \Theta_{Init}, q_{trgt})$. As stated in Lem. 4, this problem is in general undecidable for strengthened translations. The idea here [10] is to force monotonicity with respect to a well-quasi ordering \preceq on the set of its configurations. This is apparent at line 7 of the classical working list algorithm Alg. 1. We start with the natural component wise preorder $\theta \preceq \theta'$ defined as $\bigwedge_{c \in C} \theta(c) \leq \theta'(c)$. Intuitively, $\theta \preceq \theta'$ holds if θ' can be obtained by “adding more processes to” θ . The algorithm requires that we can compute membership (line 5), upward closure (line 7), minimal elements (line 7) and entailment (lines 9, 13, 15) wrt. to preorder \preceq , and predecessor computations of an upward closed set (line 7).

If no run is found, then `not_reachable` is returned. Otherwise a run is obtained and simulated on M . If the run is possible, it is sent to the fourth step of our approach (described in Sect. 5.4). Otherwise, the upward closure step $\mathbf{Up}_{\preceq}((q, \theta))$ responsible for the spurious trace is identified and an interpolant I (with $\text{vars}(I) \subseteq C$) is used to refine the preorder as follows: $\preceq_{i+1} := \{(\theta, \theta') \mid \theta \preceq_i \theta' \wedge (\theta \models I \Leftrightarrow \theta' \models I)\}$. Although stronger, the new preorder is again a well quasi ordering and the trace is guaranteed to be eliminated in the next round. We refer the reader to [3] for more details.

Lemma 5 (CMA [3]). *All steps involved in Alg. 1 are effectively computable and each instantiation of Alg. 1 is sound and terminates given the preorder is a well quasi ordering.*

5.4 Simulation on the original concurrent program

A given run of the extended counter machine $(Q, C, \Delta, Q_{Init}, \Theta_{Init}, q_{trgt})$ is simulated by this step on the original concurrent program $P = (S, L, T)$. This is possible because to each step of the counter machine run corresponds a unique and concrete transition of P . This step is classical in counter example guided abstraction refinement approaches. In our case, we need to differentiate the variables belonging to different processes during the simulation. As usual in such frameworks, if the trace turns out to be possible then we have captured a concrete run of P that violates an assertion and we report it. Otherwise, we deduce predicates that make the run infeasible and send them to step 1 (Sect. 5.1).

Theorem 1 (predicated constrained monotonic abstraction). *Assume an effective and sound predicate abstraction. If the constrained monotonic abstraction step returns `not_reachable`, then no configuration satisfying ω_{trgt} is*

```

input : A machine  $(Q, C, \Delta, Q_{Init}, \Theta_{Init}, q_{trgt})$  and a preorder  $\preceq$ 
output: not_reachable or a run  $(q_1, \theta_1); \delta_1; (q_2, \theta_2); \delta_2; \dots \delta_n; (q_{trgt}, \theta)$ 
1 Working :=  $\cup_{e \in \text{Min}_{\preceq}(\mathbb{N}^{|C|})} \{((q_{trgt}, e), (q_{trgt}, e))\}$ , Visited :=  $\{\}$ ;
2 while Working  $\neq \{\}$  do
3    $((q, \theta), \rho)$  =pick a member from Working;
4   Visited  $\cup = \{((q, \theta), \rho)\}$ ;
5   if  $(q, \theta) \in Q_{Init} \times \Theta_{Init}$  then return  $\rho$ ;
6   foreach  $\delta \in \Delta$  do
7      $pre = \text{Min}_{\preceq}(\text{Pre}_{\delta}(\text{Up}_{\preceq}((q, \theta))))$ ;
8     foreach  $(q', \theta') \in pre$  do
9       if  $\theta'' \preceq \theta'$  for some  $((q', \theta''), -)$  in Working  $\cup$  Visited then
10        | continue;
11        | else
12          foreach  $((q', \theta''), -) \in \text{Working}$  do
13            | if  $\theta' \preceq \theta''$  then Working = Working  $\setminus \{((q', \theta''), -)\}$ ;
14            | foreach  $((q', \theta''), -) \in \text{Visited}$  do
15              | if  $\theta' \preceq \theta''$  then Visited = Visited  $\setminus \{((q', \theta''), -)\}$ ;
16            | Working  $\cup = \{((q', \theta'), (q', \theta')); \delta; \rho\}$ 
17 return not_reachable;

```

Algorithm 1: Constrained monotonic abstraction

reachable in P . If a P run is returned by the simulation step, then it reaches a configuration where ω_{trgt} holds. Every iteration of the outer loop terminates given the inner loop terminates. Every iteration of the inner loop terminates.

Notice that there is no general guaranty that we establish or refute the safety property. For instance, it may be the case that one of the loops does not terminate (although each one of their iterations does) or that we need to add predicates relating local variables of two different processes (something the predicate abstraction framework we use cannot express).

6 Experimental results

We report on experiments with our prototype PACMAN(for predicated constrained monotonic abstraction). We have conducted our experiments on an Intel Xeon 2.67GHz processor with 8GB of RAM. To the best of our knowledge, the reported examples which need refinement of monotonic abstraction's preorder cannot be verified by previous techniques; either because the examples require stronger orderings than the usual preorder, or because they involve counting target predicates that are not expressed in terms of violations of program assertions.

All predicate abstraction predicates and counting invariants have been derived automatically. For the counting invariants, we implemented a thread modular analysis operating on the polyhedra numerical domain. This took less than 11 seconds for all the examples we report here. For each example, we report on the

Table 1. Checking assertion violation with PACMAN

example	P	ECM	outer loop		inner loop		results	
			num.	preds.	num.	preds.	time(s)	output
max	5:2:8	18:16:104	4	5	6	2	192	correct
max-bug	5:2:8	18:8:55	3	4	5	2	106	trace
max-nobar	5:2:8	18:4:51	3	3	3	0	24	trace
readers-writers	3:3:10	9:64:121	5	6	5	0	38	correct
readers-writers-bug	3:3:10	9:7:77	3	3	3	0	11	trace
parent-child	2:3:10	9:16:48	3	4	5	2	73	correct
parent-child -nobar	2:3:10	9:1:16	2	1	2	0	3	trace
simp-bar	5:2:9	8:16:123	3	3	5	2	93	correct
simp-nobar	5:2:9	8:7:67	3	2	3	0	13	trace
dynamic-barrier	5:2:8	8:8:44	3	3	3	0	8	correct
dynamic-barrier-bug	5:2:8	8:1:14	2	1	2	0	3	trace
as-many	3:2:6	8:4:33	3	2	6	3	62	correct
as-many-bug	3:2:6	8:1:9	2	1	2	0	2	trace

number of transitions and variables both in P and in the resulting counter machine. We also state the number of refinement steps and predicates automatically obtained in both refinement loops. We made use of several optimizations. For instance, we discarded boolean mappings corresponding to unsatisfiable combinations of predicates, we used automatically generated invariants (such as $(wait \leq count) \wedge (wait \geq 0)$ for the max example in Fig.1) to filter the state space. Such heuristics dramatically helps our state space exploration algorithms.

We report on experiments checking assertion violations in Tab.1 and deadlock freedom in Tab.2. For both cases we consider correct and buggy (by removing the barriers for instance) programs. PACMAN establishes correctness and exhibits faulty runs as expected. The tuples under the P column respectively refer to the number of variables, procedures and transitions in the original program. The tuples under the ECM column refer to the number of counters, states and transitions in the extended counter machine.

Table 2. Checking deadlock with PACMAN

example	P	ECM	outer loop		inner loop		results	
			num.	preds.	num.	preds.	time(s)	output
bar-bug-no.1	4:2:7	7:16:66	4	4	6	2	27	trace
bar-bug-no.2	4:3:8	9:16:95	4	3	4	0	33	trace
bar-bug-no.3	3:2:6	6:16:78	5	4	6	1	21	trace
correct-bar	4:2:7	7:16:62	4	4	6	2	18	correct
ddlck bar-loop	4:2:10	8:8:63	3	2	3	0	16	trace
no-ddlck bar-loop	4:2:9	7:16:78	4	3	4	0	19	correct

7 Conclusions and Future Work

We have presented a technique, predicated constrained monotonic abstraction, for the automated verification of concurrent programs whose correctness depends on synchronization between arbitrary many processes, for example by means of barriers implemented using integer counters and tests. We have introduced a new logic and an iterative method based on combination of predicate, counter and monotonic abstraction. Our prototype implementation gave encouraging results and managed to automatically establish or refute program assertions deadlock freedom. To the best of our knowledge, this is beyond the capabilities of current automatic verification techniques. Our current priority is to improve scalability by leveraging on techniques such as cartesian and lazy abstraction, partial order reduction, or combining forward and backward explorations. We are also aim to generalize to richer variable types.

References

1. P. Abdulla, F. Haziza, and L. Holk. All for the price of few. In R. Giacobazzi, J. Berdine, and I. Mastroeni, editors, *Verification, Model Checking, and Abstract Interpretation*, volume 7737 of *Lecture Notes in Computer Science*, pages 476–495. Springer Berlin Heidelberg, 2013.
2. P. A. Abdulla, K. Čerāns, B. Jonsson, and Y.-K. Tsay. Algorithmic analysis of programs with well quasi-ordered domains. *Information and Computation*, 160:109–127, 2000.
3. P. A. Abdulla, Y.-F. Chen, G. Delzanno, F. Haziza, C.-D. Hong, and A. Rezine. Constrained monotonic abstraction: A cegar for parameterized verification. In *Proc. CONCUR 2010, 21th Int. Conf. on Concurrency Theory*, pages 86–101, 2010.
4. E. Clarke, M. Talupur, and H. Veith. Environment abstraction for parameterized verification. In *Proc. VMCAI '06, 7th Int. Conf. on Verification, Model Checking, and Abstract Interpretation*, volume 3855 of *Lecture Notes in Computer Science*, pages 126–141, 2006.
5. A. F. Donaldson, A. Kaiser, D. Kroening, M. Tautschnig, and T. Wahl. Counterexample-guided abstraction refinement for symmetric concurrent programs. *Formal Methods in System Design*, 41(1):25–44, 2012.
6. A. Downey. *The Little Book of SEMAPHORES (2nd Edition): The Ins and Outs of Concurrency Control and Common Mistakes*. Createspace Independent Pub, 2009.
7. A. Finkel and P. Schnoebelen. Well-structured transition systems everywhere! *Theoretical Computer Science*, 256(1-2):63–92, 2001.
8. A. Kaiser, D. Kroening, and T. Wahl. Dynamic cutoff detection in parameterized concurrent programs. In *Proceedings of CAV*, volume 6174 of *LNCS*, pages 654–659. Springer, 2010.
9. A. Pnueli, J. Xu, and L. Zuck. Liveness with (0,1,infinity)-counter abstraction. In *Proc. 14th Int. Conf. on Computer Aided Verification*, volume 2404 of *Lecture Notes in Computer Science*, 2002.
10. A. Rezine. *Parameterized Systems: Generalizing and Simplifying Automatic Verification*. PhD thesis, Uppsala University, 2008.

A Appendix

In this section the examples of Sec.6 are demonstrated. For simplicity the property which is going to be checked in the input program is reformulated as a statement that goes to lc_{err} which denotes the error location.

A.1 Readers and Writers

```

int readcount := 0
bool lock := tt, writing := ff

main :
   $lc_{ent} \blacktriangleright lc_{ent} :$  spawn(writer)
   $lc_{ent} \blacktriangleright lc_{ent} :$  readcount = 0  $\wedge$  lock; spawn(reader); readcount := readcount + 1
   $lc_{ent} \blacktriangleright lc_{ent} :$  readcount! = 0; spawn(reader); readcount := readcount + 1

reader :
   $lc_{ent} \blacktriangleright lc_{err} :$  writing
   $lc_{ent} \blacktriangleright lc_{ext} :$  readcount = 1; readcount := readcount - 1; lock := tt
   $lc_{ent} \blacktriangleright lc_{ext} :$  readcount! = 1; readcount := readcount - 1

writer :
   $lc_{ent} \blacktriangleright lc_1 :$  lock; lock := ff
   $lc_1 \blacktriangleright lc_2 :$  writing := tt
   $lc_2 \blacktriangleright lc_3 :$  writing := ff
   $lc_3 \blacktriangleright lc_{ext} :$  lock := tt
    
```

Fig. 8. The readers and writers example.

The readers and writers problem is a classical problem. In this problem there is a resource which is shared between several processes. There are two type of processes, one that only read from the resource *reader* and one that read and write to it *writer*. At each time there can either exist several *readers* or only one *writer*. *readers* and *writers* can not exist at the same time.

In Fig.8 a solution to the readers and writers problem with preference to readers is shown. In this approach *readers* wait until there is no *writer* in the critical section and then get the lock that protects that section. We simulate a lock with a boolean variable *lock*. Considering the fact that in our model the transitions are atomic, such simulation is sound. When a *writer* wants to access the critical section, it first waits for the *lock* and then gets it (by setting it to **ff**). Before starting writing, a *writer* sets a flag *writing* that we check later on in a *reader* process. At the end a *writer* unsets *writing* and frees *lock*.

An arbitrary number of *reader* processes can also be spawned. The number of *readers* is being kept track of by the variable *readcount*. When the first *reader*

is going to be spawned (i.e. $readcount = 0$) flag $lock$ must hold. $readcount$ is incremented after spawning each $reader$. Whenever a $reader$ starts execution, it checks flag $writing$ and goes to error if it is set, because it shows that at the same time a $writer$ is writing to the shared resource. When a $reader$ wants to exit, it decrements the $readcount$. The last $reader$ frees the $lock$.

In this example we need a counting invariant to capture the relation between number of $readers$, i.e. $readcount$ and the number of processes in different locations of process $reader$.

A.2 Parent and Child

```

int  $i := 0$ 
bool  $allocated := \mathbf{ff}$ 

 $\underline{main} :$ 
 $lc_{ent} \blacktriangleright lc_{ent} : spawn(parent); i := i + 1$ 
 $lc_{ent} \blacktriangleright lc_{ent} : join(parent); i := i - 1$ 

 $\underline{parent} :$ 
 $lc_{ent} \blacktriangleright lc_1 : allocated := \mathbf{tt}$ 
 $lc_1 \blacktriangleright lc_2 : spawn(child)$ 
 $lc_2 \blacktriangleright lc_3 : join(child)$ 
 $lc_3 \blacktriangleright lc_{ext} : i = 1; allocated := \mathbf{ff}$ 
 $lc_1 \blacktriangleright lc_3 : \mathbf{tt}$ 

 $\underline{child} :$ 
 $lc_{ent} \blacktriangleright lc_{ext} : allocated$ 
 $lc_{ent} \blacktriangleright lc_{err} : \neg allocated$ 

```

Fig. 9. The Parent and Child example.

In the example of Fig.9 a sample nested spawn/join is demonstrated. In this example two types of processes exist. One is $parent$ which is spawned by $main$ and the other one is called $child$ which is spawned by $parent$. The shared variable i is initially 0 and is incremented and decremented respectively when a $parent$ process is spawned and joined. A $parent$ process first sets the shared flag $allocated$ and then either spawns and joins a $child$ process or just moves from lc_1 to lc_3 without doing anything. The $parent$ that sees $i = 1$ unsets the flag $allocated$. A child process goes to error if $allocated$ is not set. This example is error free because one can see that $allocated$ is unset when only one $parent$ exists and that $parent$ has already joined its $child$ or did not spawn any $child$, i.e. no $child$ exists. Such relation between number of $child$ and $parent$ processes as well as variable i can only be captured by appropriate counting invariants and predicate abstraction is incapable of that.

A.3 Simple Barrier

```

int  wait := 0, count := 0
bool enough := ff, flag := *, barrierOpen := ff

main :
  lcent ▶ lc1 : ¬enough; spawn(proc); count := count + 1
  lc1  ▶ lcent : enough := ff
  lc1  ▶ lcent : enough := tt

proc :
  lcent ▶ lc1 : flag := tt
  lc1  ▶ lc2 : flag := ff
  lc2  ▶ lc3 : wait := wait + 1
  lc3  ▶ lc4 : (enough ∧ wait = count); barrierOpen := tt : wait := wait - 1
  lc3  ▶ lc4 : barrierOpen; wait := wait - 1
  lc4  ▶ lcerr : flag
    
```

Fig. 10. Simple Barrier example.

In the example of Fig.10 a simple application of a barrier is shown. *main* process spawns an arbitrary number of *procs* and increments a shared variable *count* that is initially zero and counts the number of *procs* in the program before shared flag *enough* is set. Each *proc* first sets and then unsets shared flag *flag*. The statements in *lc*₂ to *lc*₄ simulate a barrier. Each *proc* first increments a shared variable *wait* which is initially zero. Then the first *proc* that finds out that the condition $(enough \wedge wait = count)$ holds, sets a shared flag *barrierOpen* and goes to *lc*₄. Other *procs* that want to traverse the barrier can the transition *lc*₃ ▶ *lc*₄ : *barrierOpen*. After the barrier a *proc* goes to error if *flag* is unset. One can see that the error state is not reachable in this program because all *procs* have to unset *flag* before any of them can traverse the barrier. To prove that this example is error free, it must be shown that the barrier implementation does not let any process be in locations *lc*_{ent}, *lc*₁ or *lc*₂ where there are processes after barrier, i.e. in locations *lc*₄ and *lc*_{err}. Proving such property requires the relation between number of processes in program locations and variables *wait* and *count* be kept. This is possible when we use counting invariants as introduced in this paper.

A.4 Dynamic Barrier

In a dynamic barrier the number of processes that have to wait at a barrier can change. The way we implemented barriers in this paper makes it easy to capture characteristics of such barriers. In the example of Fig.11 the variables

```

int  $N := *, wait := *, count := *, i := 0$ 
bool  $done := \mathbf{ff}$ 

main :
 $lc_{ent} \blacktriangleright lc_1 : count, wait := N, 0$ 
 $lc_1 \blacktriangleright lc_1 : i! = N, spawn(proc); i := i + 1$ 
 $lc_2 \blacktriangleright lc_3 : i = N \wedge wait = count$ 
 $lc_3 \blacktriangleright lc_3 : join(proc); i := i - 1$ 
 $lc_3 \blacktriangleright lc_4 : i = 0; done := \mathbf{tt}$ 

proc :
 $lc_{ent} \blacktriangleright lc_{ext} : count := count - 1$ 
 $lc_{ent} \blacktriangleright lc_{err} : done$ 

```

Fig. 11. dynamic barrier

corresponding to barrier i.e. $count$ and $wait$ are respectively set to N and 0 in the *main*'s first statement. Then *procs* are spawned as long as the counter i is not equal to N which denotes the total number of *procs* in the system. Each created *proc* decrements $count$ and by doing so it decrements the number of processes that have to wait at the barrier. In this example the barrier is in lc_2 of *main* and can be traversed as usual when $wait = count$ holds and no more *proc* is going to be spawned, i.e. $i = N$. Then *main* can non-deterministically join a *proc* or set flag $done$ if no more *proc* exists.

A.5 As Many

```

int  $count1 := 0, count2 := 0$ 
bool  $enough := \mathbf{ff}$ 

main :
 $lc_{ent} \blacktriangleright lc_1 : spawn(proc1); count1 := count1 + 1$ 
 $lc_1 \blacktriangleright lc_{ent} : spawn(proc2); count2 := count2 + 1$ 
 $lc_{ent} \blacktriangleright lc_2 : enough := \mathbf{tt}$ 

proc1 :
 $lc_{ent} \blacktriangleright lc_1 : enough$ 
 $lc_1 \blacktriangleright lc_{err} : count1 \neq count2$ 

proc2 :
 $lc_{ent} \blacktriangleright lc_1 : enough$ 

```

Fig. 12. As Many

In the example of Fig.12 process *main* spawns as many processes *proc1* as *proc2* and it increments their corresponding counters *count1* and *count2* accordingly. At some point *main* sets flag *enough* and does not spawn any other processes. Processes in *proc1* and *proc2* start execution after *enough* is set. A process in *proc1* goes to error location if $count1 \neq count2$. One can see that error is not reachable because the numbers of processes in the two groups are the same and respective counter variables are initially zero and are incremented with each spawn to represent the number of processes. To verify this example obviously the relation between *count1*, *count2* and number of processes in different locations of *proc1* and *proc2* must be captured.

A.6 Barriers causing deadlock

```

int wait := 0, count := 0, open := 0
bool proceed := ff

main :
  lcent ▶ lcent : spawn(proc); count := count + 1
  lcent ▶ lc1 : proceed := tt

proc :
  lcent ▶ lc1 : wait := wait + 1
  lc1 ▶ lc2 : proceed ∧ wait = count; open := open + 1
  lc1 ▶ lc2 : proceed ∧ wait ≠ count
  lc2 ▶ lc3 : open > 0; open := open - 1
  lc2 ▶ lcerr : open = 0 ∧ (proc@lcent)# = 0 ∧ (proc@lc1)# = 0
    
```

Fig. 13. Buggy Barrier No.1

In Fig.13 a buggy implementation of barrier is demonstrated. This example is based on an example in [6]. The barrier implementation in the book is based on semaphores and in our example the shared variable *open* which is initialized to zero plays the role of a semaphore. A buggy barrier is implemented in program locations *lc_{ent}* to *lc₃*. First process *main* spawns a number of process *proc*, increments the shared variable *count* which is supposed to count the number of *procs* and at the end sets flag *proceed*. A *proc* increments shared variable *wait* which is aimed to count the number of *procs* accumulated at the barrier. *procs* must wait for the flag *proceed* to be set before they can proceed to *lc₂*. Each *proc* that finds out that condition $proceed \wedge wait = count$ holds increments *open*. This lets another process which is waiting at *lc₂* to take the transition $lc_2 \blacktriangleright lc_3$, i.e. traverse the barrier. A deadlock situation is possible to happen in this implementation and that is when one or more processes are waiting for the condition $open > 0$ to hold, but there is no process left at *lc_{ent}* or *lc₁* of *process* which may eventually increment *open*. In this case a *process* goes to error state.

```

int wait := 0, count := 0
bool proceed := ff

main :
  lcent ▶ lcent : spawn(proc1); count := count + 1
  lcent ▶ lcent : spawn(proc2)
  lcent ▶ lcext : proceed := tt

proc1 :
  lcent ▶ lc1 : wait := wait + 1
  lc1 ▶ lc2 : proceed ∧ wait = count
  lc1 ▶ lcerr : proceed ∧ wait ≠ count ∧ (proc1@lcent)# = 0

proc2 :
  lcent ▶ lc1 : wait > 0; wait := wait - 1

```

Fig. 14. Buggy Barrier No.2

In Fig.14 another buggy implementation of a barrier is demonstrated which makes deadlock possible. Process *main* non-deterministically either spawns a *proc1* and increments *count* or spawns a *proc2* or sets flag *proceed*. *proc1* contains a barrier. Each process in *proc1* increments *wait* and then waits at *lc₁* for the barrier condition to hold. A *proc2* decrements *wait* if *wait* > 0. A deadlock happens when at least a *proc2* decrements *wait* which causes the condition in *lc₁ ▶ lc₂* of *proc1* to never hold. We check a deadlock situation in *lc₁ ▶ lc_{err}* of *proc1* which is equivalent to the situation where *proceed ∧ wait ≠ count* does not hold but there exists no process in *lc_{ent}* of *proc1* that can increment *wait*.

```

int wait := 0, count := 0
bool proceed := ff

main :
  lcent ▶ lcent : spawn(proc); count := count + 1
  lcent ▶ lc1 : proceed := tt

proc :
  lcent ▶ lc1 : wait := wait + 1
  lcent ▶ lc1 : wait > 0; wait := wait - 1
  lc1 ▶ lc2 : proceed ∧ wait = count
  lc1 ▶ lcerr : proceed ∧ wait ≠ count ∧ (proc@lcent)# = 0

```

Fig. 15. Buggy Barrier No.3

The buggy implementation of a barrier in Fig.15 is similar to Fig.14, just that this time the *proc* itself may decrement the *wait* and thus make the barrier

condition $proceed \wedge wait = count$ never hold. A deadlock situation is detected similar to the Fig.14.

```

int wait := 0, count := 0, open := 0
bool proceed := ff

main :
   $lc_{ent} \blacktriangleright lc_{ent} : spawn(proc); count := count + 1$ 
   $lc_{ent} \blacktriangleright lc_1 : proceed := \mathbf{tt}$ 

proc :
   $lc_{ent} \blacktriangleright lc_1 : wait := wait + 1$ 
   $lc_1 \blacktriangleright lc_2 : proceed \wedge wait = count; open := open + 1$ 
   $lc_1 \blacktriangleright lc_2 : proceed \wedge wait \neq count$ 
   $lc_2 \blacktriangleright lc_3 : open \geq 1$ 
   $lc_3 \blacktriangleright lc_4 : wait := wait - 1;$ 
   $lc_4 \blacktriangleright lc_{err} : wait = 0 \wedge open = 0$ 
   $lc_4 \blacktriangleright lc_{ent} : wait = 0 \wedge open \geq 1; open := open - 1$ 
   $lc_4 \blacktriangleright lc_{ent} : wait \neq 0$ 
    
```

Fig. 16. Buggy Barrier in Loop

The example in Fig.16 is based on an example in [6]. It demonstrates a buggy implementation of a reusable barrier. Reusable barriers are needed when a barrier is inside a loop. In Fig.16 the loop is formed by backward edges from lc_3 to lc_{ent} . Process *main* spawns *proc* and increments *count* accordingly. Program locations lc_{ent} to lc_3 in *proc* correspond the barrier implementation and are similar to example in Fig.13 and the other transitions make the barrier ready to be reused in the next loop iteration. The example is buggy first because deadlock is possible and second because a processes can continue to next loop iteration while others are still in previous iterations. Deadlock will happen when processes are not able to proceed from lc_4 because $wait = 0$ but $open = 0$, thus they can never take any of the $lc_4 \blacktriangleright lc_{ent}$ edges. For detecting such a deadlock scenario it is essential to capture the relation between shared variables *count* and *wait* with number of *procs* in different locations.