Counting dynamically synchronizing processes

Zeinab Ganjei, Ahmed Rezine, Petru Eles and Zebo Peng

Linköping University Post Print

N.B.: When citing this work, cite the original article.

The original publication is available at www.springerlink.com:

Zeinab Ganjei, Ahmed Rezine, Petru Eles and Zebo Peng, Counting dynamically synchronizing processes, 2016, International Journal on Software Tools for Technology Transfer (STTT), 1-18.

<http://dx.doi.org/10.1007/s10009-015-0411-0>

Copyright: Springer Verlag (Germany)

<http://www.springerlink.com/?MUD=MP>

Postprint available at: Linköping University Electronic Press

<http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-124406>

Counting Dynamically Synchronizing Processes

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

Linköping University, Sweden

Received: date / Revised version: date

Abstract. We address the problem of automatically establishing correctness for programs generating an arbitrary number of concurrent processes and manipulating variables ranging over an infinite domain. The programs we consider can make use of the shared variables to count and synchronize the spawned processes. This allows them to implement intricate synchronization mechanisms, such as barriers. Automatically verifying correctness, and deadlock freedom, of such programs is beyond the capabilities of current techniques. For this purpose, we make use of *counting predicates* that mix counters 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 a nested counter example based refinement scheme for establishing correctness (expressed as non reachability of configurations satisfying counting predicates formulas.) We have implemented a tool (PACMAN, for predicted constrained monotonic abstraction) and used it to perform parameterized verification on several programs whose correctness crucially depends on precisely capturing the number of processes synchronizing using shared variables.

Key words: parameterized verification, counting predicate, barrier synchronization, deadlock freedom, multi-threaded programs, counter abstraction, predicate abstraction, constrained monotonic abstraction

1 Introduction

We focus on automatically establishing synchronization related parameterized correctness. For this, we consider

programs spawning arbitrarily many concurrent processes that use barriers or integer shared variables for counting the number of processes at different stages of their computations. Correctness is stated in terms of Safety properties expressed using *counting predicates*. Counting predicates make it possible to express statements about program variables and counters for the number of processes satisfying some predicates on 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 at certain states.

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, are spawned in order to perform heavy computations in phases. For this reason, when not implemented directly using shared variables, constructs such as (dynamic) barriers are made available in mainstream libraries and programming languages such as `java.util.concurrent`, `java.util.concurrent.Phaser`, `X10` 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 (i.e. the number of processes participating in the barrier is fixed but arbitrarily large), is a non trivial task even in the case where all processes only manipulate boolean variables. In order 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 all processes are waiting at the barrier before any of them is allowed to cross it. This amounts to testing that the number of processes at certain locations is

* In part supported by the 12.04 CENIT project.

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 technique 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 monotonic counter machines, 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 [4, 14]. 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 [6] is to modify the semantics of the entailed tests (e.g., zero tests for barriers), such that processes not satisfying the tests are removed (e.g., zero tests 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 concurrent boolean programs with broadcasts as abstractions of concurrent programs manipulating integer variables (e.g., [11]). Such boolean approximations are monotonic even when the original program (before abstraction) can encode tests on the number of processes and has therefore a non monotonic invariant. 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).

To sum up, our approach consists in combining two nested counter example guided abstraction refinement loops. Each loop operates at a different level of abstraction. We summarize our contributions in the following points.

1. We introduce and propose to use *counting predicates* to express statements about program variables and about the number of processes satisfying given predicates on program variables.
2. We implement the outer loop by leveraging on existing symmetric predicate abstraction techniques. We encode resulting boolean programs in terms of counter machines where reachability of configurations satisfying a *counting predicates formula* is captured as a state reachability problem.
3. We explain how to strengthen the counter machine using *counting invariants*, i.e. counting predicates that hold on all runs. In this work, we automatically generate these invariants using classical thread modular analysis techniques.
4. We leverage on existing constrained monotonic abstraction techniques in order to implement the inner loop and to address the counter machine 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 predicate formulas such as deadlock freedom and assertions. All programs we report on may generate arbitrary many processes.
6. We include all proofs and make use of several examples to clarify the contributions.

The present article is an extended version of a previous conference paper [15].

Related work. Several works consider parameterized verification for concurrent programs. See [20] for a good survey. We report on some relevant recent techniques. The works in [17, 2] explore finite instances and automatically check for cutoff conditions. Except for checking larger instances, it is unclear how to refine entailed abstractions. [7] strengthens the approach of [2], but can not capture global properties that involve relations between number of processes and program variables, as we do in this work. In [12], the authors target verification of Petri nets which are inherently monotonic and generate invariants by weakening SMT formulas. Although we also target coverability, we do it for counter machines obtained by strengthening monotonic systems into non-monotonic ones. It is unclear how to apply this approach to our systems. The work in [19] gives a generalization of the IC3 algorithm and tries to build inductive invariants for well-structured transition systems. It is unclear how to adapt it to the kind of non-monotonic systems that we work with. Similar to [1], we combine auxiliary invariants obtained on certain variables in order to strengthen a reachability analysis. In [13], the authors propose an approach to synthesise counters in order to automatically build correctness proofs from program traces. The approach repeatedly builds safe counting automata and tries to establish that their language includes traces of a program given as a control flow net. Such nets can model arbitrarily many processes sharing global variables. We can also handle local variables and automatically discover relevant predicates by nesting symmetric predicate abstraction loop with a constrained monotonic abstraction loop. In [9], the authors present a highly optimized coverability checking approach for VASs with broadcasts. We need more than coverability of monotonic systems. In [18], the authors adopt symbolic representations that can track inter-thread predicates. This yields a non monotonic system and the authors force

```

int wait, count := 0, 1;
int cross, read := 0, 0;

process :
t0. pc0 → pc0 : spawn
t1. pc0 → pc1 : cross = 0; count := count + 1
t2. pc1 → pc2 : read := 1
// do some reading before the barrier
t3. pc2 → pc3 : read := 0
t4. pc3 → pc4 : wait := wait + 1
t5. pc4 → pc5 : (wait = count); cross := 1
//assert(read = 0)
t6. ...

```

Fig. 1. No matter how many processes are spawned, no process can be at pc_5 and witness $read > 0$.

monotonicity as in [6,5]. They however do not explain how to refine the obtained decidable monotonic abstraction for an undecidable problem. In [8], the authors prove termination for depth-bounded systems by instrumenting a given over-approximation with counters and sending the numerical abstraction to existing termination provers. We automatically generate the abstractions on which we establish safety properties. In addition, and as stated earlier, over-approximating the concurrent programs we target with (monotonic) well structured transition systems would result in spurious runs.

The closest works to ours are [5,11]. We introduced (constrained) monotonic abstraction in [6,5]. Monotonic abstraction was not combined with predicate abstraction, nor did it explicitly target counting properties or dynamic barrier based synchronization. In [11], 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 [11] in order to handle shared and local integer variables.

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 predicates 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 listed in Fig. 1. In this example, processes share four integer variables, namely `wait`, `count`, `read` and `cross`. Variable `count` is initialized to 1, and variables `wait`, `cross` and `read` are initialized to 0. A single process starts executing the program. Arbitrarily many processes get spawned at location pc_0 by transition t_0 .

Each process executes t_1 and increments `count` if `cross` is 0, meaning that no process has crossed the barrier. Otherwise, no process can execute the transition t_1 . It then sets and rests `read`. Intuitively, `read` is greater than zero when there is at least a process doing some reading before the barrier between transitions t_2 and t_3 . Transitions t_4 and t_5 essentially implement a barrier in the sense that all processes must have reached pc_4 in order for any of them to move to location pc_5 . The first process that crosses the barrier changes `cross` from 0 to 1. As a result, no other process can take transition t_1 and start working. After the barrier, no process should be left behind. We capture this by asserting that no process at location pc_5 should witness `read > 0`. We write $@pc_5$ to mean the predicate satisfied by all processes at location pc_5 . A process that satisfies $@pc_5 \wedge (read > 0)$ is at location pc_5 and witnesses `read > 0`. A configuration of the program satisfies the predicate $(@pc_5 \wedge (read > 0))^{\#} \geq 1$ if the number of such processes is greater than or equal to one. We call such a predicate a *counting predicate* (introduced in Sec 4). Counting predicates can be used to capture configurations violating other properties than assertions, e.g., deadlock freedom.

The assertion $(@pc_5 \wedge (read > 0))^{\#} \geq 1$ is never violated under any run starting from a configuration in which a single process starts executing from location pc_0 . In order to establish this fact, any verification procedure needs to take into account the barrier at t_5 as well as the two sources of infiniteness; namely, the infinite domain of the shared and local variables and the number of processes that may participate in the run. Apart from [13] that cannot handle local variables, 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 captures the number of processes at locations pc_1 or later, and that `wait` represents the number of processes at locations pc_4 or later, will not be able to discard scenarios where a process executes `read := 1` although one of them is at pc_5 . Such an analysis will therefore fail to show that `read = 0` each time a process is at pc_5 .

Our tool, 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 which consists of two nested CEGAR loops in the remaining on the example of Fig. 1.

From concurrent programs to boolean concurrent programs. We build on the recent predicate abstraction techniques for concurrent programs. Such techniques first

Fig. 2. Predicated Constrained Monotonic Abstraction

discard all variables and predicates and only keep the control flow. This leads to a number of counter example guided abstraction refinement steps (the outer CEGAR loop in Fig. 2) that will result in the addition of new predicates. Our implementation automatically adds the predicates `cross_leq_0`, `read_leq_0`, `wait_leq_count` and `count_leq_wait`.

It is worth noticing that all variables of the obtained concurrent program are finite state (in fact boolean). Hence, one would need a finite number of counters in order to faithfully capture the behaviour of the abstracted program using counter abstraction.

From concurrent boolean programs to counter machines.

Given a concurrent boolean program and a property to be checked, we generate a counter machine that essentially boils down to a vector addition system with transfers (with additional tests for global properties such as deadlock freedom). Each counter in the machine counts the number of processes at some location with some specific value combination for the local variables. One state in the counter machine represents reaching a configuration violating the property we want to check. The other states correspond to the possible combinations of the global variables. Such a machine cannot relate the number of processes in certain locations to the predicates that are valid at certain states (for instance that `count` = `wait`). These are essential for verification of programs where counters are used to synchronize processes. 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.

$$\text{count} = \sum_{i \geq 1} (\text{@pc}_i)^{\#}$$

$$\text{wait} = \sum_{i>4} (@pc_i)^{\#}$$

We automatically generate such invariants using a simple thread modular analysis that tracks the number of processes satisfying some property.

Example 1 (Thread modular analysis). To recover the two counting invariants above, we can perform a classi-

cal thread modular analysis [16] where we add a shared instrumentation variable `_pci` to track the number of processes at location `pci` for $i : 0 \leq i \leq 6$. We use a suitable abstract numerical domain (in this case polyhedral domain). For the program of Fig. 1, we obtained the counting invariants mentioned earlier as well as other invariants such as $0 \leq \text{count}$, $0 \leq \text{wait}$ and $\text{wait} \leq \text{count}$ that helped for pruning the state space as mentioned in Sec. 6.

Given such counting invariants, we constrain the counter machine and generate a more precise machine that may not be a vector addition system anymore. We explain in Sect. 5 that the resulting state reachability problem is now undecidable in general.

Constrained monotonic abstraction. We monotonically abstract the resulting counter machine in order to answer the state reachability problem. Spurious traces are now possible. Essentially, monotonic abstraction closes upwards the obtained sets of predecessor configurations. This over-approximation might add larger configurations that did not belong to the set of predecessor configurations. Intuitively, the effect of monotonic abstraction “in forward” on the example of Fig.1 is that it “deletes” processes violating the constraint imposed by the barrier [6]. This example illustrates a situation where such approximations yield false positives. To see this, suppose two processes exist. A first process gets to $\mathbf{pc}_4$ and waits. The second process moves to $\mathbf{pc}_2$. Deleting the second process, is allowed by the monotonic abstraction and opens the barrier for the first process. However, the assertion can now be violated because the deleted process did not have time to reset the variable `read`. 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 $\mathbf{pc}_1$ is zero, then closing upwards will not alter this fact. By doing so, the process that was deleted in forward at $\mathbf{pc}_2$ 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 perform 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 reachability problem 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 $\mathbb{N}$ and $\mathbb{Z}$ to mean the sets of natural and integer values respectively. Given two natural numbers $i, j \in \mathbb{N}$, we use $[i, j]$ to denote the set $\{k \in \mathbb{N} \mid i \leq k \leq j\}$. We let $\mathbb{B} = \{\mathbf{tt}, \mathbf{ff}\}$ be the set of boolean values. In this section, we write V and V^b to respectively mean a set of integer and boolean variables. We write X to mean some set V or V^b . In a similar manner, we write v and v^b to respectively mean an integer or a boolean variable. We also write x to mean a variable of some type.

We write $\mathbf{exprsOf}(V)$ to mean the set of arithmetic expressions over the integer variables V . An arithmetic expression e (or expression for short) in $\mathbf{exprsOf}(V)$ is an integer constant k , an integer variable v , or the sum or difference of two expressions as described below:

$$e ::= k \mid v \mid (e + e) \mid (e - e) \mid k e \quad v \in V$$

We let $\sim$ be a comparator in $\{<, \leq, =, \geq, >\}$. We write $\mathbf{predsOf}_E^{V^b}$ to mean the set of predicates (i.e., boolean expressions) over boolean variables V^b and arithmetic expressions E . A predicate π in $\mathbf{predsOf}_E^{V^b}$ is either a boolean value $\mathbf{b}$, a variable v^b in V^b , a comparison of two expressions in E or a boolean combination of predicates. It takes the following form:

$$\pi ::= \mathbf{b} \mid v^b \mid (e \sim e) \mid \neg \pi \mid \pi \wedge \pi \mid \pi \vee \pi \quad v^b \in V^b, e \in E$$

We write $\mathbf{vars}(e)$ to mean all integer variables appearing in an expression e , and $\mathbf{vars}(\pi)$ to mean all variables appearing in π , namely both boolean variables appearing in π and all integer variables in $\mathbf{vars}(e)$ for each e appearing in π . We also write $\mathbf{comparisonsOf}(\pi)$ to mean all comparisons $(e \sim e)$ appearing in π . We assume in the following an arithmetic or boolean expression $\mathbf{exp}$ or any indexed version of it. A mapping $\mathbf{x} : X \rightarrow Y$ associates an expression to each variable in X . Expressions in Y have the same type as the variables in X . We often write a mapping $\mathbf{x} : X \rightarrow Y$ as the set $\{x \leftarrow \mathbf{x}(x) \mid x \in X\}$. We write $\mathbf{exp}[\mathbf{x}]$ to mean the evaluation of expression $\mathbf{exp}$ with respect to a mapping $\mathbf{x}$. We perform the evaluation as follows. First, an expression $\mathbf{exp}_{tmp}$ is deduced by syntactically and simultaneously replacing in $\mathbf{exp}$ each occurrence of a variable $x \in X$ by the corresponding $\mathbf{x}(x)$. Then, if $\mathbf{vars}(\mathbf{exp}_{tmp}) = \emptyset$, $\mathbf{exp}[\mathbf{x}]$ is the constant obtained by evaluating $\mathbf{exp}_{tmp}$. Otherwise, $\mathbf{exp}[\mathbf{x}]$ is taken to be $\mathbf{exp}_{tmp}$ (see Ex. 2). Let $\mathbf{x} : X \rightarrow Y$ and $\mathbf{x}' : X' \rightarrow Y'$ be two mappings. If X and X' are disjoint, we write $\mathbf{exp}[\mathbf{x}, \mathbf{x}']$ to mean the evaluation of expression $\mathbf{exp}$ with respect to $\mathbf{x} \cup \mathbf{x}'$. Larger unions of mappings with pairwise disjoint domains are handled in a similar manner. We abuse notation and write $\mathbf{x}[\mathbf{x}']$ to mean $\{x \leftarrow \mathbf{x}(x) \mid x \in X \setminus X'\} \cup \{x \leftarrow \mathbf{x}'(x) \mid x \in X \cap X'\}$.

$$\begin{aligned} \mathbf{prog} &::= (\mathbf{s} := (k \mid \star))^* \\ \mathbf{process} &: \\ &(\mathbf{l} := (k \mid \star))^* \\ &(\mathbf{pc} \rightarrow \mathbf{pc} : \mathbf{stmt})^+ \\ \mathbf{stmt} &::= \mathbf{v}_1, \dots, \mathbf{v}_n := \mathbf{e}_1, \dots, \mathbf{e}_n \\ &\mid \mathbf{spawn} \mid \mathbf{join} \mid \pi \mid \mathbf{stmt}; \mathbf{stmt} \end{aligned}$$

Fig. 3. Syntax of concurrent programs: $\mathbf{s}$ is a shared variable in S , $\mathbf{l}$ is a local variable in L , $\mathbf{v}_1, \dots, \mathbf{v}_n$ are pairwise different variables in $S \cup L$, $\mathbf{e}_1, \dots, \mathbf{e}_n$ are arithmetic expressions in $\mathbf{exprsOf}(S \cup L)$ and π is a predicate in $\mathbf{predsOf}_{\mathbf{exprsOf}(S \cup L)}$.

Example 2 (Expressions). Let $\pi = v^b \wedge (v' = v' + 1)$. Then, $\mathbf{vars}(\pi) = \{v^b, v, v'\}$. If we let $\mathbf{x}^b = \{v^b \leftarrow \mathbf{tt}\}$, $\mathbf{x}_1 = \{v \leftarrow v' + 5, v' \leftarrow 3\}$ and $\mathbf{x}_2 = \{v' \leftarrow -1\}$, then $\pi[\mathbf{x}^b, \mathbf{x}_1] = \mathbf{tt} \wedge (v' + 5 = 3 + 1)$ and $\pi[\mathbf{x}^b, \mathbf{x}_1][\mathbf{x}_2] = \mathbf{tt}$.

A multiset $\mathfrak{m}$ over a set Σ is a mapping $\Sigma \rightarrow \mathbb{N}$. We sometimes write a multiset $\mathfrak{m}$ by enumerating its elements in some predefined total order on Σ , i.e. as $[\sigma, \dots]$ where each $\sigma \in \Sigma$ appears exactly $\mathfrak{m}(\sigma)$ times. A bijection from a multiset $[\sigma_1, \sigma_2, \dots, \sigma_n]$ to a multiset $[\sigma'_1, \sigma'_2, \dots, \sigma'_n]$ of the same size is a bijection from $[1, n]$ to $[1, n]$ that associates each element of the former multiset to an element of the latter. The size $|\mathfrak{m}|$ of a multiset $\mathfrak{m}$ is $\sum_{\sigma \in \Sigma} \mathfrak{m}(\sigma)$. We write $\sigma \oplus \mathfrak{m}$ to mean the multiset $\mathfrak{m}'$ such that $\mathfrak{m}'(\sigma')$ equals $\mathfrak{m}(\sigma) + 1$ if $\sigma = \sigma'$ and $\mathfrak{m}(\sigma)$ otherwise.

Example 3 (Multisets). Let $\Sigma = \{a, b, c\}$, and define $\mathfrak{m} = [a, a]$ and $\mathfrak{m}' = [c, c]$. A possible bijection from multiset $b \oplus \mathfrak{m} = [a, a, b]$ to multiset $b \oplus \mathfrak{m}' = [b, c, c]$ is the mapping $\{1 \leftarrow 2, 2 \leftarrow 1, 3 \leftarrow 3\}$ that sends the first a to the first c , the second a to b and b to the second c .

4 Concurrent Programs and Counting Predicates.

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. The procedure of a program $P = (S, L, T)$ is given in terms of a finite set T of transitions, each of the form $(\mathbf{pc}_1 \rightarrow \mathbf{pc}'_1 : \mathbf{stmt}_1)$. Transitions operate on two finite sets of integer variables, namely a set S of shared variables and a set L of local variables. Each transition $(\mathbf{pc} \rightarrow \mathbf{pc}' : \mathbf{stmt})$ involves two program locations $\mathbf{pc}$ and $\mathbf{pc}'$ and a statement $\mathbf{stmt}$. We write PC to mean the set of all locations appearing in T . We always distinguish two locations, namely an entry location $\mathbf{pc}_0$ and an exit location $\mathbf{pc}_x$. Program syntax is described in Fig. 3.

$$\begin{array}{c}
\frac{(s, \ell, m) \xrightarrow[P]{\text{stmt}} (s', \ell', m')}{(s, (pc, \ell) \oplus m) \xrightarrow[P]{(pc \rightarrow pc' : \text{stmt})} (s', (pc', \ell') \oplus m')} : \text{transition} \\
\\
\frac{(s, \ell, m) \xrightarrow[P]{\text{stmt}} (s', \ell', m') \text{ and } (s', \ell', m') \xrightarrow[P]{\text{stmt}'} (s'', \ell'', m'')}{(s, \ell, m) \xrightarrow[P]{\text{stmt}; \text{stmt}'} (s'', \ell'', m'')} : \text{sequence} \\
\\
\frac{v = \{v_i \leftarrow e_i[s, \ell] \mid 1 \leq i \leq n\}}{(s, \ell, m) \xrightarrow[P]{v_1, \dots, v_n := e_1, \dots, e_n} (s[v], \ell[v], m)} : \text{assign} \\
\\
\frac{\pi[s, \ell]}{(s, \ell, m) \xrightarrow[P]{\pi} (s, \ell, m)} : \text{assume} \\
\\
\frac{m' = (pc_0, \ell_{init}) \oplus m \text{ with } \ell_{init} \in \mathbb{L}_{init}}{(s, \ell, m) \xrightarrow[P]{\text{spawn}} (s, \ell, m')} : \text{spawn} \\
\\
\frac{m = ((pc_x, \ell') \oplus m')}{(s, \ell, m) \xrightarrow[P]{\text{join}} (s, \ell, m')} : \text{join}
\end{array}$$

Fig. 4. Semantics of concurrent programs. Executions start from some (pc_0, m_{init}) with $m_{init} \in \mathbb{M}_{init}$.

Semantics. Initially, a single process starts executing the procedure with both local and shared variables initialized as stated in their respective declarations. Executions might involve an arbitrary number of spawned processes. The execution of any process (whether initial or spawned with a **spawn** statement) starts at the entry location pc_0 with the corresponding local variables initialized as stated in their respective declarations. Any process at an exit point pc_x 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 by a single process without interruption from other processes.

More formally, a configuration is given in terms of a shared state and a processes configuration. A *shared state* $s : S \rightarrow \mathbb{Z}$ is a mapping that associates an integer to each variable in S . We write $\mathbb{S}$ to mean the set of all shared states. An *initial shared state* is a mapping in $\mathbb{S}$ that respects shared variables declarations. We write $\mathbb{S}_{init}$ to mean the set of all initial shared states. A *process state* is a pair (pc, ℓ) where the location pc belongs to PC and the *local state* $\ell : L \rightarrow \mathbb{Z}$ maps each local variable to an integer number. We also write $\mathbb{L}$ and $\mathbb{L}_{init}$ to respectively mean the sets of all local states and all initial local states. A *processes configuration* is a multiset m over process states. An *initial processes configuration* maps all (pc, ℓ) to 0 except for a single (pc_0, ℓ) , with $\ell \in \mathbb{L}_{init}$, mapped to 1. We write $\mathbb{M}$ and $\mathbb{M}_{init}$ to mean the sets of all processes configurations and initial processes configurations respectively. Finally, a *configuration* is a pair (s, m) consisting of a shared state s and a

processes configuration m . We write $(s, m) \xrightarrow[P]{t} (s', m')$ to mean that the transition t of the form $(pc \rightarrow pc' : \text{stmt})$ applies atomically to configuration (s, m) and changes it to (s', m') . We introduce a relation $\xrightarrow[P]{\text{stmt}}$ in order to describe the steps involved in the semantics of transitions (Fig. 4). We write $(s, \ell, m) \xrightarrow[P]{\text{stmt}} (s', \ell', m')$, where s, s' are shared states, ℓ, ℓ' are local states, and m, m' are multisets of process configurations, in order to mean that a process of the program P at local state ℓ when the shared state is s and the other process configurations are captured by m , can execute the statement **stmt** and take the program to a configuration where the process has local state ℓ' , the shared state is s' 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 pc_x (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*). A P run ρ is a configuration starting alternating sequence of transitions and configurations $(s_0, m_0) t_1 \dots t_n (s_n, m_n)$. The run is P feasible if $(s_i, m_i) \xrightarrow[P]{t_{i+1}} (s_{i+1}, m_{i+1})$ for each $i : 0 \leq i < n$ and s_0 and m_0 are initial, i.e., $s_0 \in \mathbb{S}_{init}$ and $m_0 \in \mathbb{M}_{init}$. Each configuration (s_i, m_i) for $i : 0 \leq i \leq n$ is then said to be *reachable*.

Example 4 (Feasible run). Consider the concurrent program in Fig. 1. An initial configuration is (s_0, m_0) where $s_0 = \{\text{count} \leftarrow 1, \text{wait} \leftarrow 0, \text{cross} \leftarrow 0, \text{read} \leftarrow 0\}$ and the initial multiset m_0 associates 1 to the unique process state with location pc_0 and 0 to all other process states. A feasible run is then represented below.

s				m
count	wait	cross	read	pc
1	0	0	0	pc ₀
$\Downarrow t_0.pc_0 \rightarrow pc_0 : \text{spawn}$				
1	0	0	0	pc ₀ , pc ₀
$\Downarrow t_0.pc_0 \rightarrow pc_0 : \text{spawn}$				
1	0	0	0	pc ₀ , pc ₀ , pc ₀
$\Downarrow t_1.pc_0 \rightarrow pc_1 : \text{cross} := 0; \text{count} := \text{count} + 1$				
2	0	0	0	pc ₀ , pc ₀ , pc ₁
$\Downarrow t_2.pc_1 \rightarrow pc_2 : \text{read} := 1$				
2	0	0	1	pc ₀ , pc ₀ , pc ₂
$\Downarrow t_3.pc_2 \rightarrow pc_3 : \text{read} := 0$				
2	0	0	0	pc ₀ , pc ₀ , pc ₃

Counting Predicates. Recall that PC is the set of program locations. We make use of a set of boolean variables $\{\text{@pc} \mid pc \in PC\}$, denoted @PC . Intuitively, a process evaluates @pc to **tt** exactly when it is at location pc . This way, we can build boolean expressions in the set $\text{predsOf}_{\text{exprsOf}(S_{UL})}^{\text{@PC}}$. With these predicates we can state facts about both the location of some process, as well as its own local variables and the values of the shared variables. For instance, at the fourth step of the run depicted

in Ex. 4, there is one process for which $(@pc_2 \wedge \text{count} \geq 1)$ holds.

We associate a *counting variable* $(\pi)^\#$ to each predicate π in $\text{predsOf}_{\text{exprsOf}(S \cup L)}^{\text{PC}}$. Intuitively, in a given program configuration, the variable $(\pi)^\#$ counts the number of processes for which the predicate π holds. We denote the counting variables $\{(\pi)^\# \mid \pi \in \text{predsOf}_{\text{exprsOf}(S \cup L)}^{\text{PC}}\}$ with $\Omega_{PC,S,L}$. For example, $(@pc_2 \wedge \text{count} \geq 1)^\#$ is a variable that counts the number of processes for which $(@pc_2 \wedge \text{count} \geq 1)$ holds. Such counting a variable is evaluated with respect to a shared state s and a process configuration (pc, ℓ) . We abuse notation and write $v[s, (pc, \ell)]$ to mean the variable v participating in a counting variable, is evaluated to $s(v)$ if $v \in S$, to $\ell(v)$ if $v \in L$, or to $(pc = pc')$ if v is the boolean variable $@pc'$.

Any predicate in $\text{predsOf}_{\text{exprsOf}(S \cup \Omega_{PC,S,L})}$ is a *counting predicate*. We need a shared state s and a processes configuration m in order to evaluate a counting variable (Eq.1) or a counting predicate (Eq.2). We abuse notation and write $\omega[s, m]$ to mean the evaluation of the counting predicate ω wrt. a configuration (s, m) . The evaluation is performed as follows. Given a configuration (s, m) , a shared variable $s \in S$ is evaluated as usual to $s[s, m] = s(s)$; whereas the counting variable $(\pi)^\#$ is evaluated to the number of processes satisfying π in (s, m) .

$$(\pi)^\# [s, m] = \sum_{\{(pc, \ell) \text{ s.t. } \pi[s, (pc, \ell)]\}} m((pc, \ell)) \quad (1)$$

$$\omega[s, m] = \omega[s, \{(\pi)^\# \leftarrow (\pi)^\# [s, m] \mid (\pi)^\# \in \text{vars}(\omega)\}] \quad (2)$$

Our counting predicates are quite expressive. For instance, we can capture assertion violations, deadlocks or rich program invariants with them (see Ex. 5). For any pc , we can define a counting predicate $\text{isEnabled}(pc)$ that captures whether a process currently at location pc can fire some transition. For instance, in the running example of Fig. 1, $\text{isEnabled}(pc_0) = \text{true}$ and $\text{isEnabled}(pc_4) = (\text{wait} = \text{count})$. If there would have been only one transition from pc_6 consisting in a join operation, then $\text{isEnabled}(pc_6)$ would have been $(@pc_x \geq 1)^\#$.

Example 5 (Counting predicates). The following counting predicates capture configurations of the program of Fig. 1: ω_1 captures configurations that violate the assertion, ω_2 captures those where a deadlock occurs, and ω_3 an over-approximation of the reachable configurations (i.e., an invariant).

$$\omega_1 = (@pc_5 \wedge (\text{read} > 0))^\# \geq 1$$

$$\omega_2 = \bigwedge_{pc \in PC} (@pc \wedge \text{isEnabled}(pc))^\# = 0$$

$$\omega_3 = (\text{count} = \sum_{i \geq 1} (@pc_i)^\#) \wedge (\text{wait} = \sum_{i \geq 4} (@pc_i)^\#)$$

```

bool read.leq.0, wait.leq.count := tt, tt

process :
t0. pc0 → pc0 : spawn
t1. pc0 → pc1 : (tt); wait.leq.count :=
                ch(wait.leq.count, ff)
t2. pc1 → pc2 : read.leq.0 := ff
t3. pc2 → pc3 : read.leq.0 := tt
t4. pc3 → pc4 : wait.leq.count :=
                ch(ff, ¬wait.leq.count)
t5. pc4 → pc5 : (wait.leq.count); (tt)
t6. ...

```

Fig. 5. Predicate abstraction of the program in Fig. 1 with respect to the predicates $\Pi = \{\text{read.leq.0}, \text{wait.leq.count}\}$.

5 Relating abstraction layers

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 predicate abstraction technique in [11] in order to generate an abstraction in the form of a boolean program $\text{abst}_\Pi(P) = (S^b, L^b, T^b)$ where all shared and local variables are boolean. 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 bijection associates a predicate $\text{originOf}(v^b)$ in Π_{shr} (resp. $\Pi_{mix} \cup \Pi_{loc}$) to each boolean variable v^b in S^b (resp. L^b). The function is a bijection as each predicate will be associated to one and only one boolean variable (that tracks the value of that predicate) and vice versa.

Example 6 (Predicate abstraction). Consider the concurrent program in Fig. 1. We implemented the predicate abstraction of [11] which results, for the predicates $\Pi = \{\text{read.leq.0}, \text{wait.leq.count}\}$, in the boolean program of Fig. 5. A bijection associates the predicates $\text{wait} \leq \text{count}$ and $\text{read} \leq 0$ in Π_{shr} respectively to the boolean variables wait.leq.count and read.leq.0 in S^b of the boolean program.

In addition, there are as many transitions in T as in T^b . For each $(pc \rightarrow pc' : \text{stmt})$ in T there is a corresponding $(pc \rightarrow pc' : \text{abst}_\Pi(\text{stmt}))$ with the same source and destination locations pc, pc' , but with an abstracted statement $\text{abst}_\Pi(\text{stmt})$ that may operate on the variables $S^b \cup L^b$. Moreover, abstracted statements can mention the local variables of passive processes, i.e., processes other than the one executing the transition. For


```

prog ::= (sb := (tt | ff | *))*
      process :
      (lb := (tt | ff | *))*
      (pc → pc : stmt)+
stmt ::= v1b, ..., vnb := ch(π1, π'1), ..., ch(πn, π'n)
      spawn | join | π | stmt; stmt

```

Fig. 6. Syntax of concurrent boolean programs.

this, we make use of the variables $L_p^b = \{l_p^b | l^b \text{ in } L^b\}$ where each l_p^b denotes local variable l^b of passive processes. We use passive variables to capture broadcasts where local variables of all passive processes need to be updated. Note that such passive variables and broadcast transitions do not exist in the original concurrent programs to be verified, but are introduced after predicate abstraction of those programs as presented in [11]. They are essential to capture the behaviour of the processes existing in the system other than the process that actually executes a transition (resp. passive and active processes).

Example 7 (Broadcast transition). Consider a concurrent program with shared and local variables s and l , the assignment transition $t_0 :: pc_1 \rightarrow pc_2 : s := l$ and the mixed predicate $mx :: (s = l)$ to be used for the predicate abstraction. Recall from Sec. 5.1 that such predicates are called *mixed predicates* as they contain both local and shared variables. Each process will have its own copy of a mixed predicate, similar to local predicates. However, unlike local predicates, mixed predicates are updated only by broadcast transitions. Before the assignment t_0 , consider two passive processes P_1 and P_2 having mixed predicates $mx_{p_1} :: (s = l_{p_1})$ evaluates to tt and $mx_{p_2} :: (s = l_{p_2})$ evaluates to ff and an active process P_a having $mx_{p_a} :: (s = l_a)$ evaluates to tt . The active process will execute t_0 , hence, $s = l_a$ will hold after the transition. At this point, mx_{p_1} will hold only if $l_a = l_{p_1}$. But, mx_2 will *not* hold. In fact, after the transition, all passive processes will be notified to update their corresponding mixed predicate w.r.t their own valuation and that of the active process. This corresponds to a broadcast (More details in Ex. 8).

Syntax and semantics of boolean programs. We describe the syntax of boolean programs in Fig. 6. Variables s^b and l^b are some variables respectively in S^b and L^b . Variables $v_1^b, \dots, v_n^b$ are pairwise different and belong to $S^b \cup L^b \cup L_p^b$. Predicate π is in $\text{predsOf}^{S^b \cup L^b}$ (i.e., a boolean combination of boolean variables in $S^b \cup L^b$) and predicates π_i, π'_i are in $\text{predsOf}^{S^b \cup L^b \cup L_p^b}$. By construction, the predicate abstraction of concurrent programs [11] does not involve assignments of passive variables to non-passive ones.

We describe semantics of boolean programs in Fig. 7. We add the superscript b to mean the boolean program variant. For instance, we use S^b, L^b and M^b to respectively mean the sets of shared states, local states and processes configurations of boolean programs. The $\text{ch}(\pi, \pi')$ operator evaluates to tt if π evaluates to tt , to ff if π' evaluates to tt , and non-deterministically to either tt or ff otherwise. Apart from the fact that all variables are now boolean and that we make use of the ch operator, the main difference of Fig. 7 with Fig. 4 is the **assign** statement as it may involve passive variables in order to capture broadcasts.

Example 8 (ch operator). Consider once again Ex. 7. Using mx , the assignment transition t_0 will be abstracted to $t_0^b :: pc_1 \rightarrow pc_2 : mx, mx_p := tt, \text{ch}(mx_p \wedge mx, mx_p \oplus mx)$ ($\oplus$ is exclusive-or). Based on the semantics of ch operator, the variable mx_p will evaluate to tt if both mx_p and mx held before the assignment, will evaluate to ff if $mx_p \oplus mx$ held before the assignment, and evaluates to a non deterministic boolean value otherwise.

Let $v_1^b, \dots, v_n^b := \text{ch}(\pi_1, \pi'_1), \dots, \text{ch}(\pi_n, \pi'_n)$ be an assignment. When describing the semantics in Fig. 7, we write $(s^b, l^b, l_p^b) \xrightarrow[v_1^b, \dots, v_n^b := \text{ch}(\pi_1, \pi'_1), \dots, \text{ch}(\pi_n, \pi'_n)]{\text{abst}_{II}(P)} (s^{b'}, l^{b'}, l_p^{b'})$ in order

to mean that a process other than the one performing the assignment (i.e. a passive process) with local state l_p^b can move to $l_p^{b'}$ when the active process moves from l^b to $l^{b'}$ and the shared state changes from s^b to $s^{b'}$. The new shared state and the new local state of the active process are obtained from their older versions according to (3) and (4). The local state $l_p^{b'}$ of the passive process is obtained as follows. First, we change the domain of l_p^b from L^b to L_p^b and obtain $l_{p,1}^b$ as described in (5). Then we apply the assignment to obtain $l_{p,2}^b$ according to (6). Finally, we obtain $l_p^{b'}$ by changing the domain of $l_{p,2}^b$ from L_p^b back to L^b . Notice that at least one $l_p^{b'}$ exists for each l_p^b . Intuitively, this step corresponds to updating, during a broadcast, where an active process changes local state together with possibly many passive processes.

$$s^{b'} = s^b[\{v_i^b \leftarrow \text{ch}(\pi_i, \pi'_i) [s^b, l^b] | v_i^b \in S^b \wedge i : 1 \leq i \leq n\}] \quad (3)$$

$$l^{b'} = l^b[\{v_i^b \leftarrow \text{ch}(\pi_i, \pi'_i) [s^b, l^b] | v_i^b \in L^b \wedge i : 1 \leq i \leq n\}] \quad (4)$$

$$l_{p,1}^b = \{l_p^b \leftarrow l_p^b(l^b) | l^b \in L^b\} \quad (5)$$

$$l_{p,2}^b = l_{p,1}^b[\{v_i^b \leftarrow \text{ch}(\pi_i, \pi'_i) [s^b, l^b, l_{p,1}^b] | v_i^b \in L_p^b \wedge i : 1 \leq i \leq n\}] \quad (6)$$

Example 9 (Updating passive local state). Consider once again Ex. 7, Ex. 8, the active local state $l^b = \{mx \leftarrow tt\}$ and a passive local state $l_p^b = \{mx \leftarrow tt\}$. The assignment in boolean transition t_0^b involves both variables mx and mx_p . We use l^b and l_p^b for evaluating respectively mx and mx_p , but the domain of the passive local state l_p^b is L^b based on the syntax of boolean concurrent programs. We need to change its domain to L_p^b in order to be able

to distinguish between passive and non-passive variables. By doing this, we obtain $\mathbb{p}_{p,1} = \{\mathbf{mx}_p \leftarrow \mathbf{tt}\}$. Then, we update $\mathbb{p}_{p,1}$ by the result of the assignment and will have $\mathbb{p}_{p,2} = \{\mathbf{mx}_p \leftarrow \mathbf{ff}\}$. Finally, we obtain $\mathbb{p}_{p,2}'$ by changing the domain of $\mathbb{p}_{p,2}$ from $\mathbb{L}^b$ back to $\mathbb{L}$ as $\mathbb{p}_{p,2}' = \{\mathbf{mx} \leftarrow \mathbf{ff}\}$.

An $\mathbf{abst}_{\Pi}(P)$ run $(s_0^b, m_0^b)t_1^b \dots t_n^b(s_n^b, m_n^b)$ is a configuration starting sequence of alternating transitions and configurations. This run is considered *feasible* if we have that $(s_i^b, m_i^b) \xrightarrow[t_{i+1}^b]{\mathbf{abst}_{\Pi}(P)} (s_{i+1}^b, m_{i+1}^b)$ for each $i : 0 \leq i < n$ and s_0^b, m_0^b are in $\mathbb{S}_{init}^b$ and $\mathbb{M}_{init}^b$ respectively. Configurations (s_i^b, m_i^b) , for $i : 0 \leq i \leq n$, are then said to be reachable.

Evaluating (counting) predicates in $\mathbf{abst}_{\Pi}(P)$. Given a shared configuration s^b , we write $\mathbf{originOf}(s^b)$ to mean the predicate $\bigwedge_{s^b \in \mathbb{S}^b} (s^b \Leftrightarrow \mathbf{originOf}(s^b))$. We write $\mathbf{originOf}(\mathbb{L}^b)$ to mean $\bigwedge_{l^b \in \mathbb{L}^b} (l^b \Leftrightarrow \mathbf{originOf}(l^b))$. Observe that $\mathbf{vars}(\mathbf{originOf}(s^b)) \subseteq S$ and the variables $\mathbf{vars}(\mathbf{originOf}(\mathbb{L}^b)) \subseteq S \cup L$. We abuse notation and write $s^b[s]$ (resp. $\mathbb{L}^b[s, \mathbb{L}]$) to mean that $\mathbf{originOf}(s^b)[s]$ (resp. $\mathbf{originOf}(\mathbb{L}^b)[s, \mathbb{L}]$) holds.

Let π be a predicate in $\mathbf{predsOf}^{\text{@PC} \cup \Pi}$, where all boolean variables are either predicates in Π or of the form @pc for some $\mathbf{pc} \in PC$. We write $\pi[s^b, (\mathbf{pc}, \mathbb{L}^b)]$ to mean the boolean value obtained by evaluating the result of replacing each boolean variable $\text{@pc}'$ with $\mathbf{tt}$ if $\mathbf{pc} = \mathbf{pc}'$ and with $\mathbf{ff}$ otherwise, and by replacing each $\mathbf{p}'$ in Π with $\mathbb{L}^b(\mathbf{v}^b)$ or $s^b(\mathbf{v}^b)$ where $\mathbf{originOf}(\mathbf{v}^b) = \pi'$.

We can now evaluate a counting variable $(\pi)^\#$, with $\pi \in \mathbf{predsOf}^{\text{@PC} \cup \Pi}$, wrt. a configuration (s^b, m^b) . We do this by counting the number of process states $(\mathbf{pc}, \mathbb{L}^b)$ in m^b for which $\pi[s^b, (\mathbf{pc}, \mathbb{L}^b)]$ holds (see Eq. 7). We can also replace each counting variable $(\pi)^\#$ in a counting predicate ω with its value in (s^b, m^b) , and each shared predicate in Π with its value in s^b (see Eq. 8). Observe that the obtained expression might still involve shared variables as these can participate in comparisons with counting variables. Such comparisons do not correspond to any predicate in Π and can therefore not be mapped.

$$(\pi)^\#[s^b, m^b] = \sum_{\{(\mathbf{pc}, \mathbb{L}^b) \mid \pi[s^b, (\mathbf{pc}, \mathbb{L}^b)]\}} m^b((\mathbf{pc}, \mathbb{L}^b)) \quad (7)$$

$$\omega[s^b, m^b] = \omega[\{(\pi)^\# \leftarrow (\pi)^\#[s^b, m^b] \mid (\pi)^\# \in \mathbf{vars}(\omega)\}] \quad (8)$$

Relation between P and $\mathbf{abst}_{\Pi}(P)$. We let $m^b[s, m]$ mean that there is a bijection h from $m = [\sigma_1, \sigma_2, \dots, \sigma_n]$ to $m^b = [\sigma'_1, \sigma'_2, \dots, \sigma'_n]$ s.t. for each i in $[1, n]$, if $\sigma_i = (\mathbf{pc}, \mathbb{L})$ then $\sigma'_{h(i)} = (\mathbf{pc}, \mathbb{L}^b)$ and $\mathbb{L}^b[s, \mathbb{L}]$. The *concretization* of an $\mathbf{abst}_{\Pi}(P)$ configuration (s^b, m^b) is $\gamma((s^b, m^b)) = \{(s, m) \mid s^b[s] \wedge m^b[s, m]\}$. The *abstraction* of (s, m) is the singleton $\alpha((s, m)) = \{(s^b, m^b) \mid s^b[s] \wedge m^b[s, m]\}$. We initialize variables in $\mathbf{abst}_{\Pi}(P)$ such that for each pair (s_{init}, m_{init}) of P , there are (s_{init}^b, m_{init}^b) that are initial

in $\mathbf{abst}_{\Pi}(P)$ so that $\alpha((s_{init}, m_{init})) = \{(s_{init}^b, m_{init}^b)\}$. The abstraction $\alpha(\rho)$ of a P run $\rho = (s_0, m_0)t_1 \dots t_n(s_n, m_n)$ is the set of $\mathbf{abst}_{\Pi}(P)$ runs $\{(s_0^b, m_0^b)t_1^b \dots t_n^b(s_n^b, m_n^b)\}$ where $\alpha((s_i, m_i)) = \{(s_i^b, m_i^b)\}$ and $t_i^b = \mathbf{abst}_{\Pi}(t_i)$. Concretizations of abstract runs are defined in a straightforward manner.

Example 10. Consider the program in Fig. 1, its corresponding abstraction wrt. the set of predicates $\Pi = \{\mathbf{read_leq_0}, \mathbf{wait_leq_count}\}$ in Fig. 5, and the feasible run in Example 4. The initial shared state s_0 defined as $\{\mathbf{count} \leftarrow 1, \mathbf{wait} \leftarrow 0, \mathbf{cross} \leftarrow 0, \mathbf{read} \leftarrow 0\}$ in original program will be encoded as the boolean shared state $s_0^b = \{\mathbf{read_leq_0} \leftarrow \mathbf{tt}, \mathbf{wait_leq_count} \leftarrow \mathbf{tt}\}$ in the boolean program. We have that $\mathbf{originOf}(s_0^b) = (((\mathbf{read} \leq 0) = \mathbf{tt}) \wedge ((\mathbf{wait} \leq \mathbf{count}) = \mathbf{tt}))$. Since there are no local variables in this example, we get that $\alpha((s_0, m_0)) = \{(s_0^b, m_0)\}$. Many other states have the same encoding, e.g. $s_1 = \{\mathbf{count} \leftarrow 10, \mathbf{wait} \leftarrow 2, \mathbf{cross} \leftarrow 0, \mathbf{read} \leftarrow 0\}$ which satisfies $s_0^b[s_1]$, although it is not initial in $\mathbf{prog}$. We get that $\gamma((s_0^b, m_0)) = \{(s_0, m_0), (s_1, m_0), \dots\}$.

Definition 1 (predicate abstraction). Let the abstraction of the program $P = (S, L, T)$ wrt. Π be the boolean program $\mathbf{abst}_{\Pi}(P) = (S^b, L^b, T^b)$. The abstraction is said to be effective and sound if $\mathbf{abst}_{\Pi}(P)$ can be effectively computed and the abstract run in the singleton $\alpha(\rho)$ of any feasible P run ρ is $\mathbf{abst}_{\Pi}(P)$ feasible.

5.2 Encoding into a counter machine

Assume a program $P = (S, L, T)$, a set of predicates $\Pi_0 \subseteq \mathbf{predsOf}_{\text{exprsOf}(S \cup L)}$ and two counting predicates, an invariant ω_{inv} in the set $\mathbf{predsOf}_{\text{exprsOf}(S \cup \Omega_{PC, S, L})}$ and a target predicate ω_{trgt} in $\mathbf{predsOf}_{\text{exprsOf}(\Omega_{PC, S, L})}$. In the following, we write $\mathbf{abst}_{\Pi}(P) = (S^b, L^b, T^b)$ to mean the abstraction of P with respect to the predicates Π defined as

$$\Pi = \bigcup_{(\pi)^\# \in \mathbf{vars}(\omega_{inv}) \cup \mathbf{vars}(\omega_{trgt})} \mathbf{comparisonsOf}(\pi) \cup \Pi_0$$

Intuitively, this step results in the formulation of a state reachability problem of a counter machine $\mathbf{enc}(\mathbf{abst}_{\Pi}(P))$ that captures reachability of abstractions of ω_{trgt} configurations with $\mathbf{abst}_{\Pi}(P)$ runs that take into account the invariant ω_{inv} .

A tuple $M = (Q, C, \Delta, Q_{init}, C_{init}, q_{trgt})$ is a counter machine where Q is a finite set of states, C is a finite set of counters (i.e., variables ranging over the natural numbers $\mathbb{N}$), Δ is a finite set of transitions, $Q_{init} \subseteq Q$ is a set of initial states, C_{init} is a set of initial 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

$$\begin{array}{c}
\frac{(\mathbb{s}^b, \mathbb{l}^b, \mathbb{m}^b) \xrightarrow[\text{abst}_{\Pi(P)}]{\text{stmt}} (\mathbb{s}^{b'}, \mathbb{l}^{b'}, \mathbb{m}^{b'})}{(\mathbb{s}^b, (\mathbf{pc}, \mathbb{l}^b) \oplus \mathbb{m}^b) \xrightarrow[\text{abst}_{\Pi(P)}]{(\mathbf{pc} \rightarrow \mathbf{pc}'; \text{stmt})} (\mathbb{s}^{b'}, (\mathbf{pc}', \mathbb{l}^{b'}) \oplus \mathbb{m}^{b'})} : \text{transition} \quad \frac{\pi[\mathbb{s}^b, \mathbb{l}^b] \text{ is true}}{(\mathbb{s}^b, \mathbb{l}^b, \mathbb{m}^b) \xrightarrow[\text{abst}_{\Pi(P)}]{\pi} (\mathbb{s}^b, \mathbb{l}^b, \mathbb{m}^b)} : \text{assume} \\
\\
\frac{(\mathbb{s}^b, \mathbb{l}^b, \mathbb{m}^b) \xrightarrow[\text{abst}_{\Pi(P)}]{\text{stmt}} (\mathbb{s}^{b'}, \mathbb{l}^{b'}, \mathbb{m}^{b'}) \text{ and } (\mathbb{s}^{b'}, \mathbb{l}^{b'}, \mathbb{m}^{b'}) \xrightarrow[\text{abst}_{\Pi(P)}]{\text{stmt}'} (\mathbb{s}^{b''}, \mathbb{l}^{b''}, \mathbb{m}^{b''})}{(\mathbb{s}^b, \mathbb{l}^b, \mathbb{m}^b) \xrightarrow[\text{abst}_{\Pi(P)}]{\text{stmt}; \text{stmt}'} (\mathbb{s}^{b''}, \mathbb{l}^{b''}, \mathbb{m}^{b''})} : \text{sequence} \\
\\
\frac{\mathbb{m}^{b'} = ((\mathbf{pc}_0, \mathbb{l}_{init}^b) \oplus \mathbb{m}^b) \text{ with } \mathbb{l}_{init}^b \in \mathbb{L}_{init}^b}{(\mathbb{s}^b, \mathbb{l}^b, \mathbb{m}^b) \xrightarrow[\text{abst}_{\Pi(P)}]{\text{spawn}} (\mathbb{s}^b, \mathbb{l}^b, \mathbb{m}^{b'})} : \text{spawn} \quad \frac{\mathbb{m}^b = ((\mathbf{pc}_x, \mathbb{l}^{b'}) \oplus \mathbb{m}^{b'})}{(\mathbb{s}^b, \mathbb{l}^b, \mathbb{m}^b) \xrightarrow[\text{abst}_{\Pi(P)}]{\text{join}} (\mathbb{s}^b, \mathbb{l}^b, \mathbb{m}^{b'})} : \text{join} \\
\\
\frac{\begin{array}{l} \text{a bijection } h \text{ from } \mathbb{m}^b = [\sigma_1, \sigma_2, \dots, \sigma_n] \text{ to } \mathbb{m}^{b'} = [\sigma'_1, \sigma'_2, \dots, \sigma'_n] \text{ s.t.} \\ \forall i \in [1, n], \text{ if } \sigma_i = (\mathbf{pc}_p, \mathbb{l}_p^b) \text{ and } \sigma'_{h(i)} = (\mathbf{pc}'_p, \mathbb{l}_p^{b'}) \text{ then } \mathbf{pc}_p = \mathbf{pc}'_p \\ \text{and } (\mathbb{s}^b, \mathbb{l}^b, \mathbb{l}_p^b) \xrightarrow[\text{abst}_{\Pi(P)}]{v_1^b, \dots, v_n^b := \text{ch}(\pi_1, \pi'_1), \dots, \text{ch}(\pi_n, \pi'_n)} (\mathbb{s}^{b'}, \mathbb{l}^{b'}, \mathbb{l}_p^{b'}) \end{array}}{(\mathbb{s}^b, \mathbb{l}^b, \mathbb{m}^b) \xrightarrow[\text{abst}_{\Pi(P)}]{v_1^b, \dots, v_n^b := \text{ch}(\pi_1, \pi'_1), \dots, \text{ch}(\pi_n, \pi'_n)} (\mathbb{s}^{b'}, \mathbb{l}^{b'}, \mathbb{m}^{b'})} : \text{assign}
\end{array}$$

Fig. 7. Semantics of boolean concurrent programs. Executions start from some $\mathbb{m}_{init}^b$ in $\mathbb{M}_{init}^b$.

$$\begin{array}{c}
\frac{\delta = (q : \text{op} : q') \text{ and } \mathbb{c} \xrightarrow[M]{\text{op}} \mathbb{c}'}{(q, \mathbb{c}) \xrightarrow[M]{\delta} (q', \mathbb{c}')} : \text{transition} \\
\\
\frac{}{\mathbb{c} \xrightarrow[M]{\text{nop}} \mathbb{c}} : \text{nop} \quad \frac{\mathbb{c} \xrightarrow[M]{\text{op}} \mathbb{c}' \text{ and } \mathbb{c}' \xrightarrow[M]{\text{op}'} \mathbb{c}''}{\mathbb{c} \xrightarrow[M]{\text{op}; \text{op}'} \mathbb{c}''} : \text{sequence} \\
\\
\frac{\exists A. \text{grad}[\mathbb{c}] \wedge \mathbb{c}' = \mathbb{c}[\{c_i \leftarrow \mathbf{e}_i[\mathbb{c}] \mid i : 1 \leq i \leq n\}]}{\mathbb{c} \xrightarrow[M]{\text{grad} \Rightarrow (c_1 \dots c_n := \mathbf{e}_1 \dots \mathbf{e}_n)} \mathbb{c}'} : \text{command}
\end{array}$$

Fig. 8. Semantics of a counter machine

in the guarded command rule in Fig. 8. A guard grad is a predicate in $\mathbf{predsOf}_{\mathbf{exprsOf}(A \cup C)}$ and a command cmd is a multiple assignment $c_1, \dots, c_n := \mathbf{e}_1, \dots, \mathbf{e}_n$ that involves $\mathbf{e}_1, \dots, \mathbf{e}_n$ in $\mathbf{exprsOf}(A \cup C)$ and pairwise different $c_1, \dots, c_n$ in C .

A *machine configuration* is a pair $(q, \mathbb{c})$ where q is a state in Q and $\mathbb{c}$ is a mapping $C \rightarrow \mathbb{N}$. Semantics are given in Fig. 8. A configuration $(q, \mathbb{c})$ is *initial* if $q \in Q_{init}$ and $\mathbb{c} \in C_{init}$. An M run ρ_M is a sequence $(q_0, \mathbb{c}_0) \delta_1 \dots (q_n, \mathbb{c}_n)$. It is *feasible* if $(q_0, \mathbb{c}_0)$ is initial and $(q_i, \mathbb{c}_i) \xrightarrow[M]{\delta_{i+1}} (q_{i+1}, \mathbb{c}_{i+1})$ for $i : 0 \leq i < n$. The machine state reachability problem is to decide whether there is an M feasible run $(q_0, \mathbb{c}_0) \delta_1 \dots (q_n, \mathbb{c}_n)$ s.t. $q_n = q_{trgt}$.

Encoding. We describe in the following a counter machine $\text{enc}(\text{abst}_{\Pi}(P))$ obtained as an encoding of the boolean program $\text{abst}_{\Pi}(P)$. Recall $\text{abst}_{\Pi}(P)$ results from the predicate abstraction of the concurrent program P wrt. some initial predicates Π_0 as well as all predicates $\mathbf{comparisonsOf}(\pi)$ for every counting variable $(\pi)^\#$ in $\text{vars}(\omega_{inv})$ and $\text{vars}(\omega_{trgt})$. The machine $\text{enc}(\text{abst}_{\Pi}(P))$ is a tuple $(Q, C, \Delta, Q_{init}, C_{init}, q_{trgt})$. A state in Q is either the target state q_{trgt} or is associated to a shared configuration $\mathbb{s}^b$ of $\text{abst}_{\Pi}(P)$. We write $q_{\mathbb{s}^b}$ to make the association explicit. There is a bijection that associates a process configuration $(\mathbf{pc}, \mathbb{l}^b)$ to each counter $c_{(\mathbf{pc}, \mathbb{l}^b)}$ in C . Only the assign rule makes use of auxiliary variables. In case of broadcasts this rule associates an auxiliary variable $a_{(\mathbf{pc}, \mathbb{l}_p^b, \mathbb{l}_p^{b'})}$ to each possible move from process configuration $(\mathbf{pc}, \mathbb{l}_p^b)$ to $(\mathbf{pc}, \mathbb{l}_p^{b'})$.

We will write $\sum_{\{(\mathbf{pc}, \mathbb{l}^b) \mid \pi[\mathbb{s}^b, (\mathbf{pc}, \mathbb{l}^b)]\}} c_{(\mathbf{pc}, \mathbb{l}^b)}$ to mean the sum of all counters $c_{(\mathbf{pc}, \mathbb{l}^b)}$ in C such that $\pi[\mathbb{s}^b, (\mathbf{pc}, \mathbb{l}^b)]$ evaluates to **tt**. We can define the mapping $\mathbb{w}_{\mathbb{s}^b, \omega}$ that maps each counting variable appearing in a counting predicate ω to a corresponding counters sum under $\mathbb{s}^b$. More formally, $\mathbb{w}_{\mathbb{s}^b, \omega}((\pi)^\#) = \sum_{\{(\mathbf{pc}, \mathbb{l}^b) \mid \pi[\mathbb{s}^b, (\mathbf{pc}, \mathbb{l}^b)]\}} c_{(\mathbf{pc}, \mathbb{l}^b)}$ for each counting predicate $(\pi)^\#$ in $\text{vars}(\omega)$. As a result, $\omega_{trgt}[\mathbb{w}_{\mathbb{s}^b, \omega}]$ is the predicate obtained from ω_{trgt} after replacing all counting variables appearing in it by the corresponding counters sums. Observe that if the predicate ω is in $\mathbf{predsOf}_{\mathbf{exprsOf}(\Omega_{PC, S, L})}$ as it is the case for ω_{trgt} , then $\omega[\mathbb{w}_{\mathbb{s}^b, \omega}]$ does not mention any shared variables in S . The target predicate ω_{trgt} does not mention any shared variables, because it is to be evaluated in a configura-

$$\begin{array}{c}
\frac{t = (\mathbf{pc} \rightarrow \mathbf{pc}' : \mathbf{stmt}_\Pi) \text{ and } \left[(s^b, \mathbb{l}^b) : op : (s^{b'}, \mathbb{l}^{b'}) \right]_{\mathbf{stmt}_\Pi}}{(q_{s^b} : c_{(\mathbf{pc}, \mathbb{l}^b)} \geq 1; c_{(\mathbf{pc}, \mathbb{l}^b)}^{--}; op; c_{(\mathbf{pc}', \mathbb{l}^{b'})}^{++} : q_{s^{b'}}) \in \Delta_t} : \text{transition} \\
\\
\frac{}{(q_{s^b} : \omega_{tgt}[\mathbb{w}_{s^b, \omega_{tgt}}] : q_{tgt}) \in \Delta_{tgt}} : \text{target} \\
\\
\frac{\left[(s^b, \mathbb{l}^b) : op : (s^{b'}, \mathbb{l}^{b'}) \right]_{\mathbf{stmt}} \text{ and } \left[(s^{b'}, \mathbb{l}^{b'}) : op' : (s^{b''}, \mathbb{l}^{b''}) \right]_{\mathbf{stmt}'}}{\left[(s^b, \mathbb{l}^b) : op; op' : (s^{b''}, \mathbb{l}^{b''}) \right]_{\mathbf{stmt}; \mathbf{stmt}'}} : \text{sequence} \\
\\
\frac{\pi[s^b, \mathbb{l}^b]}{\left[(s^b, \mathbb{l}^b) : nop : (s^b, \mathbb{l}^b) \right]_\pi} : \text{assume} \quad \frac{}{\left[(s^b, \mathbb{l}^b) : c_{(\mathbf{pc}_0, \mathbb{l}^{b_{init}})}^{++} : (s^b, \mathbb{l}^b) \right]_{\text{spawn}}} : \text{spawn} \\
\\
\frac{}{\left[(s^b, \mathbb{l}^b) : c_{(\mathbf{pc}_x, \mathbb{l}^{b'})} \geq 1; c_{(\mathbf{pc}_x, \mathbb{l}^{b'})}^{--} : (s^b, \mathbb{l}^b) \right]_{\text{join}}} : \text{join} \\
\\
\begin{array}{l}
s^{b'} = s^b[\{v_i^b \leftarrow \pi_i[s^b, (-, \mathbb{l}^b)] | i : 1 \leq i \leq n\}] \quad \mathbb{l}^{b'} = \mathbb{l}^b[\{v_i^b \leftarrow \pi_i[s^b, (-, \mathbb{l}^b)] | i : 1 \leq i \leq n\}] \\
\mathbf{tf} = \left\{ (\mathbb{l}_p^b, \mathbb{l}_p^{b'}) | (s^b, \mathbb{l}^b, \mathbb{l}_p^b) \xrightarrow{v_1^b, \dots, v_n^b := \text{ch}(\pi_1, \pi'_1), \dots, \text{ch}(\pi_n, \pi'_n)} (s^{b'}, \mathbb{l}^{b'}, \mathbb{l}_p^{b'}) \right\} \\
\qquad \qquad \qquad \text{abst}_\Pi(P)
\end{array} \\
\\
\frac{}{\left[(s^b, \mathbb{l}^b) : \left(\bigwedge_{\mathbf{pc} \in PC, \mathbb{l}_p^b \in \mathbb{L}^b} \left(c_{(\mathbf{pc}, \mathbb{l}_p^b)} = \sum_{\{(\mathbf{pc}, \mathbb{l}_p^{b'}) | (\mathbb{l}_p^b, \mathbb{l}_p^{b'}) \in \mathbf{tf}\}} a_{(\mathbf{pc}, \mathbb{l}_p^b, \mathbb{l}_p^{b'})} \right) \Rightarrow \right.} : \text{assign} \\
\left. \left. \text{cmdOf} \left(\left\{ c_{(\mathbf{pc}, \mathbb{l}_p^b)} \leftarrow \sum_{\{(\mathbf{pc}, \mathbb{l}_p^{b'}) | (\mathbb{l}_p^b, \mathbb{l}_p^{b'}) \in \mathbf{tf}\}} a_{(\mathbf{pc}, \mathbb{l}_p^b, \mathbb{l}_p^{b'})} | \mathbf{pc} \in PC, \mathbb{l}_p^b \in \mathbb{L}^b \right\} \right) : (s^{b'}, \mathbb{l}^{b'}) \right]_{\text{assign}}}
\end{array}$$

Fig. 9. Encoding of the transitions of a boolean program (S^b, L^b, T^b) , given a counting target ω_{tgt} , to the transitions $\Delta = \cup_{t \in T^b} \Delta_t \cup \Delta_{tgt}$ of a counter machine. In rule *assign*, we write $\text{cmdOf} \left(\left\{ c_{(\mathbf{pc}, \mathbb{l}^b)} \leftarrow e_{(\mathbf{pc}, \mathbb{l}^b)} | \mathbf{pc} \in PC, \mathbb{l}^b \in \mathbb{L}^b \right\} \right)$ to mean the multiple assignment that simultaneously assigns each $e_{(\mathbf{pc}, \mathbb{l}^b)}$ to $c_{(\mathbf{pc}, \mathbb{l}^b)}$. A transition $c_1 \xrightarrow[\text{abst}_\Pi(P)]{\text{grd} \Rightarrow \text{cmd}} c_2$ ensures that there is a mapping $\mathfrak{a} : A \rightarrow \mathbb{N}$ s.t. for any $c_{(\mathbf{pc}, \mathbb{l}_p^b)}$ we have $c_1(c_{(\mathbf{pc}, \mathbb{l}_p^b)}) = \sum_{(\mathbb{l}_p^b, \mathbb{l}_p^{b'}) \in \mathbf{tf}} \mathfrak{a}(a_{(\mathbf{pc}, \mathbb{l}_p^b, \mathbb{l}_p^{b'})})$ and for any $c_{(\mathbf{pc}, \mathbb{l}_p^{b'})}$ we have $c_2(c_{(\mathbf{pc}, \mathbb{l}_p^{b'})}) = \sum_{(\mathbb{l}_p^b, \mathbb{l}_p^{b'}) \in \mathbf{tf}} \mathfrak{a}(a_{(\mathbf{pc}, \mathbb{l}_p^b, \mathbb{l}_p^{b'})})$.

tion of the counter machine where no concrete value for shared variables is available. The predicate $\omega_{tgt}[\mathbb{w}_{s^b}, \omega]$ in $\text{predsOf}_{\text{exprsOf}(C)}$ is used in the target rule of the encoding in Fig. 9.

The set of transitions Δ is exactly the set $\cup_{t \in T^b} \Delta_t \cup \Delta_{tgt}$ as described in Fig. 9. We abuse notation and associate to each statement $\mathbf{stmt}$ appearing in $\text{abst}_\Pi(P)$ the set $\text{enc}(\mathbf{stmt})$ of tuples $\left[(s^b, \mathbb{l}^b) : op : (s^{b'}, \mathbb{l}^{b'}) \right]_{\mathbf{stmt}}$ generated in Fig. 9. Given a processes configuration $\mathfrak{m}^b$, we write $\mathbb{C}_{\mathfrak{m}^b}$ to mean the mapping associating $\mathfrak{m}^b((\mathbf{pc}, \mathbb{l}^b))$ to each counter $c_{(\mathbf{pc}, \mathbb{l}^b)}$ in C . We let Q_{init} be the set $\{q_{s_0^b} | s_0^b \in S_{init}^b \text{ of } \text{abst}_\Pi(P)\}$, and $\mathbb{C}_{init}$ be the mapping $\{\mathbb{C}_{\mathfrak{m}^b} | \mathfrak{m}^b((\mathbf{pc}_0, \mathbb{l}^b)) = 1 \text{ for a single } \mathbb{l}^b \in \mathbb{L}_{init}^b \text{ in } \text{abst}_\Pi(P) \text{ and } 0 \text{ otherwise}\}$. We associate a program configuration $(s^b, \mathfrak{m}^b)$ to each machine configuration $(q_{s^b}, \mathbb{C}_{\mathfrak{m}^b})$. The machine $\text{enc}(\text{abst}_\Pi(P))$ encodes $\text{abst}_\Pi(P)$ as specified in Lem. 3.

We state in Lem. 3 that the reachability problem of the obtained counter machine is equivalent to the reachability in $\text{abst}_\Pi(P)$ of boolean configurations satisfying ω_{tgt} . For this, we make use of Lem. 1 and Lem. 2. Intuitively, these relate executions of the boolean abstraction to the ones of the counter machine encoding.

Lemma 1 (translation). *For any $\mathbf{stmt}$ appearing in the program $\text{abst}_\Pi(P)$, $(s^b, \mathbb{l}^b, \mathfrak{m}^b) \xrightarrow[\text{abst}_\Pi(P)]{\mathbf{stmt}} (s^{b'}, \mathbb{l}^{b'}, \mathfrak{m}^{b'})$ iff $\mathbb{C}_{\mathfrak{m}^b} \xrightarrow[\text{enc}(\text{abst}_\Pi(P))]{op} \mathbb{C}_{\mathfrak{m}^{b'}}$ for a $\left[(s^b, \mathbb{l}^b) : op : (s^{b'}, \mathbb{l}^{b'}) \right]_{\mathbf{stmt}}$ in $\text{enc}(\mathbf{stmt})$.*

Proof. By induction on the number of atomic statements in $\mathbf{stmt}$. $\square$

Lemma 2 (translation and abstraction). *Any configuration $(s^b, \mathfrak{m}^b)$ is reachable in $\text{abst}_\Pi(P)$ iff $(q_{s^b}, \mathbb{C}_{\mathfrak{m}^b})$ is reachable in $\text{enc}(\text{abst}_\Pi(P))$.*

Proof. We show that $(s^b, \mathfrak{m}^b)$ is reachable via a run of length n in $\text{abst}_\Pi(P)$ iff $(q_{s^b}, \mathbb{C}_{\mathfrak{m}^b})$ is reachable via a run of the same length in $\text{enc}(\text{abst}_\Pi(P))$. We proceed by induction on the number of P transitions appearing in the runs. By construction, s_0^b and $\mathfrak{m}_0^b$ are initial iff $q_{s_0^b}$ and $\mathbb{C}_{\mathfrak{m}_0^b}$ are also initial. Let $(s^b, \mathfrak{m}^b)$ be a reachable $\text{abst}_\Pi(P)$ configuration and $(q_{s^b}, \mathbb{C}_{\mathfrak{m}^b})$ be the corresponding $\text{enc}(\text{abst}_\Pi(P))$ configuration. We will not consider runs in $\text{enc}(\text{abst}_\Pi(P))$ that involve Δ_{tgt} transitions as these lead to error states and not to configurations of the form $(q_{s^b}, \mathbb{C}_{\mathfrak{m}^b})$. We show $(s^b, \mathfrak{m}^b) \xrightarrow[\text{abst}_\Pi(P)]{\mathbf{pc} \rightarrow \mathbf{pc}' : \mathbf{stmt}}$

$$\begin{aligned}
(s^{b'}, m^{b'}) \text{ iff } (q_{s^b}, c_{m^b}) &\xrightarrow[\text{enc}(\text{abst}_\Pi(P))]{(q_{s^b} : c_{(pc, \mathbb{I}^b)} \geq 1 ; c_{(pc, \mathbb{I}^b)}^{--}; op; c_{(pc', \mathbb{I}^{b'})}^{++}; q_{s^{b'}})} \\
(q_{s^{b'}}, c_{m^{b'}}) \text{ for some } &\left[(s^b, \mathbb{I}^b) : op : (s^{b'}, \mathbb{I}^{b'}) \right]_{\text{stmt}} \text{ in } \text{enc}(\text{stmt}). \\
\text{Semantics of boolean programs in Fig. 7 ensure that} & \\
(s^b, m^b) \xrightarrow[\text{abst}_\Pi(P)]{(pc \rightarrow pc' : \text{stmt})} (s^{b'}, m^{b'}) &\text{ iff } m^b = (pc, \mathbb{I}^b) \oplus m_1^b \text{ and} \\
m^{b'} = (pc', \mathbb{I}^{b'}) \oplus m_1^{b'} \text{ and } (s^b, \mathbb{I}^b, m_1^b) &\xrightarrow[\text{abst}_\Pi(P)]{\text{stmt}} (s^{b'}, \mathbb{I}^{b'}, m_1^{b'}).
\end{aligned}$$

Lem. 1 ensures this is equivalent to $c_{m_1^b} \xrightarrow[\text{enc}(\text{abst}_\Pi(P))]{op} c_{m_1^{b'}}$

for some $\left[(s^b, \mathbb{I}^b) : op : (s^{b'}, \mathbb{I}^{b'}) \right]_{\text{stmt}}$ in $\text{enc}(\text{stmt})$.

Observe that $c_{m_1^b} \xrightarrow[\text{enc}(\text{abst}_\Pi(P))]{op} c_{m_1^{b'}}$ is equivalent to

$$c_{(pc, \mathbb{I}^b) \oplus m_1^b} \xrightarrow[\text{enc}(\text{abst}_\Pi(P))]{c_{(pc, \mathbb{I}^b)} \geq 1; c_{(pc, \mathbb{I}^b)}^{--}; op; c_{(pc', \mathbb{I}^{b'})}^{++}} c_{(pc', \mathbb{I}^{b'}) \oplus m_1^{b'}}. \quad \square$$

Lemma 3 (translation reachability). *Target q_{trgt} is $\text{enc}(\text{abst}_\Pi(P))$ reachable iff a configuration (s^b, m^b) is reachable in $\text{abst}_\Pi(P)$ such that $\omega_{\text{trgt}}[s^b, m^b]$ holds.*

Proof. Lem. 2 ensures that (s^b, m^b) is $\text{abst}_\Pi(P)$ reachable iff (q_{s^b}, c_{m^b}) is $\text{enc}(\text{abst}_\Pi(P))$ reachable. We conclude by observing that $\omega_{\text{trgt}}[s^b, m^b]$ holds iff the evaluation of the target predicate in the counter machine holds, i.e. $\omega_{\text{trgt}}[q_{s^b}, c_{m^b}]$ holds. $\square$

5.3 Encoding precision

We argue in the following that the obtained counter machine often results in a monotonic transition system for which the reachability problem is decidable. In fact, predicate abstraction forces monotonicity. For example, in Fig. 1, transitions t_4 and t_5 correspond to a barrier which is non-monotonic. But, the abstracted boolean program in Fig. 5 that corresponds to it, consists of only monotonic transitions. This happens because the relation between the number of processes in different program locations and the program variables is lost. This corresponds to a loss of precision that makes it impossible to establish correctness of programs such as the one depicted in Fig. 1. We explain how we do retrieve some of that precision by strengthening the abstraction.

Consider the boolean program $\text{abst}_\Pi(P)$ obtained after predicate abstraction. If a configuration $(s^{b'}, m^{b'})$ is obtained from (s^b, m^b) using some transition, then the same transition can obtain a larger configuration (i.e., has the same shared state $s^{b'}$ and more processes at the same process states in $m^{b'}$) than $(s^{b'}, m^{b'})$ from any configuration larger than (s^b, m^b) . Lem. 4 shows that indeed, all transitions in Fig. 9 (except for rule *target*) are monotonic with respect to the ordering $\sqsubseteq$ defined by $(q, c) \sqsubseteq (q', c')$ iff $q = q'$ and $c \preceq c'$ ¹.

Lemma 4 (monotonicity). *Transitions $(q : op : q')$ generated by all rules in Fig. 9, except for the target rule, are monotonic wrt. $\sqsubseteq$.*

Proof. Let op be some operation appearing in a generated transition $(q : op : q')$ of $\text{enc}(\text{abst}_\Pi(P))$. We say that an operation op is monotonic wrt. $\preceq$ if for each c_1, c_2, c_3 s.t. $c_1 \xrightarrow[\text{enc}(\text{abst}_\Pi(P))]{op} c_2$ and $c_1 \preceq c_3$ there exists an c_4 s.t. $c_3 \xrightarrow[\text{enc}(\text{abst}_\Pi(P))]{op} c_4$ and $c_2 \preceq c_4$. Observe that $(q : op : q')$ is monotonic wrt. $\sqsubseteq$ iff op is monotonic wrt. $\preceq$. In addition, observe that if both op and op' are monotonic, then so is $op; op'$. It is therefore enough to show monotonicity of nop , $c \geq 1$, c^{++} , c^{--} and $grd \Rightarrow cmd$. The first four cases are straightforward. We show $grd \Rightarrow cmd$ is monotonic. Suppose we are given

c_1, c_2, c_3 s.t. $c_1 \xrightarrow[\text{enc}(\text{abst}_\Pi(P))]{grd \Rightarrow cmd} c_2$ and $c_1 \preceq c_3$. We exhibit

a c_4 s.t. $c_3 \xrightarrow[\text{enc}(\text{abst}_\Pi(P))]{grd \Rightarrow cmd} c_4$ and $c_2 \preceq c_4$. The operation

$grd \Rightarrow cmd$, resulted from the assign rule in Fig. 9. This was defined wrt. to two pairs $(s^b, \mathbb{I}^b)$ and $(s^{b'}, \mathbb{I}^{b'})$. We fix these two pairs. By the semantics of counter machines (Fig. 8) and of the translation of the *assign* statements in Fig. 9, the fact that $c_1 \xrightarrow[\text{enc}(\text{abst}_\Pi(P))]{grd \Rightarrow cmd} c_2$, ensures that there is a mapping $\sigma : A \rightarrow \mathbb{N}$ s.t. for any $c_{(pc, \mathbb{I}^b)}$ we have $c_1(c_{(pc, \mathbb{I}^b)}) = \sum_{(\mathbb{I}_p^b, \mathbb{I}_p^{b'}) \in \text{tf}} \sigma(a_{(pc, \mathbb{I}^b, \mathbb{I}_p^b, \mathbb{I}_p^{b'})})$ and for any $c_{(pc, \mathbb{I}_p^{b'})}$ we have $c_2(c_{(pc, \mathbb{I}_p^{b'})}) = \sum_{(\mathbb{I}_p^b, \mathbb{I}_p^{b'}) \in \text{tf}} \sigma(a_{(pc, \mathbb{I}_p^b, \mathbb{I}_p^{b'})})$. Since $c_1 \preceq c_3$, then for all $c_{(pc, \mathbb{I}^b)} \in C$, we have that $c_3(c_{(pc, \mathbb{I}^b)}) = c_1(c_{(pc, \mathbb{I}^b)}) + r_{(pc, \mathbb{I}^b)} = \sum_{(\mathbb{I}_p^b, \mathbb{I}_p^{b'}) \in \text{tf}} \sigma(a_{(pc, \mathbb{I}_p^b, \mathbb{I}_p^{b'})}) + r_{(pc, \mathbb{I}^b)}$ where $r_{(pc, \mathbb{I}^b)} \geq 0$. The idea is to send these “excedents” along the enabled transfers. Fix such a $(pc, \mathbb{I}_p^b)$.

By definition of $(s^b, \mathbb{I}^b, \mathbb{I}_p^b) \xrightarrow[\text{abst}_\Pi(P)]{v_1^b, \dots, v_n^b := \text{ch}(\pi_1, \pi_1'), \dots, \text{ch}(\pi_n, \pi_n')}$

$(s^{b'}, \mathbb{I}^{b'}, \mathbb{I}_p^{b'})$ in Sect. 5.1 and of **tf** in Fig. 9, we know there is at least a $\mathbb{I}_p^{b'}$ such that $(\mathbb{I}_p^b, \mathbb{I}_p^{b'}) \in \text{tf}$. We define $c_4(c_{(pc, \mathbb{I}^b)}) := \sum_{(\mathbb{I}_p^b, \mathbb{I}_p^{b'}) \in \text{tf}} \sigma'(a_{(pc, \mathbb{I}_p^b, \mathbb{I}_p^{b'})})$ where we have $\sigma'(a_{(pc, \mathbb{I}_p^b, \mathbb{I}_p^{b'})}) := \sigma(a_{(pc, \mathbb{I}_p^b, \mathbb{I}_p^{b'})}) + r_{(pc, \mathbb{I}_p^b)} \delta_{\mathbb{I}_p^b, \mathbb{I}_p^{b'}}$ with $\delta_{\mathbb{I}_1^b, \mathbb{I}_2^b}$ iff $\mathbb{I}_1^b$ is identical to $\mathbb{I}_2^b$. So, $c_2(c_{(pc, \mathbb{I}^b)}) \preceq c_4(c_{(pc, \mathbb{I}^b)})$. We repeat the process for each counter $c_{(pc, \mathbb{I}_p^b)}$ in C . This results in a c_4 where $c_2 \preceq c_4$ and for which the same transition (i.e. assign for the two pairs $(s^b, \mathbb{I}^b)$ and $(s^{b'}, \mathbb{I}^{b'})$) is possible using the mapping σ' . $\square$

In fact, even rule *target* results in monotonic machine transitions for all counting predicates ω_{trgt} that denote upward closed sets of processes (since the intersection of two upward closed sets is also upward closed). This is for instance the case of predicates capturing assertion violations but not of those capturing deadlocks (see Sec. 4). An encoding $\text{enc}(\text{abst}_\Pi(P))$ is said to be monotonic if all its transitions are monotonic. Checking assertion violations on abstractions obtained as in Sec. 5.1 always results in monotonic encodings.

¹ $c \preceq c'$ iff $c(c) \leq c'(c)$ for each $c \in C$.

Lemma 5 (decidability). *State reachability of all monotonic encodings is decidable.*

Proof. The ordering $\sqsubseteq$ is a well quasi ordering on the set of configuration of $enc(\mathbf{abst}_\Pi(P))$ [10]. Monotonicity of the transition rules means the obtained counter machines result in well structured transition systems [3, 14]. It is well known that state reachability is decidable for such systems. $\square$

Strengthening. Monotonic encodings correspond to coarse over-approximations. Intuitively, bad configurations (such as those where a deadlock occurs, or those obtained in a backward exploration for a barrier based program such as the one in Fig. 1) are no more guaranteed to be upward closed. This loss of precision makes verification out of the reach of techniques solely based on monotonic encodings. To regain some of the lost precision, we constrain the runs using counting invariants. This is done by strengthening the counter machine transitions in order to only allow configurations allowed by ω_{inv} .

Consider again the program of Fig. 1 with shared state $s = \{\text{count} \leftarrow 4, \text{wait} \leftarrow 0, \text{cross} \leftarrow 0, \text{read} \leftarrow 0\}$ and process configuration $m = \{\text{pc}_3 \leftarrow 1, \text{pc}_1 \leftarrow 3\}$. After t_4 , $s' = \{\text{count} \leftarrow 4, \text{wait} \leftarrow 1, \text{cross} \leftarrow 0, \text{read} \leftarrow 0\}$ and $m' = \{\text{pc}_4 \leftarrow 1, \text{pc}_1 \leftarrow 3\}$. Note that executing transition t_5 (barrier) at (s', m') is impossible because of the barrier condition. Abstraction of the program wrt. the predicates $\pi_1 = \text{read.leq.0}$, $\pi_2 = \text{wait.leq.count}$, $\pi_3 = \text{count.leq.wait}$ and $\pi_4 = \text{cross.leq.0}$ yields abstract transitions t_4^b and t_5^b . It turns out that starting from the abstract configuration (s^b, m^b) the two transitions can be executed as follows:

s^b				m^b
π_1	π_2	π_3	π_4	pc
true	true	false	true	pc_2, pc_3
$\Downarrow t_4^b.\text{wait.leq.count} :=$ $\text{ch}(\text{wait.leq.count} \wedge \neg \text{count.leq.wait}, \text{count.leq.wait})$ $\text{count.leq.wait} := \text{ch}(\text{count.leq.wait}, \text{ff})$				
true	true	true	true	pc_2, pc_4
$\Downarrow t_5^b.\text{wait.leq.count} \wedge \text{count.leq.wait}; \text{cross.leq.0} := \text{ff}$				
true	true	true	false	pc_2, pc_5

The steps above are not feasible in the original program, because t_4^b changed predicate count.leq.wait to tt although $\text{count} \leq \text{wait}$ does not hold in the original program. This results in a false positive.

Our solution to recover some of the lost precision is to strengthen the transitions of the counter machine using counting predicates (i.e. predicates that relate shared variables to the number of processes) that are valid in all runs, for example:

$$\omega_{inv} : \text{count} = \sum_{i \geq 1} (@\text{pc}_i)^\# \wedge \text{wait} = \sum_{i \geq 4} (@\text{pc}_i)^\#.$$

Note that the counting invariant ω_{inv} as opposed to the target predicates ω_{trgt} (Sec. 5.2) contains shared variables. This is to regain the information about the shared variables that was lost due to abstraction. Given an abstraction of a shared state s^b , we first conjunct ω_{inv} with $\text{originOf}(s^b)$ and project away the shared

$$\frac{(q_{s^b} : op : q_{s^b'})}{\left(q_{s^b} : str(s^b, \omega_{inv}) ; op ; str(s^b, \omega_{inv}) : q_{s^b'} \right)}$$

$$\frac{\left(q_{s^b} : \omega_{trgt}[\omega_{s^b, \omega_{trgt}}] : q_{trgt} \right)}{\left(q_{s^b} : str(s^b, \omega_{inv}) ; \omega_{trgt}[\omega_{s^b, \omega_{trgt}}] : q_{trgt} \right)}$$

Fig. 10. Strengthening of a counter machine transition given a counting invariant ω_{inv} . For this, let $str(s^b, \omega_{inv}) = (\exists S. (\text{originOf}(s^b) \wedge \omega_{inv}))[\omega_{s^b, \omega_{inv}}]$.

variables. We get a predicate that only contains counting variables. Then, we substitute the counting variables with their corresponding counters to obtain a predicate that only mentions counter machine variables. We finally strengthen all counter machine transition that involve the state q_{s^b} .

Example 11 (Strengthening). Consider the boolean shared state s^b with $\text{originOf}(s^b) = ((\text{read} \leq 0) \wedge (\text{wait} \leq \text{count}) \wedge (\text{count} \leq \text{wait}) \wedge \neg(\text{cross} \leq 0))$. After conjunction with the invariant and projection of the shared variables, we obtain: $\exists S. (\text{originOf}(s^b) \wedge \omega_{inv}) = ((@\text{pc}_1)^\# + (@\text{pc}_2)^\# + (@\text{pc}_3)^\#) = 0$. After substitution, we get $str(s^b, \omega_{inv}) = (c_{(\text{pc}_1, \emptyset)} + c_{(\text{pc}_2, \emptyset)} + c_{(\text{pc}_3, \emptyset)} = 0)$. Each counter $c_{(\text{pc}_i, \emptyset)}$ is the counter for processes at location pc_i where the processes have no local state. In other words, in the corresponding boolean program, there exists no local or mixed predicates. We add the condition $(c_{(\text{pc}_1, \emptyset)} + c_{(\text{pc}_2, \emptyset)} + c_{(\text{pc}_3, \emptyset)} = 0)$ as a statement at the beginning (resp. at the end) of each transitions outgoing from (resp. incoming to) q_{s^b} .

Strengthening is described in Fig. 10. We need Lem. 6 and Lem. 7 in order to establish soundness of the strengthening step in Lem. 8. Lem. 6 states that including local and mixed predicates in the predicate abstraction step makes it possible to track the number of processes satisfying boolean combinations of those predicates in the boolean program abstraction.

Lemma 6 (abstraction and counting variables). *If $(s^b, m^b) \in \alpha((s, m))$ then for each predicate π belonging to $\text{predsOf}_{\text{exprsOf}(S \cup L)}$ s.t. $\text{comparisonsOf}(\pi) \subseteq \Pi$, $\sum_{\{(pc, \emptyset) | \pi[s, (pc, \emptyset)]\}} m((pc, \emptyset)) = \sum_{\{(pc, \emptyset) | \pi[s^b, (pc, \emptyset)]\}} m^b((pc, \emptyset))$.*

Proof. By definition of $\alpha((s, m))$, $m^b[s, m]$ holds. This means that there is a bijection h from $m = [\sigma_1, \sigma_2, \dots, \sigma_n]$ to $m^b = [\sigma'_1, \sigma'_2, \dots, \sigma'_n]$ s.t. we can associate to each $\sigma_i = (pc, \emptyset)$ in m a $\sigma_{h(i)} = (pc', \emptyset)$ in m^b such that $pc = pc'$ and $\emptyset[s, \emptyset]$. Let π by a boolean combination of predicates in $\Pi_{mix} \cup \Pi_{loc}$. By construction we have that $\pi[s^b, (pc, \emptyset)] \Leftrightarrow \pi[s, (pc, \emptyset)]$. In fact, $\pi[s^b, (pc, \emptyset)]$ and $\pi[s, (pc, \emptyset)]$ coincide on each π s.t. $\text{comparisonsOf}(\pi) \subseteq \Pi$. This implies that for π with $\text{comparisonsOf}(\pi) \subseteq \Pi$,

$$\sum_{\square} \{(\text{pc}, \ell) \mid \pi[s, (\text{pc}, \ell)]\} \mathbb{m}(\text{pc}, \ell) = \sum_{\square} \{(\text{pc}, \ell^b) \mid \pi[s^b, (\text{pc}, \ell^b)]\} \mathbb{m}^b(\text{pc}, \ell^b).$$

Lemma 7 (abstraction and strengthening). *Assume $(s^b, m^b) \in \alpha((s, m))$. If $\omega_{inv}[s, m]$ holds, then $(\exists S. (\text{originOf}(s^b) \wedge \omega_{inv})) [m^b]$ also holds.*

Proof. We know that $\omega_{inv}[s, m]$ holds by assumption and $\text{originOf}(s^b)[s]$ holds since $(s^b, m^b) \in \alpha((s, m))$. Then, $(\text{originOf}(s^b) \wedge \omega_{inv})[s, m]$ also holds, and hence the predicate $((\text{originOf}(s^b) \wedge \omega_{inv})[s]) [m]$. Since for each $(\pi)^\#$ appearing in $\text{vars}(\omega_{inv})$, $\text{comparisonsOf}(\pi) \subseteq \Pi$, we can replace by Lem. 6 m for m^b and obtain the predicate $((\text{originOf}(s^b) \wedge \omega_{inv})[s]) [m^b]$. Finally, the existence of s ensures that $(\exists S. (\text{originOf}(s^b) \wedge \omega_{inv})) [m^b]$ holds. $\square$

Lemma 8 (soundness of strengthening). *If a P configuration satisfying $\omega_{trgt} \in \text{predsOf}_{\text{exprsOf}(\Omega_{PC, S, L})}$ is reachable in P , then q_{trgt} is reachable in $\text{enc}(\text{abst}_\Pi(P))$ after any strengthening wrt. some P invariant $\omega_{inv} \in \text{predsOf}_{\text{exprsOf}(S \cup \Omega_{PC, S, L})}$.*

Proof. By contradiction. Assume a ω_{trgt} configuration is reachable in P . In addition, assume q_{trgt} is not reachable in a counter machine M obtained as the strengthening (wrt. a P counting invariant ω_{inv}) of an encoding of an abstraction $\text{abst}_\Pi(P)$ of P as explained in Section 5.2. Let a feasible run ρ_P without an $\text{enc}(\text{abst}_\Pi(P))_{str}$ feasible run $\rho_{\text{enc}(\text{abst}_\Pi(P))_{str}}$ where $\text{enc}(\text{abst}_\Pi(P))_{str}$ is a strengthening of machine $\text{enc}(\text{abst}_\Pi(P))$ with respect to an invariant ω_{inv} .

According to Def. 1, for each run ρ_P , a non-empty set $\alpha(\rho_P)$ of $\text{abst}_\Pi(P)$ feasible runs exist. Moreover, based on Lem. 3, for each run $\rho_{\text{abst}_\Pi(P)} \in \alpha(\rho_P)$ there exists an $\text{enc}(\text{abst}_\Pi(P))$ feasible run $\rho_{\text{enc}(\text{abst}_\Pi(P))}$ (before strengthening). So, if the run $\rho_{\text{enc}(\text{abst}_\Pi(P))_{str}}$ does not exist, it is because the run $\rho_{\text{enc}(\text{abst}_\Pi(P))}$ was not possible after the strengthening phase.

Let $\rho_{\text{enc}(\text{abst}_\Pi(P))} = (q_{s_0^b}, c_{m_0^b}, \delta_1, \dots, (q_{s_n^b}, c_{m_n^b}))$ and $\rho_P = (s_0, m_0), t_1, \dots, (s_n, m_n)$ with $(s_i^b, m_i^b) \in \alpha((s_i, m_i))$ for each $i : 0 \leq i \leq n$. Because $\rho_{\text{enc}(\text{abst}_\Pi(P))}$ is removed after strengthening, then there exists a step $(q_{s^b}, c_{m^b}) \xrightarrow{(q_{s^b} : \text{op} : q_{s^b'})} (q_{s^b'}, c_{m^b'})$ in $\rho_{\text{enc}(\text{abst}_\Pi(P))}$ such that its corresponding step is impossible. According to strengthening rule in Fig. 10, $\text{str}(s^b, \omega_{inv}) = (\exists S. \text{originOf}(s^b) \wedge \omega_{inv})[s^b, \omega_{inv}]$. So, the fact that the mentioned step is not possible after strengthening implies that either (q_{s^b}, c_{m^b}) does not satisfy $\text{str}(s^b, \omega_{inv})$, or that $(q_{s^b'}, c_{m^b'})$ does not satisfy $\text{str}(s^b, \omega_{inv})$. Both alternatives violate the fact that ω_{inv} is an invariant, that (s, m) and (s', m') are both reachable, and Lem. 7. $\square$

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

Lemma 9. *State reachability is in general undecidable after strengthening.*

Proof. Sketch. Let $CM = (Q, q_0, \Delta, q_F)$ be an arbitrary two counters Minsky machine. Q is a finite set of states, q_0 and q_F are respectively the initial and the final states. The two counters x_1, x_2 range over the natural numbers and start both from 0. Transitions in Δ are either an increment $(q : x_i^{++} : q')$, a decrement $(q : x_i \geq 1; x_i^{--} : q')$ or a test for zero $(q : x_i = 0 : q')$ for $i \in \{1, 2\}$. We describe a concurrent program (in Fig. 11) together with predicates Π_0 , a counting invariant ω_{inv} and a counting predicate ω_{trgt} . The obtained strengthened encoding captures the counter machine CM in the sense that solving the reachability of the target state of the encoding is equivalent to checking reachability of q_F for CM , which is undecidable in general.

Each process of the concurrent program manipulates one local variable tl and four shared variables $\text{c}_1, \text{c}_2, \text{st}, \text{ts}$. The shared variable ts and the local variable tl range both over the three distinct constants $\text{k}_{pm}, \text{k}_{p1}, \text{k}_{p2}$. These are used to identify three types of process instances. Exactly one instance (called main instance henceforth) of the process executes transitions starting with the guard $\text{ts} = \text{k}_{pm}$. The main instance spawns (and joins with) two kinds of auxiliary instances, those that execute the transition with the guard $\text{ts} = \text{k}_{p1}$ and those that execute the transition with the guard $\text{ts} = \text{k}_{p2}$. Henceforth k_{p1} and k_{p2} instances. The variable st ranges over $|Q|$ distinct constants $\{k_q \mid q \in Q\}$ and encodes the state $q \in Q$ of the counter machine CM . Finally, the variable c_1 is used to capture the number of k_{p1} instances, and c_2 the number of k_{p2} instances.

The main instance consists of transitions of the form $(\text{pc}_0 \rightarrow \text{pc}_0 : \text{stmt}_{op})$ that correspond to each transition $(q : \text{op} : q')$ in Δ . More precisely,

1. if op is x_i^{++} , for example t_{inc} in Fig.11, then $\text{stmt}_{op} = (\text{ts} = \text{k}_{pm}; \text{ts} := \text{proc}_i; \text{st} = \text{k}_q; \text{st} := \text{k}_{q'}; \text{c}_1 := \text{c}_1 + 1; \text{spawn});$
2. if op is x_i^{--} , for example t_{dec} in Fig. 11, then $(\text{ts} = \text{k}_{pm}; \text{st} = \text{k}_q; \text{st} := \text{k}_{q'}; \text{c}_2 \geq 1; \text{c}_2 := \text{c}_2 - 1; \text{join});$
3. if op is $(x_i = 0)$, for example t_{tst} in Fig. 11, then $\text{stmt}_{op} = (\text{ts} = \text{k}_{pm}; \text{st} = \text{k}_q; \text{st} := \text{k}_{q'}; \text{c}_1 = 0);$

Finally, we let $\Pi_0 = \{c_1 = 0, c_2 = 0\} \cup \{st = k_q \mid q \in Q\}$, $\omega_{trgt} = (\text{st} = \text{k}_q)^\# \geq 1$ and:

$$\omega_{inv} = \left(\begin{array}{l} (\text{ts} = \text{k}_{pm} \wedge (\text{c}_1 = (@\text{pc}_x \wedge \text{tl} = \text{k}_{p1})^\#) \wedge (\text{c}_2 = (@\text{pc}_x \wedge \text{tl} = \text{k}_{p2})^\#)) \\ \vee (\text{ts} = \text{k}_{p1} \wedge (\text{c}_1 = 1 + (@\text{pc}_x \wedge \text{tl} = \text{k}_{p1})^\#) \wedge (\text{c}_2 = (@\text{pc}_x \wedge \text{tl} = \text{k}_{p2})^\#)) \\ \vee (\text{ts} = \text{k}_{p2} \wedge (\text{c}_1 = (@\text{pc}_x \wedge \text{tl} = \text{k}_{p1})^\#) \wedge (\text{c}_2 = 1 + (@\text{pc}_x \wedge \text{tl} = \text{k}_{p2})^\#)) \end{array} \right)$$

Predicate abstraction will maintain the predicates $c_i = 0$ but will loose their connection with the number of the corresponding instances. Strengthening re-establishes this connection and introduces tests for zero

```

int  c1 := 0
int  c2 := 0
int  st := kq0
int  ts := kpm

process :
int  tl := kpm

// auxiliary processes :
tpr1 : pc0 → pcx :  ts = kp1; ts := kpm; tl := kp1
tpr2 : pc0 → pcx :  ts = kp2; ts := kpm; tl := kp2

// main process :
//(q, x1++, q')
tinc : pc0 → pc0 :  ts = kpm; ts := kp1; st = kq;
                        st := kq'; c1 := c1 + 1; spawn
//(q, x2 >= 1, x2--, q')
tdec : pc0 → pc0 :  ts = kpm; st = kq; st := kq';
                        c2 ≥ 1; c2 := c2 - 1; join
//(q, x1 = 0, q')
ttst : pc0 → pc0 :  ts = kpm; st = kq; st := kq'; c1 = 0
...

```

Fig. 11. Encoding a two counter machine.

on the counters tracking k_{p1}, k_{p2} instances, exactly at the encodings of the transitions t_{tst} . $\square$

5.4 Constrained monotonic abstraction and preorder refinement

This step addresses the state reachability problem for a counter machine $M = (Q, C, \Delta, Q_{init}, C_{init}, q_{trgt})$. As stated in Lem. 9, this problem is in general undecidable for strengthened encodings. The idea here [6] 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 $c \preceq c'$ defined as $\bigwedge_{c \in C} c(c) \leq c'(c)$. Intuitively, $c \preceq c'$ holds if c' can be obtained by “adding more processes to” c . 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 Sec. 5.5). Otherwise, the upward closure $\mathbf{Up}_{\preceq}((q, c))$ 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:

$$\preceq_{i+1} := \{(c, c') \mid c \preceq_i c' \wedge (c \models I \Leftrightarrow c' \models I)\}.$$

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 [5] for more details.

```

input : A machine  $M = (Q, C, \Delta, Q_{init}, C_{init}, q_{trgt})$ 
        and a preorder  $\preceq$ 
output: not_reachable or a run
         $(q_1, c_1), \delta_1, (q_2, c_2), \delta_2, \dots, \delta_n, (q_{trgt}, c)$ 
1 Working :=  $\bigcup_{e \in \text{Min}_{\preceq}(\mathbf{N}^{|C|})} \{((q_{trgt}, e), (q_{trgt}, e))\}$ ,
   Visited :=  $\{\}$ ;
2 while Working  $\neq \{\}$  do
3   pick and remove  $((q, c), \rho)$  from Working;
4   Visited  $\cup = \{((q, c), \rho)\}$ ;
5   if  $(q, c) \in Q_{init} \times C_{init}$  then return
     Check( $M, \preceq, \rho$ );
6   foreach  $\delta \in \Delta$  do
7      $pre = \text{Min}_{\preceq}(\text{Pre}_{\delta}(\mathbf{Up}_{\preceq}((q, c))))$ ;
8     foreach  $(q', c') \in pre$  do
9       if  $c'' \preceq c'$  for some  $((q', c''), -)$  in
         Working  $\cup$  Visited then
10        continue;
11      else
12        foreach  $((q', c''), -) \in \text{Working}$  do
13          if  $c' \preceq c''$  then
            Working = Working  $\setminus \{((q', c''), -)\}$ ;
14        foreach  $((q', c''), -) \in \text{Visited}$  do
15          if  $c' \preceq c''$  then
            Visited = Visited  $\setminus \{((q', c''), -)\}$ ;
16        Working  $\cup = \{((q', c'), (q', c'); \delta; \rho)\}$ 
17 return not_reachable;

```

Algorithm 1: Monotonic abstraction. The procedure **Check**($M, \preceq, \rho$) checks feasibility of the run ρ in the counter machine M . In case the trace is spurious, it refines the preorder $\preceq$ to eliminate that trace.

Lemma 10 (CMA [5]). *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.*

5.5 Simulation on the original concurrent program

A run of the counter machine $(Q, C, \Delta, Q_{init}, C_{init}, q_{trgt})$ 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).

example	P	$enc(\mathbf{abst}_{\Pi}(P))$	outer loop		inner loop		time(s)
			num.	preds.	num.	preds.	
cyclic barrier	5:3:12	13:31:79	4	5	5	1	30
		13:4:35	2	2	2	0	15
		13:8:64	4	3	6	2	35
dynamic barrier	5:2:8	8:8:44	3	3	3	0	20
		8:1:14	2	1	2	0	3
		8:1:8	1	0	1	0	1
flag	5:2:9	8:16:123	4	4	6	2	34
		8:8:67	3	2	3	0	5
		8:4:22	3	2	5	2	18
semaphore	4:2:7	7:32:89	5	5	5	0	68
		7:7:24	3	3	3	0	14
		7:16:66	4	4	6	2	17
maximum	5:2:8	18:16:172	5	5	8	3	489
		18:4:51	3	3	3	0	18
		8:8:48	4	3	6	2	21
parent-child	2:3:10	9:16:48	3	4	5	2	76
		9:1:16	2	1	2	0	2
		9:4:35	3	2	3	0	5
as-many	3:2:6	8:4:34	3	2	6	3	68
		8:1:9	2	1	2	0	2
		8:2:8	2	1	2	0	2
locals	6:3:13	14:8:47	3	3	3	0	16
		14:2:24	2	1	2	0	8
		14:16:95	4	4	4	0	58
shareds	7:3:11	13:32:130	4	5	10	6	160
		13:2:21	2	1	2	0	6
		12:2:20	2	1	2	0	17

Table 1: Checking assertion violation and deadlock with PACMAN

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 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 (predicate abstraction refinement) terminates given the inner loop (monotonic abstraction refinement) terminates. Every iteration of the inner loop terminates.*

Proof. By Lem. 2 and Lem. 6, if the concretization of a run of an encoding is P feasible, then ω_{trgt} is P reachable. Soundness is given by soundness of the predicate abstraction and strengthening (Lem. 8), and soundness of the CMA step (Lem. 10). Termination of each iteration of the inner loop is by well quasi ordering (Lem. 10). Termination of each iteration of the outer loop, given the inner one terminates, is by effectiveness of predicate abstraction (Def. 1) and construction of the encoding (Fig. 9) and strengthening (Fig. 10).

Notice that there is no general guaranty that we establish or refute the safety property (the problem is undecidable). 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. The reported examples that require refinements of the natural preorder are challenging for existing 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. We make use of several optimizations as explained in the following.

- Invariants and unsatisfiable combinations. We discard boolean mappings corresponding to unsatisfi-

able combinations of predicates, for example $\neg(\text{wait} \leq \text{count}) \wedge \neg(\text{count} \leq \text{wait})$ in Fig.1. We also use automatically generated invariants (such as $(\text{wait} \leq \text{count}) \wedge (\text{wait} \geq 0)$ in the same example) to filter the state space. Such heuristics dramatically help our state space exploration algorithm.

- Auxiliary variables. To make the analysis possible, we try to use as few auxiliary variables as possible. An auxiliary variable $a_{(\text{pc}, 1^b, 1^{b'})}$ captures how many processes in local state $(\text{pc}, 1^b)$ will move to state $(\text{pc}, 1^{b'})$ because of the broadcast transition due to a multiple assignment. Each time this number is constant, which turns out to often happen, we discard the auxiliary variable $a_{(\text{pc}, 1^b, 1^{b'})}$.

For each example, we give the number of transitions and variables both in P and in the resulting counter machine.

The tuples under the P column respectively refer to the size of the original program, i.e. number of variables, procedures and transitions in the original program. The tuples under the $\text{enc}(\text{abst}_{\Pi}(P))$ column refer to the size of the counter machine, i.e. number of counters, states and transitions in the extended counter machine. We also state the number of refinement steps and predicates automatically obtained in both refinement loops. We mention the required total time for the analysis. Tool and examples are accessible online².

We report on experiments checking assertion violation and deadlock freedom in Tab.1. For each example, we check three versions, the first one is correct and PACMAN establishes that. The second one is modified so that an assertion can be violated, and the third one is modified so that deadlock is possible. PACMAN for the latter two cases exhibits faulty runs as expected.

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 counting predicates and used them in a framework that combines predicate, counter and constrained monotonic abstraction. Our prototype implementation gave encouraging results and managed to automatically establish or refute program assertions and 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 also aim to generalize to richer variable types.

References

1. P. Abdulla, A. Annichini, S. Bensalem, A. Bouajjani, P. Habermehl, and Y. Lakhnech. Verification of infinite-state systems by combining abstraction and reachability analysis. In N. Halbwachs and D. Peled, editors, *Computer Aided Verification*, volume 1633 of *Lecture Notes in Computer Science*, pages 146–159. Springer Berlin Heidelberg, 1999.
2. 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.
3. P. A. Abdulla, K. Čerāns, B. Jonsson, and Y.-K. Tsay. General decidability theorems for infinite-state systems. In *Proc. LICS '96, 11th IEEE Int. Symp. on Logic in Computer Science*, pages 313–321, 1996.
4. 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.
5. 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.
6. P. A. Abdulla, G. Delzanno, N. B. Henda, and A. Rezine. Regular model checking without transducers (on efficient verification of parameterized systems). In *Tools and Algorithms for the Construction and Analysis of Systems*, pages 721–736. Springer, 2007.
7. P. A. Abdulla, F. Haziza, and L. Holk. Block me if you can! In *Static Analysis*, pages 1–17. Springer, 2014.
8. K. Bansal, E. Koskinen, T. Wies, and D. Zufferey. Structural counter abstraction. In N. Piterman and S. Smolka, editors, *Tools and Algorithms for the Construction and Analysis of Systems*, volume 7795 of *Lecture Notes in Computer Science*, pages 62–77. Springer Berlin Heidelberg, 2013.
9. G. Basler, M. Hague, D. Kroening, C.-H. Ong, T. Wahl, and H. Zhao. Boom: Taking boolean program model checking one step further. In J. Esparza and R. Majumdar, editors, *Tools and Algorithms for the Construction and Analysis of Systems*, volume 6015 of *Lecture Notes in Computer Science*, pages 145–149. Springer Berlin Heidelberg, 2010.
10. L. E. Dickson. Finiteness of the odd perfect and primitive abundant numbers with n distinct prime factors. *Amer. J. Math.*, 35:413–422, 1913.
11. A. Donaldson, A. Kaiser, D. Kroening, and T. Wahl. Symmetry-aware predicate abstraction for shared-variable concurrent programs. In G. Gopalakrishnan and S. Qadeer, editors, *Computer Aided Verification*, volume 6806 of *Lecture Notes in Computer Science*, pages 356–371. Springer Berlin Heidelberg, 2011.
12. J. Esparza, R. Ledesma-Garza, R. Majumdar, P. Meyer, and F. Niksic. An smt-based approach to coverability analysis. In *Computer Aided Verification*, pages 603–619. Springer, 2014.
13. A. Farzan, Z. Kincaid, and A. Podolski. Proofs that count. In *Proceedings of the 41st ACM SIGPLAN*

² <https://gitlab.ida.liu.se/apv/pacman>

- SIGACT Symposium on Principles of Programming Languages*, POPL '14, pages 151–164, New York, NY, USA, 2014. ACM.
14. A. Finkel and P. Schnoebelen. Well-structured transition systems everywhere! *Theoretical Computer Science*, 256(1-2):63–92, 2001.
 15. Z. Ganjei, A. Rezine, P. Eles, and Z. Peng. Abstracting and counting synchronizing processes. In *Verification, Model Checking, and Abstract Interpretation*, pages 227–244. Springer, 2014.
 16. T. A. Henzinger, R. Jhala, R. Majumdar, and S. Qadeer. Thread-modular abstraction refinement. In *Computer Aided Verification*, pages 262–274. Springer, 2003.
 17. 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.
 18. A. Kaiser, D. Kroening, and T. Wahl. Lost in abstraction: Monotonicity in multi-threaded programs. In P. Baldan and D. Gorla, editors, *CONCUR 2014 Concurrency Theory*, volume 8704 of *Lecture Notes in Computer Science*, pages 141–155. Springer Berlin Heidelberg, 2014.
 19. J. Kloos, R. Majumdar, F. Niksic, and R. Piskac. Incremental, inductive coverability. In *Computer Aided Verification*, pages 158–173. Springer, 2013.
 20. L. Zuck and A. Pnueli. Model checking and abstraction to the aid of parameterized systems (a survey). *Computer Languages, Systems & Structures*, 30(34):139 – 169, 2004. Analysis and Verification.