



## 1 Introduction

*Three-valued abstraction* (3VA) [1] is a well-established technique in software verification. It proceeds by generating an abstract state space model of the system to be analysed over predicates with the possible truth values *true*, *false* and *unknown*, where the latter value is used to represent the loss of information due to abstraction. For concurrent software systems composed of many processes, 3VA does not only replace concrete variables by predicates. It also abstracts away entire processes by summarising them into a single approximative component [2], which allows for a substantial reduction of the state space. The evaluation of temporal logic properties on models constructed via three-valued abstraction is known as *three-valued model checking* (3MC) [3]. In three-valued model checking there exist three possible outcomes: *true* and *false* results can be immediately transferred to the modelled system, whereas an *unknown* result does not allow to draw any conclusions about the properties of the system.

Verification techniques based on three-valued abstraction and model checking typically assume that an *explicit* three-valued state space model corresponding to the system to be analysed is constructed and explored [3]. However, explicit-state model checking is known for its high memory demands in comparison to *symbolic* model checking techniques like BDD-based model checking [4] and satisfiability-based bounded model checking (BMC) [5]. The benefits of bounded model checking are that its compressed state space representation as a propositional logic formula allows to handle larger systems than explicit-state techniques, and that its performance profits from the advancements in the SAT solver technology. Although there exist a few works on *three-valued bounded model checking*, these approaches are either solely defined for hardware systems [6], or they require an explicit state space model as input which is then symbolically encoded in propositional logic [7]. It is however not efficient to first translate a given system into an *explicit* state space model before encoding it symbolically for bounded model checking.

In [8] we presented a verification technique for concurrent software systems that allows to directly transfer an abstracted input system  $Sys$ , a temporal logic property  $\psi$  and a bound  $b \in \mathbb{N}$  into a propositional logic formula  $\llbracket Sys, \psi \rrbracket_b$  that encodes the corresponding three-valued bounded model checking problem.  $\llbracket Sys, \psi \rrbracket_b$  is defined over a set of Boolean atoms and the constants *true*, *false* and *unknown*, where the latter only occurs non-negated in the formula. The result of the encoded model checking problem can be obtained via two satisfiability checks. The first check considers an over-approximation  $\llbracket Sys, \psi \rrbracket_b^+$  of the encoding where all *unknowns* are assumed to be *true*. The second check considers an under-approximation  $\llbracket Sys, \psi \rrbracket_b^-$  where all *unknowns* are assumed to be *false*. Unsatisfiability of the over-approximation implies that the bounded model checking result is *false*. Satisfiability of the under-approximation implies that the result is *true*. If only the over-approximation is satisfiable, then the model checking result is *unknown*, which indicates that the current abstraction is too coarse for a definite outcome. While our technique proposed in [8] allows for a compact state space encoding and for the efficient verification of safety and

1 liveness properties under fairness via SAT solving, it does not offer a concept  
2 for abstraction refinement in case of an *unknown* result.

3 In this article, we extend our previous work by introducing a *cause-guided*  
4 *abstraction refinement* procedure for SAT-based three-valued bounded model  
5 checking. For this, we enhanced our propositional logic encoding by adding the  
6 *cause of uncertainty* to each *unknown* that occurs in the formula  $\llbracket Sys, \psi \rrbracket_b$ . A  
7 cause refers to missing information in the current abstraction. We developed  
8 a technique for determining whether this information is relevant for the verifi-  
9 cation task to be solved: If there exists a truth assignment  $\alpha$  that satisfies the  
10 over-approximation  $\llbracket Sys, \psi \rrbracket_b^+$  but not the under-approximation  $\llbracket Sys, \psi \rrbracket_b^-$ , then  
11  $\alpha$  characterises an unconfirmed witness path for the temporal logic property  
12  $\psi$ , i.e. a path with some *unknown* transitions or predicates. Hence, our ap-  
13 proach operates with *implicit* paths given by truth assignments. In contrast to  
14 most counterexample-guided abstraction refinement (CEGAR) techniques [9],  
15 explicit paths do not need to be generated. We next determine all propositional  
16 logic clauses of the under-approximation that are not satisfied under the assign-  
17 ment  $\alpha$ . Our encoding has the property that these clauses contain at least one  
18 *unknown*, and its associated *cause* refers to missing predicates that are required  
19 for a definite model checking result. Our cause-guided refinement procedure  
20 now adds these predicates to the abstraction and constructs the encoding cor-  
21 responding to the refined model checking problem. The procedure is iteratively  
22 applied until a definite result can be obtained.

23 We have integrated iterative cause-guided abstraction refinement into our  
24 SAT-based three-valued bounded model checking tool TVMC, which is available  
25 at [www.github.com/ssfm-up/TVMC](http://www.github.com/ssfm-up/TVMC). In an experimental evaluation, we show  
26 that our novel refinement approach allows to automatically and quickly reach  
27 the right level of abstraction for solving software verification tasks. We also  
28 demonstrate in a number of case studies that TVMC outperforms the similar  
29 tool 3SPOT [2] in most cases. Moreover, we present two enhancements of our  
30 verification technique based on existing work that we have implemented as well:  
31 *Temporal induction* [10] allows us to translate our results of bounded model  
32 checking of safety properties into *unbounded model checking results*, and *spot-*  
33 *light abstraction* [11] enables us to verify *parameterised systems* composed of  
34 arbitrarily many uniform processes.

35 The remainder of this article is organised as follows. In Section 2 we in-  
36 troduce the concurrent systems that we consider in our software verification  
37 approach. Section 3 provides the background on three-valued abstraction and  
38 bounded model checking. Section 4 introduces our propositional logic encod-  
39 ing of software verification tasks and presents a theorem which states that the  
40 satisfiability result for an encoded verification task is equivalent to the result of  
41 the corresponding three-valued bounded model checking problem. In Section 5  
42 we show how our encoding can be augmented with fairness constraints for the  
43 verification of liveness properties. Section 6 introduces our novel cause-guided  
44 abstraction refinement technique. In Section 7 we present the implementation  
45 of our approach, we introduce enhancements that have been implemented as  
46 well, and we present experimental results. Section 8 discusses related work. We

1 conclude this paper in Section 9 and give an outlook on future work.

## 2. Concurrent Software Systems

3 We start with a brief introduction to the systems that we consider in our  
 4 work. A *concurrent software system*  $Sys$  consists of a number of possibly non-  
 5 uniform processes  $P_1$  to  $P_n$  composed in parallel:  $Sys = \parallel_{i=1}^n P_i$ . It is defined  
 6 over a set of variables  $Var = Var_s \cup \bigcup_{i=1}^n Var_i$  where  $Var_s$  is a set of shared  
 7 variables and  $Var_1, \dots, Var_n$  are sets of local variables associated with the pro-  
 8 cesses  $P_1, \dots, P_n$ , respectively. The state space over  $Var$  corresponds to the set  
 9  $S_{Var}$  of all type-correct valuations of the variables. Given a state  $s \in S_{Var}$  and  
 10 an expression  $e$  over  $Var$ , then  $s(e)$  denotes the valuation of  $e$  in  $s$ . An example  
 11 for a concurrent system implementing mutual exclusion is depicted in Figure 1.

$$y : \text{semaphore where } y = 1;$$

$$P_1 :: \left[ \begin{array}{l} \text{loop forever do} \\ \left[ \begin{array}{l} 0: \text{acquire}(y, 1); \\ 1: \text{CRITICAL} \\ \text{release}(y, 1); \end{array} \right] \end{array} \right] \parallel P_2 :: \left[ \begin{array}{l} \text{loop forever do} \\ \left[ \begin{array}{l} 0: \text{acquire}(y, 1); \\ 1: \text{CRITICAL} \\ \text{release}(y, 1); \end{array} \right] \end{array} \right]$$

Figure 1: Concurrent system  $Sys$ .

12 Here we have two processes operating on a shared counting semaphore vari-  
 13 able  $y$ . Processes  $P_i$  can be formally represented as *control flow graphs* (CFGs)  
 14  $G_i = (Loc_i, \delta_i, \tau_i)$  where  $Loc_i = \{[0]_2, \dots, [Loc_i]_2\}$  is a set of control loca-  
 15 tions given as binary numbers,  $\delta_i \subseteq Loc_i \times Loc_i$  is a transition relation, and  
 16  $\tau_i : Loc_i \times Loc_i \rightarrow Op$  is a function labelling transitions with operations from a  
 17 set  $Op$ .

### 18 Definition 1 (Operations).

19 Let  $Var = \{v_1, \dots, v_m\}$  be a set of variables. The set of operations  $Op$  on these  
 20 variables consists of all statements of the form  $assume(e) : v_1 := e_1, \dots, v_m := e_m$   
 21 where  $e, e_1, \dots, e_m$  are expressions over  $Var$ .

22 Hence, every operation consists of a guard and a list of assignments. For  
 23 convenience, we sometimes just write  $e$  instead of  $assume(e)$ . Moreover, we  
 24 omit the *assume* part completely if the guard is *true*. The control flow graphs  
 25  $G_1$  and  $G_2$  corresponding to the processes of our example system are depicted in  
 26 Figure 2.  $G_1$  and  $G_2$  also illustrate the semantics of the operations  $acquire(y, 1)$   
 27 and  $release(y, 1)$ .

28 A concurrent system given by  $n$  individual control flow graphs  $G_1, \dots, G_n$   
 29 can be modelled by one composite CFG  $G = (Loc, \delta, \tau)$  where  $Loc = \times_{i=1}^n Loc_i$ .  
 30  $G$  is the product graph of all individual CFGs. We assume that initially all pro-  
 31 cesses of a concurrent system are at location 0. Moreover, we assume that a  
 32 deterministic initialisation of the system variables is given by an assertion  $\phi$   
 33 over  $Var$ . In our example we have that  $\phi = (y = 1)$ . Now, a computation




Figure 2: Control flow graphs  $G_1$  and  $G_2$  composed in parallel.

1 of a concurrent system corresponds to a sequence where in each step one process is non-deterministically selected and the operation at its current location is attempted to be executed. In case the execution is not blocked by a guard, the variables are updated according to the assignment part and the process advances to the consequent control location. For verifying properties of concurrent systems typically only *fair* computations where all processes infinitely often proceed are considered. We will discuss our notion of fairness in more detail in Section 5. The overall state space  $S$  of a concurrent system corresponds to the set of states over  $Var$  combined with the possible locations, i.e.  $S = Loc \times S_{Var}$ . Hence, each state in  $S$  is a tuple  $\langle l, s \rangle$  with  $l = (l_1, \dots, l_n) \in Loc$  and  $s \in S_{Var}$ .

2 Control flow graphs allow to model concurrent systems formally. For an efficient verification it is additionally required to reduce the state space complexity. For this purpose, we use *three-valued predicate abstraction* [2]. Such an abstraction is an approximation in the sense that all definite verification results (*true*, *false*) obtained for an abstract system can be transferred to the original system. Only *unknown* results necessitate abstraction refinement [12]. In abstract systems operations do not refer to concrete variables but to predicates  $Pred = \{p_1, \dots, p_m\}$  over  $Var$  with the three-valued domain  $\{true, unknown, false\}$ . *Unknown*, typically abbreviated by  $\perp$ , is a valid truth value as we operate with the three-valued *Kleene logic*  $\mathcal{K}_3$  [13] whose semantics is given by the truth tables in Figure 3.

$\wedge$	<i>true</i>	$\perp$	<i>false</i>	$\vee$	<i>true</i>	$\perp$	<i>false</i>	$\neg$	
<i>true</i>	<i>true</i>	$\perp$	<i>false</i>	<i>true</i>	<i>true</i>	<i>true</i>	<i>true</i>	<i>true</i>	<i>false</i>
$\perp$	$\perp$	$\perp$	<i>false</i>	$\perp$	<i>true</i>	$\perp$	$\perp$	$\perp$	$\perp$
<i>false</i>	<i>false</i>	<i>false</i>	<i>false</i>	<i>false</i>	<i>true</i>	$\perp$	<i>false</i>	<i>false</i>	<i>true</i>

Figure 3: Truth tables for the three-valued Kleene logic  $\mathcal{K}_3$ .

The *information order* ' $\leq_{\mathcal{K}_3}$ ' of the Kleene logic is defined as  $\perp \leq_{\mathcal{K}_3} true$ ,  $\perp \leq_{\mathcal{K}_3} false$ , and *true*, *false* incomparable. Operations in abstract systems are of the following form:

$$assume(choice(a, b)) : p_1 := choice(a_1, b_1), \dots, p_m := choice(a_m, b_m)$$

1 where  $a, b, a_1, b_1, \dots, a_m, b_m$  are logical expressions over  $Pred$  and  $choice(a, b)$ -  
 2 expressions have the following semantics:

**Definition 2 (Choice Expressions).**

Let  $s$  be a state over a set of three-valued predicates  $Pred$ . Moreover, let  $a$  and  $b$  be logical expressions over  $Pred$ . Then

$$s(choice(a, b)) = \begin{cases} true & \text{if, and only if, } s(a) \text{ is true,} \\ false & \text{if, and only if, } s(b) \text{ is true,} \\ \perp & \text{else.} \end{cases}$$

3 The application of three-valued predicate abstraction ensures that for any  
 4 state  $s$  and for any expression  $choice(a, b)$  in an abstract control flow graph the  
 5 following holds:  $s(a) = true \Rightarrow s(b) = false$  and  $s(b) = true \Rightarrow s(a) = false$ .  
 6 In particular, this implies that  $s(a)$  and  $s(b)$  are never both *true*. Moreover,  
 7 the following equivalences hold:

$$\begin{aligned} choice(true, false) &\equiv true \\ choice(false, true) &\equiv false \\ choice(false, false) &\equiv \perp \\ choice(a, \neg a) &\equiv a \\ choice(\neg a, a) &\equiv \neg a \\ choice(a, b) &\equiv (a \vee \neg b) \wedge (a \vee b \vee \perp) \\ choice(b, a) &\equiv \neg choice(a, b) \end{aligned}$$

8 A three-valued expression  $choice(a, b)$  over  $Pred$  approximates a Boolean  
 9 expression  $e$  over  $Var$ , written  $choice(a, b) \preceq e$ , if, and only if,  $a$  logically implies  
 10  $e$  and  $b$  logically implies  $\neg e$ . The three-valued approximation relation can be  
 11 straightforwardly extended to operations as described in [2]. An abstract system  
 12  $Sys'$  approximates a concrete system  $Sys$ , written  $Sys' \preceq Sys$ , if the systems  
 13 have isomorphic CFGs and the operations in the abstract system approximate  
 14 the corresponding ones in the concrete system. An example for an abstract  
 15 system that approximates the concrete system in Figure 2 is depicted in Figure  
 16 4. For illustration: the abstract operation  $(y > 0) := choice((y > 0), false)$  sets  
 17 the predicate  $(y > 0)$  to *true* if  $(y > 0)$  was *true* before, and it never sets the  
 18 predicate to *false*. This is a sound three-valued approximation of the concrete  
 19 operation  $y := y + 1$  over the predicate  $(y > 0)$ .

The state space of an abstract system is defined as  $S = Loc \times S_{Pred}$  where  
 $S_{Pred}$  is the set of all possible valuations of the three-valued predicates in  $Pred$ .  
 The state space corresponding to the abstraction of our example is thus  $S =$

$$\begin{aligned} &\{ \langle (0, 0), (y > 0) = true \rangle, \langle (0, 0), (y > 0) = \perp \rangle, \langle (0, 0), (y > 0) = false \rangle \\ &\langle (1, 0), (y > 0) = true \rangle, \langle (1, 0), (y > 0) = \perp \rangle, \langle (1, 0), (y > 0) = false \rangle \\ &\langle (0, 1), (y > 0) = true \rangle, \langle (0, 1), (y > 0) = \perp \rangle, \langle (0, 1), (y > 0) = false \rangle \\ &\langle (1, 1), (y > 0) = true \rangle, \langle (1, 1), (y > 0) = \perp \rangle, \langle (1, 1), (y > 0) = false \rangle \}. \end{aligned}$$




Figure 4: Abstract system represented by control flow graphs  $G'_1$  and  $G'_2$  corresponding to the concrete control flow graphs  $G_1$  and  $G_1$ . Transitions are labelled with abstract operations over  $Pred = \{(y > 0)\}$ .

1        So far we have seen how concurrent systems can be formally represented  
2        and abstracted. Next we will take a look on how model checking of abstracted  
3        systems is defined.

### 4        3. Three-Valued Bounded Model Checking

5        CFGs allow us to model the *control flow* of a concurrent system. The veri-  
6        fication of a system additionally requires to explore a corresponding *state space*  
7        model. Since we use three-valued abstraction, we need a model that incorpo-  
8        rates the truth values *true*, *false* and *unknown*. *Three-valued Kripke structures*  
9        are models with a three-valued domain for transitions and labellings of states:

#### 10        Definition 3 (Three-Valued Kripke Structure).

11        A three-valued Kripke structure over a set of atomic predicates  $AP$  is a tuple  
12         $M = (S, \langle l^0, s^0 \rangle, R, L)$  where

- 13        •  $S$  is a finite set of states,
- 14        •  $\langle l^0, s^0 \rangle \in S$  is the initial state,
- 15        •  $R : S \times S \rightarrow \{\text{true}, \perp, \text{false}\}$  is a transition function with  $\forall \langle l, s \rangle \in S :$   
16         $\exists \langle l', s' \rangle \in S : R(\langle l, s \rangle, \langle l', s' \rangle) \in \{\text{true}, \perp\}$ ,
- 17        •  $L : S \times AP \rightarrow \{\text{true}, \perp, \text{false}\}$  is a labelling function that associates a truth  
18        value with each atomic predicate in each state.

19        A simple example for a three-valued Kripke structure  $M$  over  $AP = \{p\}$  is  
20        depicted in Figure 5.

21        A path  $\pi$  of a Kripke structure  $M$  is a sequence of states  $\langle l^0, s^0 \rangle \langle l^1, s^1 \rangle \dots$   
22        with  $R(\langle l^k, s^k \rangle, \langle l^{k+1}, s^{k+1} \rangle) \in \{\text{true}, \perp\}$ .  $\pi(k)$  denotes the  $k$ -th state of  $\pi$ ,  
23        whereas  $\pi^k$  denotes the  $k$ -th suffix  $\pi(k)\pi(k+1)\pi(k+2)\dots$  of  $\pi$ . By  $\Pi_M$   
24        we denote the set of all paths of  $M$  starting in the initial state. Paths are considered  
25        for the evaluation of temporal logic properties of Kripke structures.

26        A concurrent system  $Sys = \parallel_{i=1}^n P_i$  abstracted over a set of predicates  $Pred$   
27        can be represented as a three-valued Kripke structure according to the following  
28        definition:




Figure 5: Three-valued Kripke structure.

1 **Definition 4 (Concurrent System as Three-Valued Kripke Structure).**

2 Let  $Sys = \parallel_{i=1}^n P_i$  over  $Var$  be a concurrent system given by a composite control flow graph  $G = (Loc, \delta, \tau)$  and an initial state predicate  $\phi$ . Moreover, let  
3 control flow graph  $G = (Loc, \delta, \tau)$  and an initial state predicate  $\phi$ . Moreover, let  
4  $Pred$  be a set of predicates over  $Var$ . The corresponding three-valued Kripke  
5 structure is a tuple  $M = (S, \langle l^0, s^0 \rangle, R, L)$  over a set of atomic predicates  
6  $AP = Pred \cup \{(loc_i = j) \mid i \in [1..n], j \in Loc_i\}$ , where  $(loc_i = j)$  denotes  
7 that the process  $P_i$  is currently at control location  $j$ , with

8 •  $S := Loc \times S_{Pred}$ ,

9 •  $\langle l^0, s^0 \rangle := \langle (0, \dots, 0), s \rangle$  where  $s \in S$  with  $s(\phi) = true$ ,

10 •  $R(\langle l, s \rangle, \langle l', s' \rangle) := \bigvee_{i=1}^n R_i(\langle l, s \rangle, \langle l', s' \rangle) :=$   
11  $\bigvee_{i=1}^n (\delta_i(l_i, l'_i) \wedge \bigwedge_{i' \neq i} (l_{i'} = l'_{i'}) \wedge s(choice(a, b)) \wedge \bigwedge_{j=1}^m s'(p_j) = s(choice(a_j, b_j)))$

12 assuming that  $l_i$  is the single location of  $P_i$  in the composite location  $l$  and

13  $\tau_i(l_i, l'_i) = assume(choice(a, b)) : p_1 := choice(a_1, b_1), \dots, p_m := choice(a_m, b_m)$ ,

14 •  $L(\langle l, s \rangle, p) := s(p)$  for each  $p \in Pred$ ,

15 •  $L(\langle l, s \rangle, (loc_i = j)) := \begin{cases} true & \text{if } l_i = j \\ false & \text{else} \end{cases}$

16 assuming that  $l_i$  is the single location of  $P_i$  in the composite location  $l$ .

18 In [2] it has been shown that there is a one-to-one correspondence between  
19 computations of a concurrent system  $Sys$  and paths of a three-valued Kripke  
20 structure  $M$  modelling the state space of  $Sys$ . Moreover, we get Lemma 1 from  
21 [2] that establishes a relation between different abstract models of a system.

22 **Lemma 1**

23 Let  $Sys = \parallel_{i=1}^n P_i$  over  $Var$  be a concurrent system. Let  $AP_a$  and  $AP_r$  be  
24 sets of atomic predicates over  $Var$  with  $AP_a \subset AP_r$ . Moreover, let  $M_a =$



1  $(S_a, \langle l_a^0, s_a^0 \rangle, R_a, L_a)$  be the three-valued Kripke structure modelling the state  
 2 space of *Sys* abstracted over  $AP_a$ , and let  $M_r = (S_r, \langle l_r^0, s_r^0 \rangle, R_r, L_r)$  be the  
 3 three-valued Kripke structure modelling the state space of *Sys* abstracted over  
 4  $AP_r$ . Then the following holds:

- 5 1. For every path  $\pi_a \in \Pi_{M_a}$  there exists a path  $\pi_r \in \Pi_{M_r}$  with  $\forall k \in \mathbb{N} :$   
 6  $R_a(\pi_a(k), \pi_a(k+1)) = \text{true} \Rightarrow R_r(\pi_r(k), \pi_r(k+1)) = \text{true}$  and  $\forall p \in$   
 7  $AP_a : L_a(\pi_a(k), p) \leq_{\mathcal{K}_3} L_r(\pi_r(k), p)$
- 8 2. For every path  $\pi_r \in \Pi_{M_r}$  there exists a path  $\pi_a \in \Pi_{M_a}$  with  $\forall k \in \mathbb{N} :$   
 9  $R_r(\pi_r(k), \pi_r(k+1)) \neq \text{false} \Rightarrow R_a(\pi_a(k), \pi_a(k+1)) \neq \text{false}$  and  $\forall p \in$   
 10  $AP_a : L_a(\pi_a(k), p) \leq_{\mathcal{K}_3} L_r(\pi_r(k), p)$

11 Hence, for each path  $\pi_a$  in the more abstract model  $M_a$  there is a path  $\pi_r$   
 12 in the finer model  $M_r$  such that  $\pi_r$  is a refinement of  $\pi_a$  in terms of the logic  
 13  $\mathcal{K}_3$ . Moreover, for each path  $\pi_r$  in the finer model  $M_r$  there is a path  $\pi_a$  in the  
 14 more abstract model  $M_a$  such that  $\pi_a$  is an abstraction of  $\pi_r$ . An example of  
 15 two paths with such an abstraction-refinement relation is depicted in Figure 6.  
 16 As we can see, every labelling with a definite value in  $\pi_a$  has the same definite  
 17 value in  $\pi_r$ , whereas labellings with *unknown* values in  $\pi_a$  may have different  
 18 values in  $\pi_r$ . Furthermore, every definite transition in  $\pi_a$  is also definite in  $\pi_r$ ,  
 19 whereas *unknown* transitions in  $\pi_a$  may be either still *unknown* or definite  $\pi_r$ .




Figure 6: Paths with an abstraction-refinement relation.

20 According to Lemma 1, definite path information is preserved if we consider  
 21 a refinement of an abstract model. Subsequently, we will see that this extends  
 22 to temporal logic properties evaluated on abstract and refined models.

23 The number of states of a Kripke structure modelling a given system is  
 24 exponential in the number of its locations and variables. State explosion is  
 25 the major challenge in software model checking. One approach to cope with  
 26 the state explosion problem is to use a symbolic and therefore more compact  
 27 representation of the Kripke structure. In SAT-based bounded model checking  
 28 [5] all possible path prefixes up to a bound  $b \in \mathbb{N}$  are encoded in a propositional  
 29 logic formula. The formula is then conjuncted with an encoding of the temporal  
 30 logic property to be checked. In case the overall formula is satisfiable, the

1 satisfying truth assignment characterises a witness path of length  $b$  for the  
 2 property in the state space of the encoded system. Hence, bounded model  
 3 checking can be performed via satisfiability solving. We now briefly recapitulate  
 4 the syntax and bounded semantics of the linear temporal logic (LTL):

**Definition 5 (Syntax of LTL).**

Let  $AP$  be a set of atomic predicates and  $p \in AP$ . The syntax of LTL formulae  $\psi$  is given by

$$\psi ::= p \mid \neg p \mid \psi \vee \psi \mid \psi \wedge \psi \mid \mathbf{G}\psi \mid \mathbf{F}\psi \mid \mathbf{X}\psi.$$

5 The temporal operator  $\mathbf{G}$  is read as *globally*,  $\mathbf{F}$  is read as *finally* (or *eventually*),  
 6 and  $\mathbf{X}$  is read as *next*. For the sake of simplicity, we omit the temporal  
 7 operator  $\mathbf{U}$  (*until*). Due to the extended domain of truth values in three-valued  
 8 Kripke structures, the bounded evaluation of LTL formulae is based on the  
 9 Kleene logic  $\mathcal{K}_3$  (compare Section 2). Based on  $\mathcal{K}_3$ , LTL formulae can be evalu-  
 10 ated on  $b$ -bounded path prefixes of three-valued Kripke structures. Such finite  
 11 prefixes  $\pi(0) \dots \pi(b)$  can still represent infinite paths if the prefix has a *loop*,  
 12 i.e. the last state  $\pi(b)$  has a successor state that is also part of the prefix.

**Definition 6 ( $b$ -Loop).**

14 Let  $\pi$  be a path of a three-valued Kripke structure  $M$  and let  $r, b \in \mathbb{N}$  with  $r \leq b$ .  
 15 Then  $\pi$  has a  $(b, r)$ -loop if  $R(\pi(b), \pi(r)) \in \{\text{true}, \perp\}$  and  $\pi$  is of the form  $v \cdot w^\omega$   
 16 where  $v = \pi(0) \dots \pi(r-1)$  and  $w = \pi(r) \dots \pi(b)$ .  $\pi$  has a  $b$ -loop if there exists  
 17 an  $r \in \mathbb{N}$  with  $r \leq b$  such that  $\pi$  has a  $(b, r)$ -loop.

18 An example for a path with a loop is depicted in Figure 7. Let  $b = 3$  be  
 19 then bound. As we can see, the path  $\pi$  has a  $(3, 1)$ -loop and therefore a 3-loop.




Figure 7: Path with a  $b$  loop.

20 For the bounded evaluation of LTL formulae on paths of Kripke structures  
 21 we have to distinguish between paths *with* and *without* a  $b$ -loop.

**Definition 7 (Three-Valued Bounded Evaluation of LTL).**

23 Let  $M = (S, \langle l^0, s^0 \rangle, R, L)$  over  $AP$  be a three-valued Kripke structure. More-  
 24 over, let  $b \in \mathbb{N}$  and let  $\pi$  be a path of  $M$  with a  $b$ -loop. Then the  $b$ -bounded  
 25 evaluation of an LTL formula  $\psi$  on  $\pi$ , written  $[\pi \models_b^k \psi]$  where  $k \leq b$  denotes  
 26 the current position along the path, is inductively defined as follows:

$$\begin{aligned}
[\pi \models_b^k p] &\equiv L(\pi(k), p) \\
[\pi \models_b^k \neg p] &\equiv \neg L(\pi(k), p) \\
[\pi \models_b^k \psi \vee \psi'] &\equiv [\pi \models_b^k \psi] \vee [\pi \models_b^k \psi'] \\
[\pi \models_b^k \psi \wedge \psi'] &\equiv [\pi \models_b^k \psi] \wedge [\pi \models_b^k \psi'] \\
[\pi \models_b^k \mathbf{G}\psi] &\equiv \bigwedge_{k' \geq k} (R(\pi(k'), \pi(k' + 1)) \wedge [\pi \models_b^{k'} \psi]) \\
[\pi \models_b^k \mathbf{F}\psi] &\equiv \bigvee_{k' \geq k} ([\pi \models_b^{k'} \psi] \wedge \bigwedge_{k''=k}^{k'-1} R(\pi(k''), \pi(k'' + 1))) \\
[\pi \models_b^k \mathbf{X}\psi] &\equiv R(\pi(k), \pi(k + 1)) \wedge [\pi \models_b^{k+1} \psi]
\end{aligned}$$

If  $\pi$  is a path without a  $b$ -loop then the  $b$ -bounded evaluation of  $\psi$  is defined as:

$$\begin{aligned}
[\pi \models_b^k \mathbf{G}\psi] &\equiv \text{false} \\
[\pi \models_b^k \mathbf{F}\psi] &\equiv \bigvee_{k'=k}^b ([\pi \models_b^{k'} \psi] \wedge \bigwedge_{k''=k}^{k'-1} R(\pi(k''), \pi(k'' + 1))) \\
[\pi \models_b^k \mathbf{X}\psi] &\equiv \text{if } k < b \text{ then } R(\pi(k), \pi(k + 1)) \wedge [\pi \models_b^{k+1} \psi] \text{ else false}
\end{aligned}$$

1 The other cases are identical to the case where  $\pi$  has a  $b$ -loop. The universal  
2 bounded evaluation of  $\psi$  on an entire Kripke structure  $M$  is  $[M \models_{U,b} \psi] \equiv$   
3  $\bigwedge_{\pi \in \Pi_M} [\pi \models_b^0 \psi]$ . The existential bounded evaluation of  $\psi$  on a Kripke structure  
4 is  $[M \models_{E,b} \psi] \equiv \bigvee_{\pi \in \Pi_M} [\pi \models_b^0 \psi]$ .

Checking temporal logic properties for three-valued Kripke structures is what is known as three-valued model checking [3]. Universal model checking can always be transformed into existential model checking based on the equation

$$[M \models_{U,b} \psi] = \neg [M \models_{E,b} \neg \psi].$$

5 From now on we only consider the existential case, since it is the basis  
6 of satisfiability-based bounded model checking. Bounded model checking [5]  
7 is typically performed incrementally, i.e.  $b$  is iteratively increased until the  
8 property can be either proven or a completeness threshold [14] is reached. In  
9 the three-valued scenario there exist three possible outcomes: *true*, *false* and  $\perp$ .  
10 For our example Kripke structure  $M$  we have that  $[M \models_{E,0} \mathbf{F}p]$  evaluates to  $\perp$   
11 and  $[M \models_{E,1} \mathbf{F}p]$  evaluates to *true*, which is witnessed by the 1-bounded path  
12 prefix  $\langle l^0, s^0 \rangle \langle l^2, s^2 \rangle$ .

13 By combining Lemma 1 with the definitions of LTL we get Corollary 1:

14 **Corollary 1**

15 Let  $Sys = \parallel_{i=1}^n P_i$  over  $Var$  be a concurrent system. Let  $AP_a$  and  $AP_r$  be sets  
16 of atomic predicates over  $Var$  with  $AP_a \subset AP_r$ . Let  $M_a = (S_a, \langle l_a^0, s_a^0 \rangle, R_a, L_a)$   
17 be the three-valued Kripke structure modelling the state space of  $Sys$  abstracted  
18 over  $AP_a$ , and let  $M_r = (S_r, \langle l_r^0, s_r^0 \rangle, R_r, L_r)$  be the three-valued Kripke structure  
19 modelling the state space of  $Sys$  abstracted over  $AP_r$ . Moreover, let  $\psi$  be an LTL  
20 formula and  $b \in \mathbb{N}$  be a bound. Then the following holds:

- 21 1.  $[M_a \models_{E,b} \psi] = \text{true} \Rightarrow [M_r \models_{E,b} \psi] = \text{true}$   
22 2.  $[M_a \models_{E,b} \psi] = \text{false} \Rightarrow [M_r \models_{E,b} \psi] = \text{false}$

1 Hence, all definite model checking results obtained under three-valued ab-  
 2 straction can be immediately transferred to any refined model, and thus, also  
 3 the concrete system  $Sys$  modelled by the three-valued Kripke structure, whereas  
 4 an *unknown* result tells us that the current level of abstraction is too coarse.  
 5 For the latter case we will present an automatic refinement procedure in Section  
 6 6 that refines the abstraction by adding predicates to  $AP$ .

7 In the next section we define a propositional logic encoding of three-valued  
 8 bounded model checking tasks for abstracted concurrent systems. Our encod-  
 9 ing allows to immediately transfer verification tasks into a propositional logic  
 10 formulae that can be then processed via a SAT solver. Thus, the expensive  
 11 construction of an explicit Kripke structure is not required in our approach.  
 12 The state space of the system under consideration as well as the property to  
 13 be checked will be implicitly contained in the propositional logic encoding, and  
 14 the model checking result will be equivalent to the result of the corresponding  
 15 satisfiability tests.

#### 16 4. Propositional Logic Encoding

In our previous work [15] we showed that the three-valued bounded model  
 checking problem  $[M \models_{E,b} \psi]$ , where  $M$  is given as an *explicit* Kripke structure,  
 can be reduced to two classical SAT problems. Here we show that for a given sys-  
 tem  $Sys$  abstracted over  $Pred$ , a temporal logic property  $\psi$ , and a bound  $b \in \mathbb{N}$ ,  
 it is not even necessary to consider the corresponding model checking problem.  
 We can immediately construct a propositional logic encoding  $\llbracket Sys, \psi \rrbracket_b$  and per-  
 form two SAT checks. One check considers an *over-approximating completion*  
 of the encoding, marked with ‘+’, where all  $\perp$ ’s are assumed to be *true*:

$$\llbracket Sys, \psi \rrbracket_b^+ := \llbracket Sys, \psi \rrbracket_b[\perp \mapsto true]$$

and the second check considers an *under-approximating completion*, marked with  
 a ‘-’, where all  $\perp$ ’s are assumed to be *false*:

$$\llbracket Sys, \psi \rrbracket_b^- := \llbracket Sys, \psi \rrbracket_b[\perp \mapsto false].$$

17 Here  $[\perp \mapsto z]$  with  $z \in \{true, false\}$  denotes the assumption that  $\perp$  mapped  $z$ .  
 18 We will show that the following holds:

$$[M \models_{E,b} \psi] = \begin{cases} true & \text{if } SAT(\llbracket Sys, \psi \rrbracket_b^-) = true \\ false & \text{if } SAT(\llbracket Sys, \psi \rrbracket_b^+) = false \\ \perp & \text{else} \end{cases}$$

19 Hence, it is not required to construct and explore an explicit Kripke structure  
 20  $M$  modelling the state space of  $Sys$ . All we need to do is to construct  $\llbracket Sys, \psi \rrbracket_b$   
 21 and check the satisfiability of its under- and over-approximation in order to  
 22 obtain the result of the corresponding three-valued model checking problem.

1 The formula  $\llbracket Sys, \psi \rrbracket_b$  is defined over a set of Boolean atoms and over *true*,  
2 *false* and  $\perp$ . We now give a step-by-step description on how  $\llbracket Sys, \psi \rrbracket_b$  can be  
3 constructed for a concurrent system  $Sys = \parallel_{i=1}^n P_i$  abstracted over a set of  
4 predicates *Pred* and given by a number of control flow graphs  $G_i = (Loc_i, \delta_i, \tau_i)$   
5 with  $1 \leq i \leq n$ , a temporal logic property  $\psi \in \text{LTL}$ , and a bound  $b \in \mathbb{N}$ . The  
6 construction of  $\llbracket Sys, \psi \rrbracket_b$  is divided into the translation of the abstract system  
7 into a formula  $\llbracket Sys \rrbracket_b$  and the translation of the property  $\psi$  into a formula  $\llbracket \psi \rrbracket_b$ .

We start with the encoding of the system, which first requires to encode its  
states as propositional logic formulae. Since a state of a concurrent system is a  
tuple  $\langle l, s \rangle$  where  $l$  is a composite control flow location and  $s$  is a valuation of  
all predicates in *Pred*, we encode  $l$  and  $s$  separately. First, we introduce a set of  
Boolean atoms for the encoding of locations. A composite location  $(l_1, \dots, l_n) \in$   
 $Loc$  is a list of single locations  $l_i \in Loc_i$  where  $Loc_i = \{0, \dots, |Loc_i|\}$  and  $i$   
is the identifier of the associated process  $P_i$ . Each  $l_i$  is a binary number from the  
domain  $\{[0]_2, \dots, [|Loc_i|]_2\}$ . We assume that all these numbers have  $d_i$  digits  
where  $d_i$  is the number required to binary represent the maximum value  $|Loc_i|$ .  
We introduce the following set of Boolean atoms:

$$LocAtoms := \{l_i[j] \mid i \in [1..n], j \in [1..d_i]\}$$

8 Hence, for each process  $P_i$  of the system we introduce  $d_i$  Boolean atoms,  
9 each referring to a distinct digit along the binary representation of its locations.  
10 The atoms now allow us to define the following encoding of locations:

**Definition 8 (Encoding of Locations).**

Let the location  $l_i \in \{0, \dots, |Loc_i|\}$  be given as a binary number. Moreover, let  
 $l_i(j)$  be a function evaluating to *true* if the  $j$ -th digit of  $l_i$  is 1, and to *false*  
otherwise. Then  $l_i$  can be encoded in propositional logic as follows:

$$enc(l_i) := \bigwedge_{j=1}^{d_i} ((l_i[j] \wedge l_i(j)) \vee (\neg l_i[j] \wedge \neg l_i(j)))$$

11 Let  $l = (l_1, \dots, l_n)$  be a composite location. Then  $enc(l) := \bigwedge_{i=1}^n enc(l_i)$ .

12 Note that since the function  $l_i(j)$  evaluates to *true* or *false* an encoding  
13  $enc(l_i)$  can be always simplified to a conjunction of literals over *LocAtoms*. For  
14 instance, the initial location  $(0, 0)$  of our example system from Section 2 will be  
15 encoded to  $\neg l_1[1] \wedge \neg l_2[1]$  and the location  $(0, 1)$  will be encoded to  $\neg l_1[1] \wedge l_2[1]$ .

Next, we encode the predicate part of states. Let  $s \in S_{Pred}$  where  $Pred =$   
 $\{p_1, \dots, p_m\}$ . We introduce the following set of Boolean atoms:

$$PredAtoms := \{p[j] \mid p \in Pred, j \in \{u, t\}\}$$

16 Hence, for each three-valued predicate  $p$  we introduce two Boolean atoms.  
17 The atom  $p[u]$  will let us indicate whether  $p$  evaluates to *unknown*, and  $p[t]$  will  
18 let us indicate whether it evaluates to *true* or *false*:

**Definition 9 (Encoding of States over Predicates).**

Let  $p \in Pred$  and let  $val \in \{true, \perp, false\}$ . Then  $(p = val)$  can be logically

encoded follows:

$$enc(p = val) := \begin{cases} \neg p[u] \wedge p[t] & \text{if } val = true \\ \neg p[u] \wedge \neg p[t] & \text{if } val = false \\ p[u] & \text{if } val = \perp \end{cases}$$

1 Let  $s$  be a state over  $Pred$ . Then  $enc(s) := \bigwedge_{p \in Pred} enc(p = s(p))$ .

For an overall state  $\langle l, s \rangle \in S$  we consequently get

$$enc(\langle l, s \rangle) := enc(l) \wedge enc(s).$$

Since  $enc(\langle l, s \rangle)$  yields a conjunction of literals, there exists exactly one satisfying truth assignment  $\alpha : LocAtoms \cup PredAtoms \rightarrow \{true, false\}$  for a state encoding. We denote the assignment characterising an encoded state  $\langle l, s \rangle$  by  $\alpha_{\langle l, s \rangle}$ . For instance, the initial state  $\langle (0, 0), (y > 0) = true \rangle$  of our abstracted example system will be encoded to

$$Init = \neg l_1[1] \wedge \neg l_2[1] \wedge \neg p[u] \wedge p[t]$$

where  $p = (y > 0)$ , i.e. we abbreviate  $(y > 0)$  by  $p$ . The assignment characterising  $Init$  is

$$\alpha_{\langle (0,0), (y>0)=true \rangle} : l_1[1] \mapsto false, l_2[1] \mapsto false, p[u] \mapsto false, p[t] \mapsto true.$$

2 .

3 The encoding function  $enc$  can be extended to *logical expressions* in negation  
4 normal form (NNF), which we require for our later transition encoding:

**Definition 10 (Encoding of Logical Expressions).**

Let  $p \in Pred$  and  $e, e'$  logical expressions in NNF over  $Pred \cup \{true, \perp, false\}$ . Let  $val \in \{true, \perp, false\}$ . Then the encoding of a logical expression is inductively defined as follows:

$$\begin{aligned} enc(val) &:= val \\ enc(\neg val) &:= \neg val \\ enc(p) &:= (p[u] \wedge \perp) \vee (\neg p[u] \wedge p[t]) \\ enc(\neg p) &:= (p[u] \wedge \perp) \vee (\neg p[u] \wedge \neg p[t]) \\ enc(e \wedge e') &:= enc(e) \wedge enc(e') \\ enc(e \vee e') &:= enc(e) \vee enc(e') \\ enc(choice(e, e')) &:= enc((e \vee NNF(\neg e')) \wedge (e \vee e' \vee \perp)) \end{aligned}$$

Next, we take a look at how the transition relation of an abstracted system can be encoded. We will construct a propositional logic formula

$$\llbracket Sys \rrbracket_b = Init_0 \wedge Trans_{0,1} \wedge \dots \wedge Trans_{b-1,b}$$

that exactly characterises path prefixes of length  $b \in \mathbb{N}$  in the state space of the

system  $Sys$  abstracted over  $Pred$ . Since we consider states as parts of such prefixes, we have to extend the encoding of states by index values  $k \in \{0, \dots, b\}$  where  $k$  denotes the position along a path prefix. For this we introduce the notion of indexed encodings. Let  $F$  be a propositional logic formula over  $Atoms = LocAtoms \cup PredAtoms$  and  $true$ ,  $false$  and  $\perp$ . Then  $F_k$  stands for  $F[a/a_k \mid a \in Atoms]$ . Our overall encoding will be thus defined over the set  $Atoms_{[0,b]} = \{a_k \mid a \in Atoms, 0 \leq k \leq b\}$ . An assignment  $\alpha_{\langle l, s \rangle}$  to the atoms in a subset  $Atoms_{[k,k]} \subseteq Atoms_{[0,b]}$  thus characterises a state  $\langle l, s \rangle$  at position  $k$  of a path prefix, whereas an assignment  $\alpha_{\langle l^0, s^0 \rangle \dots \langle l^b, s^b \rangle}$  to the atoms in  $Atoms_{[0,b]}$  characterises an entire path prefix  $\langle l^0, s^0 \rangle \dots \langle l^b, s^b \rangle$ . Since all execution paths start in the initial state of the system, we extend its encoding by the index 0, i.e. we get

$$Init_0 = \neg l_1[1]_0 \wedge \neg l_2[1]_0 \wedge \neg p[u]_0 \wedge p[t]_0.$$

1 The encoding of all possible state space transitions from position  $k$  to  $k+1$   
 2 is defined as follows:

**Definition 11 (Encoding of Transitions).**

Let  $Sys = \parallel_{i=1}^n P_i$  over  $Pred$  be an abstracted concurrent system given by the single control flow graphs  $G_i = (Loc_i, \delta_i, \tau_i)$  with  $1 \leq i \leq n$ . Then all possible transitions for position  $k$  to  $k+1$  can be encoded in propositional logic as follows:

$$Trans_{k,k+1} :=$$

$$\bigvee_{i=1}^n \bigvee_{(l_i, l'_i) \in \delta_i} (enc(l_i)_k \wedge enc(l'_i)_{k+1} \wedge \bigwedge_{i' \neq i} (idle(i')_{k,k+1}) \wedge enc(\tau_i(l_i, l'_i))_{k,k+1})$$

where

$$idle(i')_{k,k+1} := \bigwedge_{j=1}^{d_{i'}} (l_{i'}[j]_k \leftrightarrow l_{i'}[j]_{k+1})$$

and

$$enc(\tau_i(l_i, l'_i))_{k,k+1} := \begin{aligned} & enc(choice(a, b))_k \\ & \wedge \bigwedge_{j=1}^m ( (enc(a_j)_k \wedge enc(p_j = true)_{k+1}) \\ & \quad \vee (enc(b_j)_k \wedge enc(p_j = false)_{k+1}) \\ & \quad \vee (enc(\neg a_j \wedge \neg b_j)_k [\perp \mapsto true] \wedge enc(p_j = \perp)_{k+1})) \end{aligned}$$

3 assuming that  $\tau_i(l_i, l'_i) = assume(choice(a, b)) : p_1 := choice(a_1, b_1), \dots, p_m :=$   
 4  $choice(a_m, b_m)$ .

5 Thus, we iterate over the system's processes  $P_i$  and over the processes' con-  
 6 trol flow transitions  $\delta_i(l_i, l'_i)$ . Now we construct the  $k$ -indexed encoding of a  
 7 source location  $l_i$  and conjunct it with the  $(k+1)$ -indexed encoding of a desti-  
 8 nation location  $l'_i$ . This gets conjuncted with the sub formula  $\bigwedge_{i' \neq i} idle(i')_{k,k+1}$   
 9 which encodes that all processes different to the currently considered process  $P_i$   
 10 are idle, i.e. do not change their control flow location, while  $P_i$  proceeds. The  
 11 last part of the transition encoding concerns the operation associated with the

<sub>1</sub> control flow transition  $\delta_i(l_i, l'_i)$ : The sub formula  $enc(\tau_i(l_i, l'_i))_{k,k+1}$  evaluates to  
<sub>2</sub> *true* for assignments  $\alpha_{\langle l, s \rangle \langle l', s' \rangle}$  to the atoms in  $Atoms_{[k,k+1]}$  that characterise  
<sub>3</sub> pairs of states  $s$  and  $s'$  over  $Pred$  where the guard of the operation  $\tau_i(l_i, l'_i)$  is  
<sub>4</sub> *true* in  $s$  and the execution of the operation in  $s$  definitely results in the state  
<sub>5</sub>  $s'$ . The operation encoding evaluates to  $\perp$  for states  $s$  and  $s'$  where the guard  
<sub>6</sub> of the operation is  $\perp$  in  $s$  or where it is *unknown* whether the execution of  
<sub>7</sub> the operation in  $s$  results in the state  $s'$ . In all other cases  $enc(\tau_i(l_i, l'_i))_{k,k+1}$   
<sub>8</sub> evaluates to *false*. Our transition encoding requires that an operation  $\tau_i(l_i, l'_i)$   
<sub>9</sub> assigns to all predicates in  $Pred$ : Thus, if a predicate  $p$  is not modified by the  
<sub>10</sub> operation we assume that  $p := p$  is part of the assignment list.

The encoding of the control flow transition  $\delta_1(0, 1)$  of our abstract example system with

$$\tau_1(0, 1) = (assume(p) : p := choice(false, \neg p)),$$

where  $p$  abbreviates  $(y > 0)$ , yields the following:

$$\begin{aligned}
 enc(0)_k &= \neg l_1[1]_k \\
 \wedge &\quad \wedge \\
 enc(1)_{k+1} &= l_1[1]_{k+1} \\
 \wedge &\quad \wedge \\
 idle(2)_{k,k+1} &= (l_2[1]_k \leftrightarrow l_2[1]_{k+1}) \\
 \wedge &\quad \wedge \\
 enc(\tau_1(0, 1))_{k,k+1} &= ((p[u]_k \wedge \perp) \vee (\neg p[u]_k \wedge p[t]_k)) \wedge \\
 &\quad ((false \wedge (\neg p[u]_{k+1} \wedge p[t]_{k+1})) \\
 &\quad \vee (((p[u]_k \wedge \perp) \vee (\neg p[u]_k \wedge \neg p[t]_k)) \wedge (\neg p[u]_{k+1} \wedge \neg p[t]_{k+1})) \\
 &\quad \vee (((p[u]_k \wedge true) \vee (\neg p[u]_k \wedge p[t]_k)) \wedge (p[u]_{k+1})))
 \end{aligned}$$

The encoding of the operation only evaluates to *true* for assignments to the atoms in  $Atoms_{[k,k+1]}$  that characterise a predicate state  $s$  at position  $k$  with  $s(p) = true$  and a state  $s'$  at position  $k+1$  with  $s'(p) = \perp$ . An overall satisfying assignment for this encoding is

$$\begin{aligned}
 \alpha_{\langle (0,0), (y>0)=true \rangle \langle (1,0), (y>0)=\perp \rangle} : & \quad l_1[1]_k &\mapsto & \quad false, \\
 & \quad l_2[1]_k &\mapsto & \quad false, \\
 & \quad l_1[1]_{k+1} &\mapsto & \quad true, \\
 & \quad l_2[1]_{k+1} &\mapsto & \quad false, \\
 & \quad p[u]_k &\mapsto & \quad false, \\
 & \quad p[t]_k &\mapsto & \quad true, \\
 & \quad p[u]_{k+1} &\mapsto & \quad true
 \end{aligned}$$

characterising the definite transition between the states  $\langle (0, 0), (y > 0) = true \rangle$



and  $\langle(1, 0), (y > 0) = \perp\rangle$ . The assignments

$$\begin{aligned} &\alpha_{\langle(0, l_2), (y > 0) = true\rangle \langle(1, l_2), (y > 0) = false\rangle}, \\ &\alpha_{\langle(0, l_2), (y > 0) = \perp\rangle \langle(1, l_2), (y > 0) = false\rangle}, \\ &\alpha_{\langle(0, l_2), (y > 0) = \perp\rangle \langle(1, l_2), (y > 0) = \perp\rangle} \end{aligned}$$

1 with  $l_2 \in \{0, 1\}$  yield *unknown* for the encoding and hereby correctly char-  
 2 acterise  $\perp$ -transitions in the abstract state space. All other assignments yield  
 3 *false* indicating that corresponding pairs of states do not characterise valid tran-  
 4 sitions.

The encoding definitions now allow us to construct the propositional logic formula

$$\llbracket Sys \rrbracket_b = Init_0 \wedge Trans_{0,1} \wedge \dots \wedge Trans_{b-1,b}$$

5 that characterises all possible path prefixes of length  $b \in \mathbb{N}$  in the state space  
 6 of the encoded system. Each assignment  $\alpha : Atoms_{[0,b]} \rightarrow \{true, false\}$  that  
 7 satisfies the formula characterises a definite path prefix, whereas an assignment  
 8 that makes the formula evaluate to *unknown* characterises a prefix with some  
 9  $\perp$ -transitions.

10 The second part of the encoding concerns the LTL property to be checked.  
 11 The three-valued bounded LTL encoding has been defined in [15] before. Here  
 12 we adjust it to our encodings of predicates and locations. Again, we distinguish  
 13 the cases where the property is evaluated on a path prefix with and without a  
 14 loop. The LTL encoding for the evaluation on prefixes with a loop is defined as:

**Definition 12 (LTL Encoding with Loop).**

Let  $p$  and  $(loc_i = l_i) \in AP$ ,  $\psi$  and  $\psi'$  LTL formulae, and  $b, k, r \in \mathbb{N}$  with  $k, r \leq b$  where  $k$  is the current position,  $b$  the bound and  $r$  the destination position of the  $b$ -loop. Then the LTL encoding with a loop,  $r\llbracket\psi\rrbracket_b^k$ , is defined as follows:

$$\begin{aligned} r\llbracket(loc_i = l_i)\rrbracket_b^k &\equiv enc(l_i)_k \\ r\llbracket\neg(loc_i = l_i)\rrbracket_b^k &\equiv \neg enc(l_i)_k \\ r\llbracket p \rrbracket_b^k &\equiv enc(p)_k \\ r\llbracket\neg p \rrbracket_b^k &\equiv enc(\neg p)_k \\ r\llbracket\psi \vee \psi'\rrbracket_b^k &\equiv r\llbracket\psi\rrbracket_b^k \vee r\llbracket\psi'\rrbracket_b^k \\ r\llbracket\psi \wedge \psi'\rrbracket_b^k &\equiv r\llbracket\psi\rrbracket_b^k \wedge r\llbracket\psi'\rrbracket_b^k \\ r\llbracket\mathbf{G}\psi\rrbracket_b^k &\equiv \bigwedge_{k'=\min(k,r)}^b r\llbracket\psi\rrbracket_b^{k'} \\ r\llbracket\mathbf{F}\psi\rrbracket_b^k &\equiv \bigvee_{k'=\min(k,r)}^b r\llbracket\psi\rrbracket_b^{k'} \\ r\llbracket\mathbf{X}\psi\rrbracket_b^k &\equiv r\llbracket\psi\rrbracket_b^{succ(k)} \end{aligned}$$

15 where  $succ(k) = k + 1$  if  $k < b$  and  $succ(k) = r$  else.

16 For a path prefix without a loop the LTL encoding is defined as:

1 **Definition 13 (LTL Encoding without Loop).**

2 Let  $\psi$  be an LTL formula and  $b, k \in \mathbb{N}$  with  $k \leq b$  where  $k$  is the current position  
 3 and  $b$  the bound. Then the LTL encoding without a loop,  $\llbracket \psi \rrbracket_b^k$ , is defined as  
 4 follows:

$$\begin{aligned} \llbracket \mathbf{G}\psi \rrbracket_b^k &\equiv \text{false} \\ \llbracket \mathbf{F}\psi \rrbracket_b^k &\equiv \bigvee_{k'=k}^b \llbracket \psi \rrbracket_b^{k'} \\ \llbracket \mathbf{X}\psi \rrbracket_b^k &\equiv \text{if } k < b \text{ then } {}_r\llbracket \psi \rrbracket_b^{k+1} \text{ else false} \end{aligned}$$

5 The LTL encoding without a loop of the other cases is identical to the LTL  
 6 encoding with a loop.

An example encoding is  $\llbracket \mathbf{F}p \rrbracket_2^0 = \text{enc}(p)_0 \vee \text{enc}(p)_1 \vee \text{enc}(p)_2$  which expresses that a predicate  $p$  holds *eventually*, i.e. at some position 0, 1 or 2 along a 2-prefix. Remember that a prefix  $\langle l^0, s^0 \rangle \dots \langle l^b, s^b \rangle$  has a  $b$ -loop if there exists a transition from  $\langle l^b, s^b \rangle$  to a previous state  $\langle l^r, s^r \rangle$  along the prefix with  $0 \leq r \leq b$ . Hence, we can define a loop constraint based on our transition encoding: A prefix characterised by an assignment  $\alpha_{\langle l^0, s^0 \rangle \dots \langle l^b, s^b \rangle}$  has definitely resp. maybe a  $b$ -loop if the loop constraint

$$\bigvee_{r=0}^b \text{Trans}_{b,r}$$

evaluates to *true* resp. *unknown* under  $\alpha_{\langle l^0, s^0 \rangle \dots \langle l^b, s^b \rangle}$  where  $\text{Trans}_{b,r}$  is defined according to Definition 11 but with  $k$  substituted by  $b$  and  $k + 1$  by  $r$ . This now allows us to define the overall encoding of whether a concurrent system  $Sys$  satisfies an LTL formula  $\psi$ :

$$\llbracket Sys, \psi \rrbracket_b := \llbracket Sys \rrbracket_b \wedge \llbracket \psi \rrbracket_b$$

with

$$\llbracket \psi \rrbracket_b := \llbracket \psi \rrbracket_b^0 \vee \bigvee_{r=0}^b ({}_{r}\llbracket \psi \rrbracket_b^0).$$

7 We have proven the following theorem that establishes the relation between  
 8 the satisfiability result for  $\llbracket Sys, \psi \rrbracket_b$  and the result of the corresponding model  
 9 checking problem:

**Theorem 1**

Let  $M$  be a three-valued Kripke structure representing the state space of an abstracted concurrent system  $Sys$ , let  $\psi$  be an LTL formula and  $b \in \mathbb{N}$  Then:

$$\llbracket M \models_{E,b} \psi \rrbracket \equiv \begin{cases} \text{true} & \text{if } \text{SAT}(\llbracket Sys, \psi \rrbracket_b^-) = \text{true} \\ \text{false} & \text{if } \text{SAT}(\llbracket Sys, \psi \rrbracket_b^+) = \text{false} \\ \perp & \text{else} \end{cases}$$

10 *Proof sketch.*

11 We have proven Theorem 1 by showing the following: (I) For each  $b$ -bounded

1 path  $\pi$  in  $M$  there exists an assignment  $\alpha_\pi : Atoms_{[0,b]} \rightarrow \{true, false\}$  that ex-  
 2 actly characterises  $\pi$  in  $\llbracket Sys \rrbracket_b$ , i.e. the transition values along  $\pi$  and  $\alpha_\pi(\llbracket Sys \rrbracket_b)$   
 3 are identical and the labellings along  $\pi$  and  $\alpha_\pi(\llbracket Sys \rrbracket_b)$  are identical as well. (II)  
 4 The evaluation of an LTL property  $\psi$  on  $\pi$  yields the same result as  $\alpha_\pi(\llbracket \psi \rrbracket_b)$ .  
 5 The full proof can be found in [16].  $\square$

6 Hence, via two satisfiability tests, one where  $\perp$  is substituted by *true* and one  
 7 where it is substituted by *false*, we can determine the result of the corresponding  
 8 model checking problem. Our encoding can be straightforwardly built based on  
 9 the concurrent system, which saves us the expensive construction of an explicit  
 10 state space model. In the next section we show that our encoding can be  
 11 also easily augmented by fairness constraints, which allows us to check liveness  
 12 properties of concurrent systems under realistic conditions.

## 13 5. Extension to Fairness

Our approach allows to check LTL properties of concurrent software sys-  
 tems via SAT solving. While the verification of safety properties like mutual  
 exclusion does not require any fairness assumptions about the behaviour of the  
 processes of the system, fairness is essential for verifying *liveness* properties un-  
 der realistic conditions. The most common notions of fairness in verification are  
*unconditional*, *weak* and *strong* fairness: An unconditional fairness constraint  
 claims that in an infinite computation, certain operations have to be infinitely  
 often executed. A weak fairness constraint claims that in an infinite compu-  
 tation, each operation that is *continuously* enabled has to be infinitely often  
 executed. A strong fairness constraint claims that in an infinite computation,  
 each operation that is *infinitely often* enabled has to be infinitely often executed.  
 All these types of constraints can be straightforwardly expressed in LTL. We  
 now define these constraints for characterising fair, i.e. realistic, behaviour of  
 our concurrent systems  $Sys = \parallel_{i=1}^n P_i$  over  $Pred$ . Our unconditional fairness  
 constraint is defined as:

$$ufair \equiv \bigwedge_{i=1}^n \bigvee_{(l_i, l'_i) \in \delta_i} \mathbf{GF}(executed(l_i, l'_i))$$

Hence, for each process some operation has to be executed infinitely often,  
 i.e. each process proceeds infinitely often. Note that we model termination via  
 a location with a self-loop. Thus, terminated processes can still proceed. The  
 expression  $executed(l_i, l'_i)$  can be easily defined in LTL. For this we extend the  
 set  $Pred$  by a progress predicate for each process:  $Pred := Pred \cup \{progress_i \mid i \in [1..n]\}$ .  
 Moreover, we extend each operation as follows:  $\tau_i(l_i, l'_i)$  sets  $progress_i$   
 to *true* and all  $progress_{i'}$  with  $i' \neq i$  to *false*. Now  $executed(l_i, l'_i)$  is defined as

$$executed(l_i, l'_i) \equiv (loc_i = l_i) \wedge \mathbf{X}((loc_i = l'_i) \wedge progress_i).$$

14 where  $\mathbf{X}$  is the LTL operator ‘next’. An example for a system where operations  
 15 are extended with *progress* statements is given by the parallel composition of

- 1 control flow graphs depicted in Figure 8. As we can see, each time when a
- 2 process executes an operation, it sets its own progress predicate to *true* and the
- 3 progress predicate of the other process to *false*.

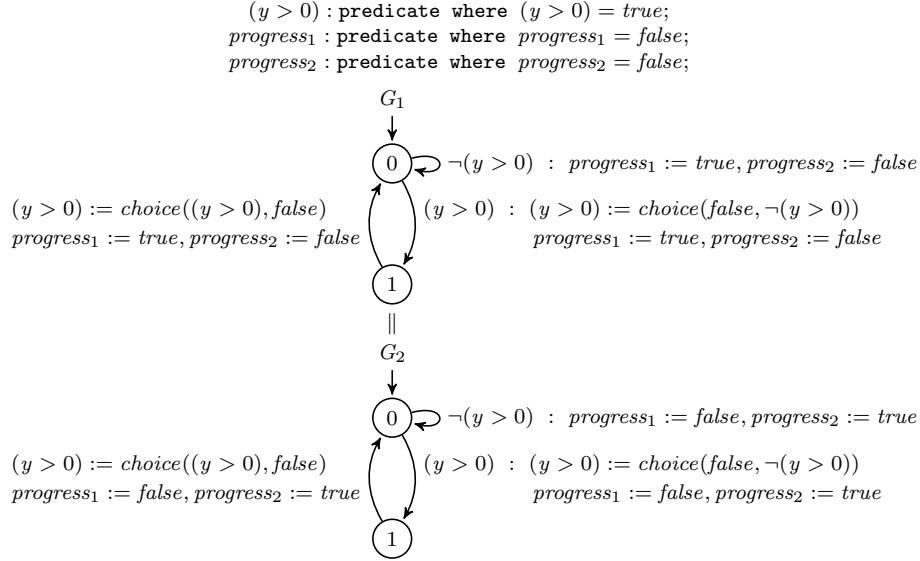


Figure 8: Parallel composition of abstract control flow graphs where operations are extended with progress statements.

An operation associated with a control flow transition  $(l_i, l'_i)$  is executed if  $(loc_i = l_i)$  holds in the current state and  $(loc_i = l'_i) \wedge \text{progress}_i$  holds in the next state. For the control flow graph  $G_1$  in Figure 8 we have for instance  $\text{executed}(0, 1) \equiv (loc_1 = 0) \wedge \mathbf{X}((loc_1 = 1) \wedge \text{progress}_1)$ . Next, we define our weak fairness constraint:

$$\text{wfair} \equiv \bigwedge_{i=1}^n \bigwedge_{(l_i, l'_i) \in \delta_i} (\mathbf{FG}(\text{enabled}(l_i, l'_i)) \rightarrow \mathbf{GF}(\text{executed}(l_i, l'_i)))$$

Hence, for each process, each continuously enabled operation has to be infinitely often executed. Instead of incorporating *each* operation in this type of constraint it is also possible to restrict the operations to crucial ones, which results in a shorter constraint and thus also restrains the complexity of model checking under fairness. For our running example it is for instance appropriate to just incorporate operations in *wfair* that correspond to the successful acquisition of the semaphore. Note that *wfair* can be easily transferred into negation normal form via the common propositional logic transformation rules such that it is conform with the definition of LTL. The expression  $\text{enabled}(l_i, l'_i)$  can be defined as an LTL formula over locations and *Pred* as follows:

$$\text{enabled}(l_i, l'_i) \equiv (loc_i = l_i) \wedge \text{choice}(a, b)$$

assuming that  $\tau_i(l_i, l'_i) = \text{assume}(\text{choice}(a, b)) : p_1 := \text{choice}(a_1, b_1), \dots, p_m := \text{choice}(a_m, b_m)$ . Thus, an operation associated with a control flow transition  $(l_i, l'_i)$  is enabled if  $(loc_i = l_i)$  holds and the guard of the operation holds as well. An example of an *enabled* expression in terms of the control flow graph  $G_1$  in Figure 8 is  $\text{enabled}(0, 1) \equiv (loc_1 = 0) \wedge (y > 0)$ . Finally, we define our strong fairness constraint:

$$\text{sfair} \equiv \bigwedge_{i=1}^n \bigwedge_{(l_i, l'_i) \in \delta_i} (\mathbf{GF}(\text{enabled}(l_i, l'_i)) \rightarrow \mathbf{GF}(\text{executed}(l_i, l'_i)))$$

Hence, for each process, each operation that is enabled infinitely often has to be executed infinitely often. In model checking under fairness we can either check properties under specific constraints or we can combine all to a general one

$$\text{fair} \equiv \text{ufair} \wedge \text{wfair} \wedge \text{sfair}.$$

Existential bounded model checking under fairness is now defined as:

$$[M \models_{E,b}^{\text{fair}} \psi] \equiv [M \models_{E,b} (\text{fair} \wedge \psi)]$$

Thus, we check whether there exists a  $b$ -bounded path that is fair and satisfies the property  $\psi$ . Such a model checking problem can be straightforwardly encoded in propositional logic based on our definitions in the previous section. We get

$$\llbracket \text{Sys}, \text{fair} \wedge \psi \rrbracket_b := \llbracket \text{Sys} \rrbracket_b \wedge \llbracket \text{fair} \wedge \psi \rrbracket_b,$$

- 1 which can be fed into a SAT solver in order to obtain the result of model checking
- 2  $\psi$  under fairness. Next, we introduce a systematic and fully-automatic approach
- 3 to the refinement of three-valued abstractions in case the corresponding three-
- 4 valued model checking problem yields *unknown*.

## 5 6. Cause-Guided Abstraction Refinement

6 In this section we present our approach to the refinement of three-valued ab-  
7 stractions in case the corresponding model checking problem yields an *unknown*  
8 result. Our abstractions represent uncertainty in the form of the constant  $\perp$ .  
9 SAT-based three-valued model checking is performed via two satisfiability tests,  
10 one where all occurrences of  $\perp$  are mapped to *true* in the propositional logic  
11 encoding  $\llbracket \text{Sys}, \psi \rrbracket_b$  and one where its occurrences are mapped to *false*. Here we  
12 introduce an enhanced encoding that comprises the *causes of uncertainty*: Each  
13  $\perp$  in the encoding gets superscripted with a *cause*, which can be missing infor-  
14 mation about a transition or a predicate. During the satisfiability tests all  $\perp$ 's  
15 are still treated the same, meaning that either all of them are mapped to *true* or  
16 all to *false* (compare Theorem 1). Once we have obtained an overall *unknown*  
17 model checking result, i.e.  $\text{SAT}(\llbracket \text{Sys}, \psi \rrbracket_b^+) = \text{true}$  for an assignment  $\alpha$  and

1 SAT( $\llbracket Sys, \psi \rrbracket_b^-$ ) = *false*, we proceed as follows: We now have that the assign-  
 2 ment  $\alpha$  characterises an unconfirmed witness path for  $\psi$  containing *unknowns*.  
 3 Thus, this path is not present if all  $\perp$ 's get mapped to *false*. We determine the  
 4 unsatisfied clauses of  $\alpha(\llbracket Sys, \psi \rrbracket_b^-)$ . All these clauses contain uncertainty in the  
 5 sense of  $\perp$ 's and we will see that we can straightforwardly derive the correspond-  
 6 ing causes. We then apply our novel *cause-guided abstraction refinement* which  
 7 rules out the causes of uncertainty by adding new predicates to the abstraction.  
 8 We will show that our fully-automatic iterative refinement approach enables us  
 9 to quickly reach the right level of abstraction in order to obtain a definite model  
 10 checking result.

11 The basis of our refinement technique is an enhanced encoding comprising  
 12 causes of uncertainty:

**Definition 14 (Causes of Uncertainty).**

Let  $\llbracket Sys, \psi \rrbracket_b$  be the propositional logic encoding of a three-valued bounded model  
 checking problem corresponding to a concurrent system  $Sys = \parallel_{i=1}^n P_i$  ab-  
 stracted over *Pred* where each process is given by a single control flow graph  
 $G_i = (Loc_i, \delta_i, \tau_i)$  with  $1 \leq i \leq n$ . Uncertainty is represented in the encoding by  
 the constant  $\perp$ . Each  $\perp$  in the encoding can be associated with a cause which  
 we define as follows:

$$cause \in \{p_k, (l_i, l'_i)_k\}$$

13 with  $p \in Pred$ ,  $0 \leq k \leq b$  and  $l_i, l'_i \in Loc_i$ .

14 We will use  $p_k$  in order to denote that the predicate  $p$  potentially evaluates  
 15 to *unknown* at position  $k$  of the encoding, and with  $(l_i, l'_i)_k$  we will denote that  
 16 missing predicates over the guard of the operation  $\tau_i(l_i, l'_i)$  potentially cause an  
 17 *unknown* transition from position  $k$  to  $k + 1$  in the encoding. Note that we  
 18 refer to *potential* uncertainty in an encoding  $\llbracket Sys, \psi \rrbracket_b$ , since  $\llbracket Sys, \psi \rrbracket_b$  always  
 19 characterises *many* possible execution paths. For a *specific* path characterised  
 20 by an assignment  $\alpha$  to the atoms of  $\llbracket Sys, \psi \rrbracket_b$  will see that we can refer to *actual*  
 21 uncertainty. Next we show how causes of uncertainty can be integrated into the  
 22 encoding in the sense of adding them as superscripts to the  $\perp$ 's. For this we  
 23 introduce an enhanced encoding of abstract operations:

**Definition 15 (Enhanced Encoding of Operations).**

Let  $Sys = \parallel_{i=1}^n P_i$  over *Pred* be an abstracted concurrent system given by  
 the single control flow graphs  $G_i = (Loc_i, \delta_i, \tau_i)$  with  $1 \leq i \leq n$ . Then the  
 encoding of abstract operations  $\tau_i(l_i, l'_i) = \text{assume}(\text{choice}(a, b)) : p_1 :=$   
 $\text{choice}(a_1, b_1), \dots, p_m := \text{choice}(a_m, b_m)$  comprising the causes of uncertainty is  
 defined as follows:

$$\begin{aligned}
 enc(\tau_i(l_i, l'_i))_{k, k+1} := & \quad enc(\text{choice}(a, b))_k \\
 & \wedge \bigwedge_{j=1}^m ( \quad (enc(a_j)_k \wedge enc(p_j = \text{true})_{k+1}) \\
 & \quad \vee (enc(b_j)_k \wedge enc(p_j = \text{false})_{k+1}) \\
 & \quad \vee (enc(\neg a_j \wedge \neg b_j)_k [\perp \mapsto \text{true}] \wedge enc(p_j = \perp)_{k+1}) )
 \end{aligned}$$

with

$$enc(choice(a, b))_k := enc((a \vee NNF(\neg b)) \wedge (a \vee b \vee \perp^{(\mathbf{l}, \mathbf{l}')_k}))_k$$

and the following inductive definition of the encoding of logical expressions  $e, e'$  over predicates  $p \in Pred$ .

$$\begin{aligned} enc(p)_k &:= (p[u]_k \wedge \perp^{\mathbf{p}_k}) \vee (\neg p[u]_k \wedge p[t]_k) \\ enc(\neg p)_k &:= (p[u]_k \wedge \perp^{\mathbf{p}_k}) \vee (\neg p[u]_k \wedge \neg p[t]_k) \\ enc(e \wedge e')_k &:= enc(e)_k \wedge enc(e')_k \\ enc(e \vee e')_k &:= enc(e)_k \vee enc(e')_k \end{aligned}$$

This definition enhances our previous encoding Definitions 9 and 10 in terms of superscripting each  $\perp$  with a corresponding cause. We will ignore the causes and treat all  $\perp$ 's the same during the satisfiability checks. Hence, the enhanced encoding is equivalent to the standard encoding. However, in case of an *unknown* model checking result, the causes will become crucial and will allow us to immediately derive expedient refinement steps. For illustration, we encode the following abstract operation:

$$\tau(l_i, l'_i) = assume(choice(false, false)) : p := choice(p, \neg p)$$

Remember that  $choice(false, false) \equiv \perp$ . For the enhanced encoding we get:

$$\begin{aligned} enc(\tau(l_i, l'_i))_{k, k+1} = \\ \perp^{(\mathbf{l}, \mathbf{l}')_k} \wedge (((p[u]_k \wedge \perp^{\mathbf{p}_k}) \vee (\neg p[u]_k \wedge p[t]_k)) \wedge \neg p[u]_{k+1} \wedge p[t]_{k+1}) \\ \vee (((p[u]_k \wedge \perp^{\mathbf{p}_k}) \vee (\neg p[u]_k \wedge \neg p[t]_k)) \wedge \neg p[u]_{k+1} \wedge \neg p[t]_{k+1}) \end{aligned}$$

1 Thus, uncertainty may be caused by the unknown guard of the abstract  
 2 operation  $\tau(l_i, l'_i)$  or by the predicate  $p$  evaluating to  $\perp$  at position  $k$ . Actual  
 3 uncertainty along a path characterised by an assignment  $\alpha$  is only present if  
 4 the  $\perp$ 's occur in clauses unsatisfied under  $\alpha[\perp \mapsto false]$ . Then we can utilise  
 5 the causes attached to the  $\perp$ 's in order to rule out the uncertainty. We will  
 6 now introduce our iterative abstraction-based model checking procedure with  
 7 cause-guided refinement:

8 **Procedure 1 (Iterative Abstraction-Based Model Checking).**

9 Let  $G = (Loc, \delta, \tau)$  be the concrete control flow graph representing a concurrent  
 10 system  $Sys$  defined over a set of variables  $Var$ . Moreover, let  $\psi$  be an LTL  
 11 formula to be checked for  $Sys$ . The corresponding bounded model checking  
 12 problem can be solved via three-valued abstraction refinement and satisfiability  
 13 solving as follows:

- 14 1. Initialise the set of predicates  $Pred$  with the atomic propositions over  $Var$   
 15 referenced in  $\psi$ . Initialise the bound  $b$  with 1.
- 16 2. Construct the abstract control flow graph  $G_a = (Loc_a, \delta_a, \tau_a)$  representing  
 17  $Sys$  abstracted over the current set  $Pred$  via a three-valued abstractor [2].

- 1     3. Encode the three-valued bounded model checking problem  $[M(G_a) \models_{E,b} \psi]$
- 2          $\psi]$  for the current bound  $b$  in propositional logic, which yields the formula
- 3          $\llbracket Sys(G_a), \psi \rrbracket_b$ .
- 4     4. Apply SAT-based three-valued bounded model checking (according to
- 5         Theorem 1):
- 6         (a) If the result is  $[M(G_a) \models_{E,b} \psi] = true$ , then there exists a  $b$ -bounded
- 7             witness path for  $\psi$  in the state space of  $Sys$ . The path is characterised
- 8             by an assignment  $\alpha$  satisfying  $\llbracket Sys(G_a), \psi \rrbracket_b^-$ . Return  $\alpha$ .
- 9         (b) If the result is  $[M(G_a) \models_{E,b} \psi] = false$ , then there does not exist a
- 10              $b$ -bounded witness path for  $\psi$  in the state space of  $Sys$ . Terminate
- 11             if  $b$  has reached the completeness threshold [14] of the verification
- 12             task. Otherwise increment  $b$  and go to 3.
- 13         (c) If the result is  $[M(G_a) \models_{E,b} \psi] = \perp$ , then it is *unknown* whether
- 14             there exists a  $b$ -bounded witness path for  $\psi$  in the state space of
- 15              $Sys$ . An unconfirmed witness path for  $\psi$  with *unknowns* is charac-
- 16             terised by an assignment  $\alpha$  satisfying  $\llbracket Sys(G_a), \psi \rrbracket_b^+$  but not satisfy-
- 17             ing  $\llbracket Sys(G_a), \psi \rrbracket_b^-$ . Apply the procedure *Cause-Guided Abstraction*
- 18             *Refinement*, which updates  $Pred$ , and go to 2.

19 **Procedure 2 (Cause-Guided Abstraction Refinement).**

20 Let  $G = (Loc, \delta, \tau)$  be the concrete composite control flow graph representing a  
 21 concurrent system  $Sys = \parallel_{i=1}^n P_i$  defined over a set of variables  $Var$  and let  $G_i =$   
 22  $(Loc_i, \delta_i, \tau_i)$  with  $1 \leq i \leq n$  the corresponding single CFGs. Let  $\psi$  be an LTL  
 23 formula to be checked for  $Sys$ . Moreover, let  $G_a = (Loc_a, \delta_a, \tau_a)$  be the abstract  
 24 composite control flow graph representing  $Sys$  abstracted over a set of predicates  
 25  $Pred$  and let  $[M(G_a) \models_{E,b} \psi]$  be the corresponding three-valued bounded model  
 26 checking problem with  $[M(G_a) \models_{E,b} \psi] = \perp$ . Then for the propositional logic  
 27 encoding  $\llbracket Sys(G_a), \psi \rrbracket_b$  the following holds:  $SAT(\llbracket Sys(G_a), \psi \rrbracket_b^+) = true$  and  
 28 the solver additionally returns a corresponding satisfying assignment  $\alpha$  charac-  
 29 terising an unconfirmed witness path.  $SAT(\llbracket Sys(G_a), \psi \rrbracket_b^-) = false$ . Now the  
 30 abstraction can be automatically refined by updating  $Pred$  as follows:

- 31     1. Determine the set  $U$  of clauses of  $\llbracket Sys, \psi \rrbracket_b$  that are unsatisfied under the
- 32         assignment  $\alpha[\perp \mapsto false]$ . (Each  $u \in U$  must contain at least one  $\perp$  since
- 33         we have that  $SAT(\llbracket Sys(G_a), \psi \rrbracket_b^+) = true$ .)
- 34     2. Determine a set  $Causes$  such that for each  $u \in U$  there exists a *cause*  $\in$
- 35          $Causes$  with  $\perp^{cause}$  is contained in  $u$ .
- 36     3. For each *cause*  $\in Causes$ :
- 37         (a) If *cause*  $= (l_i, l'_i)_k$  with  $l_i, l'_i \in Loc_i$ ,  $1 \leq i \leq n$  and  $0 \leq k \leq b$ ,
- 38             then the value of the  $k$ -th transition along the unconfirmed witness
- 39             path characterised by  $\alpha$  is *unknown*. The transition from  $k$  to  $k + 1$
- 40             is associated with an operation  $\tau_i(l_i, l'_i) = assume(e) : v_1 :=$
- 41              $e_1, \dots, v_m := e_m$  of the concrete control flow graph  $G_i$ .  $\tau_i(l_i, l'_i)$  can
- 42             be straightforwardly derived from  $G_i$ . Add the atomic propositions
- 43             occurring in the assume condition  $e$  as new predicates to  $Pred$ .



1 (b) If  $cause = p_k$  with  $p \in Pred$  and  $0 \leq k \leq b$ , then the value of  $p$  is  
2 *unknown* at position  $k$  of the unconfirmed witness path characterised  
3 by  $\alpha$ , i.e.  $\alpha(p[u]_k) = true$ . Let  $k' < k$  be the last predecessor of  $k$   
4 with  $\alpha(p[u]_{k'}) = false$ , i.e. the last position where the value of  $p$  is  
5 known. The transition from position  $k'$  to  $k' + 1$  is associated with  
6 an operation  $\tau_i(l_i, l'_i) = assume(e) : v_1 := e_1, \dots, v_m := e_m$  of a con-  
7 crete control flow graph  $G_i$ . Missing information about this concrete  
8 operation in terms of predicates is the cause of uncertainty in the  
9 current abstraction.  $\tau_i(l_i, l'_i)$  can be straightforwardly derived based  
10 on  $G_i$  and  $\alpha(LocAtoms_{[k', k'+1]})$ , which indicates the corresponding  
11 control flow locations. Let  $wp_{\tau_i(l_i, l'_i)}(p) = p[v_1 \leftarrow e_1, \dots, v_m \leftarrow e_m]$   
12 be the weakest precondition<sup>1</sup> of  $p$  with respect to the assignment part  
13 of the operation  $\tau_i(l_i, l'_i)$ . Add the atomic propositions occurring in  
14  $wp_{\tau_i(l_i, l'_i)}(p)$  as new predicates to  $Pred$ .

15 Our procedure refines the three-valued abstraction by adding further predi-  
16 cates to the set  $Pred$ . According to Corollary 1, such a refinement is sound in  
17 the sense that it preserves the validity of definite temporal logic properties.

18 We now exemplify our iterative abstraction refinement approach based on  
19 the simple system  $Sys$  and the corresponding concrete control flow graph  $G_c$   
20 depicted in Figure 9.




Figure 9: Concurrent system  $Sys$  and corresponding concrete control flow graph  $G_c$ .

21 Here we have a single process operating on the integer variable  $y$  and we  
22 want to check whether there exists an execution that finally reaches control flow  
23 location 1. Thus, the temporal logic property of interest is  $\mathbf{F}(loc = 1)$ . In the  
24 first iteration, we start with bound  $b = 1$  and  $Pred = \emptyset$ . The corresponding  
25 abstract control flow graph  $G_{a_1}$ , computable with a three-valued abstractor, is  
26 depicted in Figure 10.

In order to solve the corresponding three-valued bounded model checking  
problem  $[M(G_{a_1}) \models_{E,1} \mathbf{F}(loc = 1)]$  we construct the propositional logic encod-

<sup>1</sup>Computable via an SMT solver with built-in linear integer arithmetic theory. In our approach we use Z3 [17].

no predicates




Figure 10: Abstract control flow graph  $G_{a_1}$ .

ing that now comprises the causes of uncertainty:

$$\begin{aligned} & \llbracket Sys(G_{a_1}), \mathbf{F}(loc = 1) \rrbracket_1 = \\ & \underbrace{(\neg l_0)}_{Init_0} \wedge \underbrace{((\neg l_0 \wedge \neg l_1 \wedge \perp^{(0,0)_0}) \vee (\neg l_0 \wedge l_1 \wedge \perp^{(0,1)_0}) \vee (l_0 \wedge l_1))}_{Trans_{0,1}} \wedge \underbrace{(l_0 \vee l_1)}_{\llbracket \mathbf{F}(loc=1) \rrbracket_1^0} \end{aligned}$$

Since our system consists of a single process and only one digit is necessary to encode the control flow, we can omit the process- and digit-indices of the atoms. Solely the position index is required for the encoding. Now we run the associated satisfiability tests. We get

$$\text{SAT}(\llbracket Sys(G_{a_1}), \mathbf{F}(loc = 1) \rrbracket_1^+) = true$$

and the corresponding satisfying truth assignment  $\alpha : l_0 \mapsto false, l_1 \mapsto true$ . Moreover, we get

$$\text{SAT}(\llbracket Sys(G_{a_1}), \mathbf{F}(loc = 1) \rrbracket_1^-) = false$$

Hence,  $[M(G_{a_1}) \models_{E,1} \mathbf{F}(loc = 1)]$  yields *unknown* and  $\alpha$  characterises an unconfirmed witness path

$$\langle 0 \rangle \xrightarrow{\perp} \langle 0 \rangle$$

in the abstract state space where  $\xrightarrow{\perp}$  denotes an *unknown* transition between the states. Next we apply the procedure *Cause-Guided Abstraction Refinement*. Remember that SAT solvers always operate on formulae transferred into conjunctive normal form (CNF). The current encoding is equivalent to the following

formula in CNF<sup>2</sup>

$$\begin{aligned} & (\neg l_0) \wedge (l_0 \vee \neg l_1 \vee \perp^{(\mathbf{0},\mathbf{1})\mathbf{o}}) \wedge (l_1 \vee \perp^{(\mathbf{0},\mathbf{0})\mathbf{o}}) \\ & \wedge (l_0 \vee \perp^{(\mathbf{0},\mathbf{0})\mathbf{o}} \vee \perp^{(\mathbf{0},\mathbf{1})\mathbf{o}}) \wedge (l_1 \vee \perp^{(\mathbf{0},\mathbf{0})\mathbf{o}} \vee \perp^{(\mathbf{0},\mathbf{1})\mathbf{o}}) \wedge (l_0 \vee l_1) \end{aligned}$$

Under the assignment  $\alpha[\perp \mapsto \text{false}]$  (the assignment  $\alpha$  extended by the assignment that maps all  $\perp$ 's to *false*) we get the following set of unsatisfied clauses for our encoding:

$$U = \{(l_0 \vee \neg l_1 \vee \perp^{(\mathbf{0},\mathbf{1})\mathbf{o}}), (l_0 \vee \perp^{(\mathbf{0},\mathbf{0})\mathbf{o}} \vee \perp^{(\mathbf{0},\mathbf{1})\mathbf{o}})\}$$

A corresponding set of causes of uncertainty that covers  $U$  is

$$\text{Causes} = \{(0, 1)_0\}$$

1 since  $\perp^{(\mathbf{0},\mathbf{1})\mathbf{o}}$  occurs in all clauses of  $U$ .  $(0, 1)_0$  indicates that at the current  
 2 level of abstraction uncertainty is caused by the missing guard of the operation  
 3  $\tau(0, 1)$ . We have that  $\tau(0, 1) = \neg(y > 0)$  in the concrete system. Hence, we add  
 4  $(y > 0)$  to the set of predicates:  $\text{Pred} := \text{Pred} \cup \{(y > 0)\}$  and proceed with the  
 5 next iteration.

6 In the second iteration, we have  $b = 1$  and  $\text{Pred} = \{(y > 0)\}$ . The corresponding abstract control flow graph  $G_{a_2}$  is depicted in Figure 11.

$(y > 0) : \text{predicate where } (y > 0) = \text{true};$




Figure 11: Abstract control flow graph  $G_{a_2}$ .

7

In order to solve  $[M(G_{a_2}) \models_{E,1} \mathbf{F}(\text{loc} = 1)]$  we construct the following encoding  $((y > 0)$  abbreviated by  $p$ ):

$$\begin{aligned} & \llbracket \text{Sys}(G_{a_2}), \mathbf{F}(\text{loc} = 1) \rrbracket_1 = \\ & (\neg l_0 \wedge \neg p[u]_0 \wedge p[t]_0) \wedge ((\neg l_0 \wedge \neg l_1 \wedge \text{enc}(\tau(0, 0))_{0,1}) \\ & \vee (\neg l_0 \wedge l_1 \wedge \text{enc}(\tau(0, 1))_{0,1}) \vee (l_0 \wedge l_1 \wedge \text{enc}(\tau(1, 1))_{0,1})) \wedge (l_0 \vee l_1) \end{aligned}$$

<sup>2</sup>For the sake of simplicity we use a standard CNF transformation in this illustrating example. Note that in our implementation we use the more compact Tseitin CNF transformation which introduces additional auxiliary atoms. Hence, we would get a slightly different CNF formula and unsatisfied clauses. Nevertheless, these clauses would hint at exactly the same causes of uncertainty.

with

$$\begin{aligned}
& enc(\tau(0, 0))_{0,1} = \\
& ((p[u]_0 \wedge \perp^{\mathbf{p}0}) \vee (\neg p[u]_0 \wedge p[t]_0)) \\
& \wedge (((p[u]_0 \wedge \perp^{\mathbf{p}0}) \vee (\neg p[u]_0 \wedge \neg p[t]_0)) \wedge (\neg p[u]_1 \wedge \neg p[t]_1)) \\
& \vee ((p[u]_0 \vee (\neg p[u]_0 \wedge p[t]_0)) \wedge (p[u]_1))
\end{aligned}$$

and  $\tau(0, 1)_{0,1}$  and  $\tau(1, 1)_{0,1}$  encoded analogously. As we can see, uncertainty is now potentially caused by predicate  $p$  evaluating to  $\perp$  at position 0. Now we run the associated satisfiability tests. We get

$$\text{SAT}(\llbracket Sys(G_{a_2}), \mathbf{F}(loc = 1) \rrbracket_1^+) = false$$

and

$$\text{SAT}(\llbracket Sys(G_{a_2}), \mathbf{F}(loc = 1) \rrbracket_1^-) = false$$

1 Hence,  $[M(G_{a_2}) \models_{E,1} \mathbf{F}(loc = 1)]$  yields *false*, which indicates that there  
2 does not exist a 1-bounded witness path for  $\mathbf{F}(loc = 1)$ . Consequently, we  
3 increment the bound:  $b := b + 1$  and proceed with the next iteration.

In the third iteration, we have  $b = 2$  and still  $Pred = \{(y > 0)\}$ . Hence, we continue with the abstract the control flow graph  $G_{a_2}$ . In order to solve  $[M(G_{a_2}) \models_{E,2} \mathbf{F}(loc = 1)]$  we construct the following encoding:

$$\begin{aligned}
& \llbracket Sys(G_{a_2}), \mathbf{F}(loc = 1) \rrbracket_2 = \\
& (\neg l_0 \wedge \neg p[u]_0 \wedge p[t]_0) \wedge \bigwedge_{k=0}^1 ((\neg l_k \wedge \neg l_{k+1} \wedge enc(\tau(0, 0))_{k,k+1}) \\
& \vee (\neg l_k \wedge l_{k+1} \wedge enc(\tau(0, 1))_{k,k+1}) \vee (l_k \wedge l_{k+1} \wedge enc(\tau(1, 1))_{k,k+1})) \\
& \wedge (l_0 \vee l_1 \vee l_2)
\end{aligned}$$

Now we run the associated satisfiability tests. We get

$$\text{SAT}(\llbracket Sys(G_{a_2}), \mathbf{F}(loc = 1) \rrbracket_2^+) = true$$

and the corresponding satisfying truth assignment  $\alpha$  :

$$\begin{aligned}
& l_0 \mapsto false, \quad l_1 \mapsto false, \quad l_2 \mapsto true, \\
& p[u]_0 \mapsto false, \quad p[t]_0 \mapsto true, \\
& p[u]_1 \mapsto true, \quad p[t]_1 \mapsto true, \\
& p[u]_2 \mapsto true, \quad p[t]_2 \mapsto true.
\end{aligned}$$

Moreover, we get

$$\text{SAT}(\llbracket Sys(G_{a_2}), \mathbf{F}(loc = 1) \rrbracket_2^-) = false$$

Hence,  $[M(G_{a_2}) \models_{E,2} \mathbf{F}(loc = 1)]$  yields *unknown* and  $\alpha$  characterises an

unconfirmed witness path

$$\langle 0, p = true \rangle \rightarrow \langle 0, p = \perp \rangle \xrightarrow{\perp} \langle 1, p = \perp \rangle$$

in the abstract state space. Next we apply the procedure *Cause-Guided Abstraction Refinement*. After deriving the set  $U$  of clauses of  $\llbracket Sys(G_{a_2}), \mathbf{F}(loc = 1) \rrbracket_2$  that are unsatisfied under the assignment  $\alpha[\perp \mapsto false]$ , we determine a corresponding set of causes covering  $U$ . We get:

$$Causes = \{p_1\}$$

$p_1$  indicates that at the current level of abstraction uncertainty is caused by the predicate  $p$  evaluating to *unknown* at position 1 of the witness path characterised by  $\alpha$ . Now we determine the last predecessor position where the value of  $p$  is known, that is the greatest  $k < 1$  with  $\alpha(p[u]_k) = false$ . This holds for  $k = 0$ . Hence, the transition from position 0 to 1 along the witness path characterised by  $\alpha$  makes  $p$  *unknown*. We have that  $\alpha(l_0) = false$  and  $\alpha(l_1) = false$ , which indicates that the transition from position 0 to 1 is associated with the operation  $\tau(0,0)$ . From the concrete control flow graph we get that the assignment part of  $\tau(0,0)$  is  $y := y - 1$ . Thus, the weakest precondition of  $p = (y > 0)$  with respect to  $\tau(0,0)$  is

$$wp_{\tau(0,0)}(y > 0) = (y > 0)[y \leftarrow y - 1] = (y - 1 > 0) = (y > 1).$$

1 Hence, we add  $(y > 1)$  to the set of predicates and proceed with the next  
2 iteration.

3 In the forth iteration, we have  $b = 2$  and  $Pred = \{(y > 0), (y > 1)\}$ . The  
4 corresponding abstract control flow graph  $G_{a_3}$  is depicted in Figure 12.

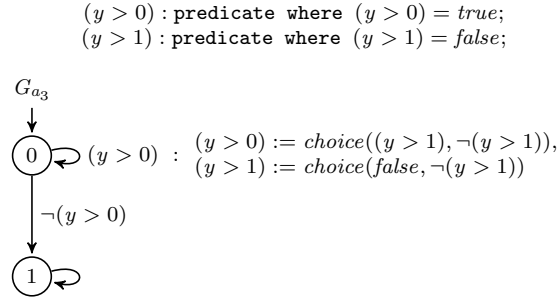


Figure 12: Abstract control flow graph  $G_{a_3}$ .

In order to solve  $[M(G_{a_3}) \models_{E,2} \mathbf{F}(loc = 1)]$  we construct the encoding  $\llbracket Sys(G_{a_3}), \mathbf{F}(loc = 1) \rrbracket_2$  and run the associated satisfiability tests. We get

$$\text{SAT}(\llbracket Sys(G_{a_3}), \mathbf{F}(loc = 1) \rrbracket_2^+) = true$$

and

$$\text{SAT}(\llbracket \text{Sys}(G_{a_3}), \mathbf{F}(\text{loc} = 1) \rrbracket_2^-) = \text{true}$$

and the corresponding satisfying truth assignment  $\alpha$  :

$$\begin{aligned} l_0 &\mapsto \text{false}, & l_1 &\mapsto \text{false}, & l_2 &\mapsto \text{true}, \\ p[u]_0 &\mapsto \text{false}, & p[t]_0 &\mapsto \text{true}, \\ p[u]_1 &\mapsto \text{false}, & p[t]_1 &\mapsto \text{false}, \\ p[u]_2 &\mapsto \text{false}, & p[t]_2 &\mapsto \text{false}, \\ q[u]_0 &\mapsto \text{false}, & q[t]_0 &\mapsto \text{false}, \\ q[u]_1 &\mapsto \text{true}, & q[t]_1 &\mapsto \text{false}, \\ q[u]_2 &\mapsto \text{true}, & q[t]_2 &\mapsto \text{false} \end{aligned}$$

where  $p$  abbreviates  $(y > 0)$  and  $q$  abbreviates  $(y > 1)$ . We can immediately conclude that  $[M(G_{a_3}) \models_{E,2} \mathbf{F}(\text{loc} = 1)]$  yields *true*, which indicates that  $\alpha$  characterises a definite 2-bounded witness path

$$\langle 0, p = \text{true}, q = \text{false} \rangle \rightarrow \langle 0, p = \text{false}, q = \perp \rangle \rightarrow \langle 1, p = \text{false}, q = \perp \rangle$$

1 for  $\mathbf{F}(\text{loc} = 1)$ . This outcome completes our verification task. Within four  
 2 iterations of cause-guided abstraction refinement resp. bound incrementation  
 3 we have automatically proven that the property of interest holds for the system.

4 Thus, given a software verification task to be solved, our cause-guided refine-  
 5 ment approach enables us to automatically reach the right level of abstraction  
 6 in order to obtain a definite result in verification. Next, we present the imple-  
 7 mentation of our encoding-based model checking technique and we report on  
 8 experimental results.

## 9 7. Implementation, Enhancements and Experiments

10 In this section we introduce the implementation of our theoretical concepts  
 11 and we present enhancements based on existing work on *temporal induction*  
 12 [10] as well as on *symmetry-based parameterised verification* [11]. Moreover, we  
 13 discuss several experimental results.

### 14 7.1. The TVMC Tool

15 We have implemented a SAT-based bounded model checker called TVMC  
 16 for three-valued abstractions of concurrent software systems.<sup>3</sup> TVMC employs a  
 17 three-valued abstractor [2] that builds abstract control flow graphs for a given  
 18 concurrent system *Sys* and a set of predicates *Pred*. It supports almost all  
 19 control structures of the C language as well as *int*, *bool* and *semaphore* as data

---

<sup>3</sup>available at [www.github.com/ssfm-up/TVMC](http://www.github.com/ssfm-up/TVMC)

1 types. Based on the CFGs and an input LTL formula  $\psi$ , our tool automatically  
 2 constructs an encoding  $\llbracket Sys, \psi \rrbracket_b$  of the corresponding verification task. TVMC  
 3 iteratively refines the abstraction in case of an *unknown* result and increments  
 4 the bound in case of a *false* result. It terminates once a *true* result is obtained  
 5 or a *false* result is obtained for a predefined threshold of the bound: In each  
 6 iteration the two instances of the encoding are processed by a solver thread  
 7 of the SAT solver Sat4j [18]. A *true* result for  $\llbracket Sys, \psi \rrbracket_b^-$  can be immediately  
 8 transferred to the corresponding model checking problem  $[M \models_{E,b} \psi]$ . The same  
 9 holds for a *false* result for  $\llbracket Sys, \psi \rrbracket_b^+$  if  $b$  represents a completeness threshold of  
 10 the verification task [14]. In case of an *unknown* result we apply cause-guided  
 11 abstraction refinement as defined in the previous section. For *true* and *unknown*  
 12 results, we additionally output a definite resp. unconfirmed witness path for the  
 13 property  $\psi$  in the form of an assignment satisfying  $\llbracket Sys, \psi \rrbracket_b$ . The tool chain of  
 14 TVMC is depicted in Figure 13.

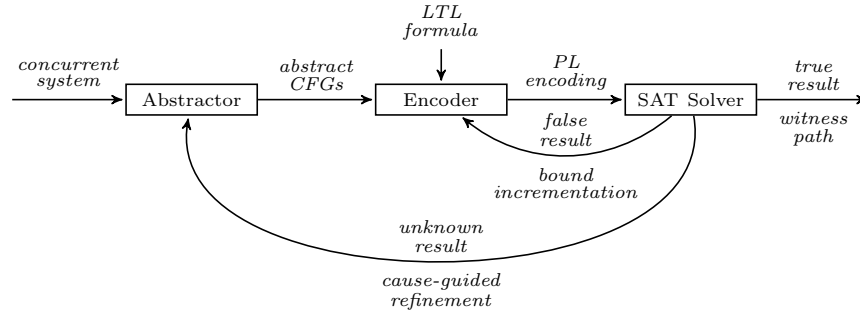


Figure 13: TVMC tool chain.

15 We now illustrate how our tool systematically solves verification tasks via  
 16 three-valued abstraction and cause-guided refinement. An example system  
 17  $Sys = \parallel_{i=1}^n P_i$  implementing a solution to the dining philosophers problem  
 18 is depicted in Figure 14. Here we have  $n \in \mathbb{N}$  philosopher processes and the  
 19 same number of binary semaphore variables modelling the forks. Processes  $P_i$   
 20 with  $i < n$  continuously attempt to first acquire the semaphore  $y_i$  and second  
 21 the semaphore  $y_{i+1}$ , whereas process  $P_n$  attempts to acquire first  $y_1$  and then  
 22  $y_n$ . Hence, all philosophers will always pick up the lower-numbered fork first  
 23 and the higher-numbered fork second. Once a process has successfully acquired  
 24 both semaphores it consecutively releases them and attempts to acquire them  
 25 again.




Figure 14: Dining philosophers system  $Sys$ .

TVMC generally searches for violations of desirable properties. For an instantiation of the dining philosophers system with  $n = 2$  the violation of deadlock-freedom can be expressed in LTL as

$$\psi = \mathbf{F}((loc_1 = 01) \wedge (loc_2 = 01)).$$

Hence, we want to check whether a state is reachable where each philosopher has picked up one fork and is waiting for the other fork. Starting with  $Pred = \emptyset$  and  $b = 1$  our tool automatically constructs the corresponding abstract control flow graphs and the encoding  $\llbracket Sys, \psi \rrbracket_b$ . Next, it iteratively increments the bound and refines the initial abstraction. For our example the bound will be increased until an unconfirmed witness path for the property of interest will be detected at  $b = 2$ :

$$\langle 00, 00 \rangle \xrightarrow{\perp} \langle 01, 00 \rangle \xrightarrow{\perp} \langle 01, 01 \rangle$$

1 Then cause-guided refinement will add the predicates  $y_1$  and  $y_2$  in a single  
 2 step, which yields the level of abstraction characterised by the abstract control  
 3 flow graphs depicted in Figure 15. Finally the bound will be further increased  
 4 until a completeness threshold is reached, which is the case for  $b = 64$  for this  
 5 verification task. A technique for computing over-approximations of complete-  
 6 ness thresholds is introduced in [14]. Completeness thresholds for checking LTL  
 7 properties that are restricted to the temporal operators  $\mathbf{F}$  and  $\mathbf{G}$  are linear in  
 8 the size of the abstraction, i.e. in the number of abstract states.

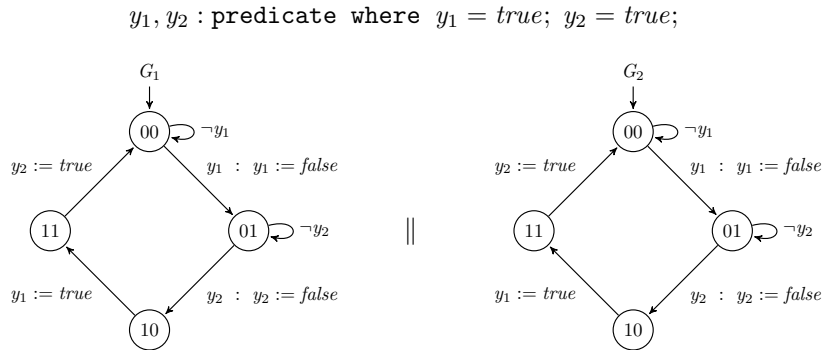


Figure 15: Abstraction of the dining philosophers system with  $n = 2$ .

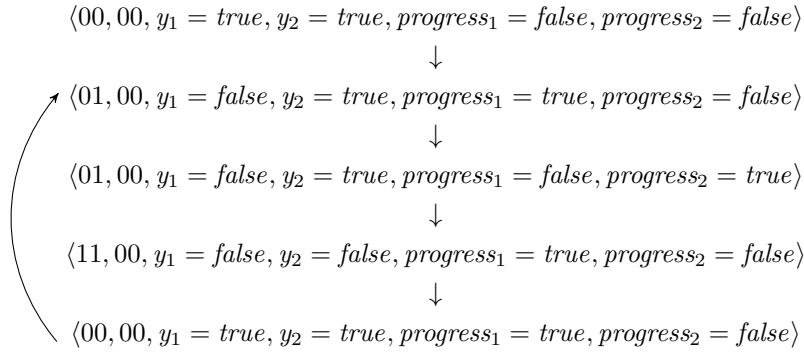


1 Via SAT solving we obtain a *false* result in the final iteration, which allows  
 2 us to conclude that the instantiation of the dining philosophers system with 2  
 3 processes is *safe* in terms of deadlock-freedom.

A distinct feature of our approach is the verification of *liveness* properties of concurrent systems under fairness assumptions. The formula

$$\psi' = \bigvee_{i=1}^n \mathbf{F}(\mathbf{G}\neg(\text{loc}_i = 10))$$

characterises the violation of a liveness property regarding our dining philosophers system. It states that eventually some philosopher process will nevermore reach its critical location. For the instantiation of the system with  $n = 2$  processes and starting with  $Pred = \{progress_1, progress_2\}$  (fairness predicates only) and  $b = 1$  our tool constructs the encoding  $\llbracket Sys, wfair \wedge \psi' \rrbracket_b$ . Within four iterations over  $b$  and a refinement step adding the predicates  $y_1$  and  $y_2$  we can already detect a satisfying assignment for the encoding that characterises a weak fair path with a loop where  $P_2$  never reaches its critical location:



4 Thus, we have proven that liveness of  $Sys$  is violated under *weak* fairness.  
 5 With our tool we could also prove that liveness of  $Sys$  holds under *strong* fair-  
 6 ness, which required to set the bound to the completeness threshold of the  
 7 verification task. Moreover, we could successfully verify generalisations of the  
 8 dining philosophers system with more philosophers and semaphores, which we  
 9 will discuss in the experimental results section.

10 The previous examples illustrate the basic functionality of the TVMC tool.  
 11 Next, we introduce an extension that allows for an improved verification of  
 12 safety properties, without setting the bound to the completeness threshold.

### 13 7.2. Complete Safety Verification via Temporal Induction

Bounded model checking is inherently incomplete since the bound restricts the length of the explored paths. Hence, it is predominantly used for detecting property violations rather than for proving their absence. As we previously outlined, a simple way of making bounded model checking complete is to determine a completeness threshold of a verification task and to set the bound to this threshold. However, determining small or minimal completeness thresholds

is costly and for many verification tasks even the minimal threshold might be too large for efficient verification. An alternative way of establishing completeness is the combination of bounded model checking with *temporal induction* [10]. Temporal induction allows to reduce an unbounded model checking problem to two bounded model checking problems: the *base case* and the *inductive step*. The approach is defined for checking safety violations of the form  $\mathbf{F}(unsafe)$  where *unsafe* is a predicate expression. In the base case it is checked whether there exists a  $b$ -bounded path, starting in the initial state, that witnesses the safety violation. In the inductive step it is checked whether there exists a loop-free  $(b+1)$ -bounded path, starting in an arbitrary state, where *unsafe* only holds in the  $(b+1)$ st state. If there exists a bound  $b$  for which neither in the base case nor in the step such witnesses can be detected, then it can be concluded that safety is globally not violated. Both the base case and the inductive step can be straightforwardly expressed as three-valued bounded model checking problems of our framework. The base case does not require any change of our encoding:

$$\llbracket Sys, \mathbf{F}(unsafe) \rrbracket_b^{base} := \llbracket Sys, \mathbf{F}(unsafe) \rrbracket_b.$$

For the step we have to remove the initial state constraint *Init*<sub>0</sub> and add a constraint *loopFree* that ensures that all states along the considered path prefixes are pairwise different. Moreover, we have to adjust the property encoding such that it is checked whether *unsafe* only holds in the  $(b+1)$ st state:

$$\llbracket Sys, \mathbf{F}(unsafe) \rrbracket_{b+1}^{step} := \bigwedge_{k=0}^b Trans_{k,k+1} \wedge loopFree \wedge \neg \llbracket \mathbf{F}(unsafe) \rrbracket_b^0 \wedge \llbracket unsafe \rrbracket_{b+1}^{b+1}$$

1 with  $loopFree = \bigwedge_{0 \leq i < j \leq b+1} (\bigvee_{p \in AP} \neg(enc(p)_i \leftrightarrow enc(p)_j))$ .

2 We have integrated the temporal induction approach into TVMC. Similar to  
 3 standard three-valued bounded model checking, we have to consider the cases  
 4 where  $\perp$  is mapped to *true* and where  $\perp$  is mapped to *false*, but now for both  
 5 the base case and the inductive step. A *true* result for the base case implies that  
 6 we have detected a safety violation. A *false* result for both the base case and the  
 7 step implies that safety is globally not violated. If we obtain an *unknown* result  
 8 for either the base case or the step, then we apply cause-guided abstraction  
 9 refinement. If we get *false* for the base case and *true* for the inductive step,  
 10 then we have to increment the bound. We checked deadlock-freedom for the  
 11 dining philosophers example in Figure 14 with the induction approach. As the  
 12 predicate expression *unsafe* we used  $(loc_1 = 01) \wedge (loc_2 = 01)$ . In comparison  
 13 with the completeness threshold approach, that required the final bound  $b = 64$ ,  
 14 temporal induction enabled us to already prove mutual exclusion with bound  
 15  $b = 3$ . In our experimental evaluation we will present a case study on the  
 16 induction-based approach.

### 17 7.3. Parameterised Verification via Spotlight Abstraction

18 So far, we have seen that our tool can be used for detecting property vi-  
 19 olations and also for proving their absence. However, in all of our previous

1 examples we verified systems with a *fixed* number of processes. The three-  
 2 valued abstractor that we use also allows to construct finite abstractions of  
 3 parameterised systems with an *unbounded* number of uniform processes [11]. A  
 4 simple example of a parameterised system with uniform processes can be de-  
 5 rived from our dining philosophers. Let  $Sys = \parallel_{i=1}^n P_i$  be a concurrent system  
 6 over  $Var = \{y_1, y_2\}$  where each  $P_i$  is a philosopher that picks up  $y_1$  first and  $y_2$   
 7 second. As long as we do not instantiate the number of processes  $n$ , this system  
 8 is *parameterised*. But in contrast to our previous philosopher example, where  
 9 all processes requested different pairs of forks in different orders, the processes  
 10 of our modified example are fully uniform in terms of fork requests.

For constructing a finite abstraction of such a parameterised uniform sys-  
 tem, three-valued predicate abstraction is combined with *spotlight abstraction*  
 [2]. The basic concept of spotlight abstraction is to partition a parallel compo-  
 sition of processes into a *spotlight* and a *shade*. The control flow of spotlight  
 processes is then explicitly considered under abstraction, whereas the processes  
 in the shade get summarised into a single abstract process  $P_\perp$  that approxi-  
 mates their behaviour with regard to three-valued logic. Hence, predicates over  
 variables that are modified by processes in the shade may be set to *unknown*  
 by  $P_\perp$ . For our parameterised uniform philosopher system  $Sys = \parallel_{i=1}^n P_i$ , spot-  
 light abstraction can be applied to the process partition  $Spotlight = \{P_1, P_2\}$   
 and  $Shade = \{P_3, \dots, P_n\}$ , i.e. we consider  $P_1$  and  $P_2$  explicitly and sum-  
 marise the parameterised number of processes  $P_3$  to  $P_n$  into  $P_\perp$ . The property  
 $\psi = \mathbf{F}((loc_1 = 01) \wedge (loc_2 = 01))$  can also be disproven for the abstraction  
 $P_1 \parallel P_2 \parallel P_\perp$  with our tool, which allows us to conclude that the processes  $P_1$   
 and  $P_2$  will never be in a deadlock situation for *any* instantiation of the uniform  
 philosopher system with  $n \geq 2$  processes. The LTL formula  $\psi$  characterises a  
 local property since it refers to particular processes of a parameterised system.  
 However, as shown in [11], symmetry arguments enable us to transfer this result  
 to arbitrary pairs of processes in the system. We can conclude that

$$\psi_{global} = \exists 1 \leq i, j \leq n, i \neq j : \mathbf{F}((loc_i = 01) \wedge (loc_j = 01))$$

11 does not hold for any instantiation of the uniform philosophers system, i.e. no  
 12 pair of processes will ever reach a deadlock. The approach also works for pa-  
 13 rameterised systems with different (but finite numbers of) classes of uniform  
 14 processes, e.g. a class of reader processes and a class of writer processes. More  
 15 details about spotlight abstraction and its combination with symmetry argu-  
 16 ments can be found in [11]. In our experimental evaluation we will present a  
 17 case study on parameterised verification with our tool.

#### 18 7.4. Experimental Evaluation

19 We have experimentally evaluated the performance of our tool in a num-  
 20 ber of case studies. In our experiments, we compared TVMC with the similar  
 21 three-valued model checking tool 3SPOT [2], which is also designed for the ver-  
 22 ification of concurrent systems. 3SPOT uses the same abstractor as TVMC that

1 yields abstract control flow graphs. After abstraction, the tool follows a differ-  
2 ent approach to solve verification tasks: The state space corresponding to the  
3 abstract control flow graphs is represented as a three-valued decision diagram  
4 (TDD), which is a generalisation of a Boolean decision diagram (BDD). The  
5 model checking algorithm of 3SPOT is therefore closely related to BDD-based  
6 CTL model checking [19]. Hence, we compare a decision diagram-based model  
7 checking approach with our SAT-based approach. Similar to TVMC, 3SPOT  
8 generates unconfirmed witness paths in case of an *unknown* result and refines  
9 the abstraction based on these paths. The witness paths of 3SPOT are explicitly  
10 generated and explored, whereas our witness paths are implicitly represented  
11 by truth assignments. Since 3SPOT is a CTL model checker, we focussed in our  
12 case studies on temporal logic properties from the common fragment of LTL  
13 and CTL. We conducted our experiments on a 1.6 GHz Intel Core i7 system  
14 with 8 GB memory. We measured the final bound, the number of refinement  
15 steps, the final number of predicates as well as the overall time for encoding  
16 and SAT-based model checking in all iterations, and we compared it with the  
17 performance results of 3SPOT.

18 In the case study PHILOSOPHERS, we verified generalisations of the dining  
19 philosopher system with  $n \leq 2$  processes. Each additional philosopher process  
20 involved an exponential growth of the state space complexity. We checked for the  
21 existence of a computation where eventually some philosopher will nevermore  
22 reach its critical location, which characterises the violation of a liveness property.  
23 Verification was performed under the assumption of weak fairness, as this is the  
24 only type of fairness constraint that is supported by 3SPOT. Since liveness of  
25 the philosopher system is not guaranteed under weak fairness, we could always  
26 detect a property violation. The performance results of the PHILOSOPHERS case  
27 study are shown in Table 1. We can see that for very small instances the time  
28 performance of 3SPOT was better, whereas for larger  $n$  our TVMC was signifi-  
29 cantly faster than 3SPOT. Hence, our SAT-based approach scaled considerably  
30 better here. Another observation is that both tools required the same number of  
31 predicates for completing the verification tasks. However, while 3SPOT needed  
32 several refinement iterations in order to reach the right level of abstraction, our  
33 cause-guided refinement detected and added all necessary predicates within a  
34 single step. More specifically, our tool generated an unconfirmed witness path  
35 in the first iteration from which all necessary predicates could be immediately  
36 derived in order to refine the path to a real witness.

37 In case study DIJKSTRA, we verified an implementation of Dijkstra’s mutual  
38 exclusion algorithm [20]. Again, we considered instances with an increasing  
39 number of processes. Dijkstra’s algorithm for two processes is depicted in Figure  
40 16. We checked whether there is no violation of mutual exclusion, i.e. whether  
41 there will be never more than one process at the critical location at the same  
42 time. In order to prove that this safety property holds, we used TVMC with  
43 the temporal induction approach discussed in Section 7.2. Since 3SPOT is an  
44 unbounded model checker, there was no need to use temporal induction with  
45 this tool. The performance results of the DIJKSTRA case study are also show in  
46 Table 1. Similar as in the previous case study we can observe that TVMC scales

case study	processes	TVMC				3SPOT		
		bound	refinements	predicates	time	refinements	predicates	time
PHILOSOPHERS	2	3	1	4	1.13s	2	4	0.41s
	3	5	1	6	2.12s	3	6	0.92s
	4	7	1	8	4.69s	4	8	47.7s
	5	9	1	10	12.4s	5	10	696s
	6	11	1	12	38.1s	6	12	152m
	7	13	1	14	379s	7	14	> 5h
DIJKSTRA	2	12	3	6	3.32s	2	2	0.18s
	3	16	4	9	22.2s	5	5	3.05s
	4	21	5	12	52m	> 6	> 6	OOM
	5	25	6	15	158m	> 6	> 6	OOM
PARAMETERISED	2 classes	4	1	2	0.85s	2	2	0.16s
	3 classes	6	1	3	1.17s	3	3	0.30s
	4 classes	8	1	4	1.82s	4	4	0.98s
	5 classes	10	1	5	8.63s	5	5	55.8s
	6 classes	12	1	6	89.2s	6	6	799s

Table 1: Experimental Results of the case studies PHILOSOPHERS, DIJKSTRA and PARAMETERISED.

1 better for larger systems than 3SPOT. Our new tool could even successfully  
2 prove safety when 3SPOT failed due to an out-of-memory exception (OOM).  
3 A second observation is that in the cases where both tools succeeded (two and  
4 three processes), TVMC required more predicates than 3SPOT. This is due to the  
5 fact that proving the inductive step of the temporal induction approach requires  
6 some extra predicates because of the arbitrary initial state. Nevertheless, TVMC  
7 still showed a good performance with these extra predicates. If we compare  
8 the PHILOSOPHERS case study where a property violation could be detected  
9 with the DIJKSTRA case study where correctness could be proven, then we  
10 can see that in the latter TVMC requires multiple iterations of cause-guided  
11 refinement in order to achieve a definite result. This is due to the fact that in  
12 the second case study each refinement iteration rules out a different unconfirmed  
13 witness path until no more witnesses exist, whereas in the first case study the  
14 first detected unconfirmed witness could be immediately refined to a real one.  
15 Another observation that we can make by comparing the two case studies is  
16 that showing the absence of errors generally requires more computational effort  
17 than detecting errors, which is not surprising.

```

turn : integer where turn = 1;
flag1, flag2 : integer where flag1 = 0; flag1 = 0;

P1 :: [
  loop forever do
    0000: flag1 := 1;
    0001: while(turn ≠ 1)
      [
        0010: if(turn = 2)
          [
            0011: if(flag2 = 0)
              [
                0100: turn := 1; ] ] ] ]
    0101: flag1 := 2;
    0110: if(flag2 = 2)
      [ 0111: goto 0000; ]
    1000: CRITICAL
    1001: flag1 := 0; ] ] || P2 :: [
  loop forever do
    0000: flag2 := 1;
    0001: while(turn ≠ 2)
      [
        0010: if(turn = 1)
          [
            0011: if(flag1 = 0)
              [
                0100: turn := 2; ] ] ] ]
    0101: flag2 := 2;
    0110: if(flag1 = 2)
      [ 0111: goto 0000; ]
    1000: CRITICAL
    1001: flag2 := 0; ] ]

```

Figure 16: Implementation of Dijkstra’s Mutual Exclusion Algorithm with  $n = 2$ .

1 In case study PARAMETERISED, we evaluated parameterised verification via  
2 spotlight abstraction. For this, we looked at parameterised variants of the dining  
3 philosopher system with an increasing number of classes of uniform processes.  
4 In the simplest case, we considered a system with two classes of philosophers:  
5 an unbounded number  $n \in \mathbb{N}$  of philosophers that pick up fork  $y_1$  first and  
6 fork  $y_2$  second, and an unbounded number  $m \in \mathbb{N}$  of philosophers that pick up  
7  $y_2$  first and  $y_1$  second. The corresponding parameterised system is depicted in  
8 Figure 17. Moreover, we considered generalised cases e.g. a system with three  
9 classes and the three different pick up orders  $y_1$  then  $y_2$ ,  $y_2$  then  $y_3$ , and  $y_3$  then  
10  $y_1$ . In these parameterised systems a deadlock in the sense of circular waiting  
11 for forks is possible. We used TVMC with the spotlight abstraction approach  
12 discussed in Section 7.3 in order to check this property. Since 3SPOT also  
13 supports spotlight abstraction, we could compare the performance of the two  
14 tools. The experimental results of the PARAMETERISED case study are shown  
15 in Table 1. As we can see, our approach is also capable of checking properties of  
16 parameterised systems. Again, TVMC scales better than 3SPOT and our cause-  
17 guided refinement technique finds all necessary predicates for detecting a real  
18 witness path within a single refinement step.

```

y1, y2 : binary semaphore where y1 = true; y2 = true;

||_{i=1}^n P_i :: [
  loop forever do
    00: acquire (y1);
    01: acquire (y2);
    10: CRITICAL
      release (y1);
    11: release (y2); ] ] || ||_{j=1}^m P_j :: [
  loop forever do
    00: acquire (y2);
    01: acquire (y1);
    10: CRITICAL
      release (y2);
    11: release (y1); ] ]

```

Figure 17: Parameterised dining philosophers system with two classes of uniform processes.

19 All in all, our experiments showed that SAT-based three-valued bounded  
20 model checking is a feasible approach to the verification of concurrent systems.  
21 We were able to detect violations of safety and liveness properties and to prove  
22 the absence of such violations. Cause-guided refinement allows to rule out un-

1 certainty in the abstraction within a small number of iterations, especially when  
2 property violations can be detected. Moreover, we could enhance the perfor-  
3 mance of proving the absence of safety violations via temporal induction. Natu-  
4 rally, our approach cannot resolve the state explosion problem in general. Each  
5 additional process involves an exponentially greater complexity of the under-  
6 lying verification task. However, we demonstrated that for larger systems our  
7 TVMC tool scales significantly better than the comparable 3SPOT tool. Fur-  
8 thermore, we were able to integrate spotlight abstraction into our approach,  
9 which facilitates the verification of parameterised systems composed of uniform  
10 processes.

## 11 8. Related Work

12 Our SAT-based software verification technique is related to a number of  
13 existing approaches in the field of bounded model checking for software. The  
14 bounded model checker CBMC [21] supports the verification of *sequential C*  
15 *programs*. It is based on a Boolean abstraction of the input program and it  
16 allows for checking buffer overflows, pointer safety and assertions, but not full  
17 LTL properties. A similar tool is F-Soft [22]. This bounded model checker  
18 for sequential programs is restricted to the verification of *reachability proper-*  
19 *ties*. While CBMC and F-Soft support a wider range of program constructs like  
20 pointers and recursion, our approach focusses on the challenges associated with  
21 *concurrency* and the verification of *liveness properties* under fairness. The tool  
22 TCBMC [23] is an extension of CBMC for verifying *safety properties* of *concur-*  
23 *rent programs*. TCBMC introduces the concept of bounding context switches  
24 between processes, which is a special abstraction technique for reducing concur-  
25 rency. Our approach supports the process summarisation abstraction of 3Spot  
26 [2], which allows us to reduce the complexity induced by concurrency in a dif-  
27 ferent way. The verification of concurrent C programs is also addressed in [24].  
28 The authors introduce a tool that translates C programs into a TLA+ [25]  
29 specification which is then model checked via an explicit-state approach.

30 In contrast to the above mentioned tools, we employ *three-valued abstraction*,  
31 which preserves *true* and *false* results in verification. *Three-valued bounded*  
32 *model checking* is addressed in [6] and [7]. However, only in the context of  
33 *hardware verification* [6] resp. assuming that an *explicit three-valued Kripke*  
34 *structure* is given [7]. To the best of our knowledge, our approach is the first  
35 that supports software verification under fairness via an immediate propositional  
36 logic encoding and SAT-based BMC.

37 Abstraction refinement for SAT-based model checking is addressed in [26,  
38 27, 28]. The hardware verification approach presented in [26] employs Boolean  
39 abstraction by means of variable hiding. Counterexamples detected in the ab-  
40 stract model are simulated on the concrete model via SAT solving. Unsatisfia-  
41 bility results correspond to spurious counterexamples. In this case abstraction  
42 refinement is applied by deriving new variables from the unsatisfiable core. Very  
43 similar approaches are used in [27] and [28]. The authors of [27] additionally

1 show that their technique yields an over-approximative abstraction that pre-  
2 serves safety and always has a completeness threshold for not only refuting but  
3 also proving properties. The authors of [28] generate refinement interpolants  
4 from simulated spurious counterexamples in SAT-based model checking. In our  
5 three-valued approach we do not have to simulate abstract counterexamples.  
6 In case of an *unknown* result, our path characterising assignment allows us to  
7 derive new predicates from the set of unsatisfied clauses that is typically sig-  
8 nificantly smaller than the unsatisfiable core. Unconfirmed witness paths in a  
9 three-valued abstract model are also used for refinement in [2]. However, the  
10 proposed approach requires to explicitly generate and analyse paths, whereas  
11 our refinement happens based on unsatisfied clauses that implicitly represent  
12 uncertainty in the abstraction.

13 The notion of causality in the context of temporal logic model checking has  
14 also been used in [29]. The authors present a technique for detecting property  
15 violations in concurrent systems via counterfactual reasoning [30]. While clas-  
16 sical model checking techniques generate witness paths resp. counterexamples,  
17 the approach of [29] derives *causal factors* that lead to a property violation.  
18 These causal factors are orders of occurrences of events in the analysed system  
19 and, according to the authors, provide a more concise explanation of why the  
20 desired property does not hold. The technique of [29] explores the concrete state  
21 space of the system, whereas our approach is based on three-valued predicate  
22 abstraction and iterative refinement: The lack of predicates over certain system  
23 variables may cause an *unknown* result in verification along with an unconfirmed  
24 witness. Hence, in our scenario a *cause of uncertainty* tells us which predicates  
25 are missing at the current level of abstraction in order to either confirm or to  
26 rule out the unconfirmed witness.

27 Our work is also related to other analysis techniques for concurrent sys-  
28 tems. The authors of [31, 32] propose a construction-based method for ensuring  
29 deadlock- and livelock-freedom of process compositions. Hence, the systems to  
30 be analysed as well as the desired properties are similar to the ones in our ap-  
31 proach. While [31] focusses on communication via message passing, we consider  
32 asynchronous systems with shared variables. [31] is based on the correctness-  
33 by-construction paradigm, whereas our technique aims at the verification of  
34 implementations of software systems. Thus, the two approaches complement  
35 each other in the sense that [31] can ensure a correct design, whereas our ap-  
36 proach can ensure the correct implementation of a design. In [32] local analysis  
37 is used for establishing correctness of the entire system. This is slightly related  
38 to the spotlight abstraction that we use in our approach. Spotlight abstraction  
39 also facilitates the verification of local properties. For uniform systems such  
40 local verification results can be transferred to the global system.

## 41 9. Conclusion and Outlook

42 We introduced a verification technique for concurrent software systems based  
43 on three-valued abstraction, cause-guided refinement and SAT-based bounded  
44 model checking. We defined a direct propositional logic encoding of software



1 verification tasks and we proved that our encoding is sound in the sense that SAT  
2 results can be straightforwardly transferred to the corresponding model checking  
3 problem. Hence, the expensive construction and exploration of an explicit state  
4 space model is not necessary. Our tool enables the verification of safety and  
5 liveness properties under fairness. With cause-guided refinement we introduced  
6 an automatic technique for systematically reaching the right level of abstraction  
7 in order to obtain a definite outcome in verification. Refinement steps are  
8 straightforwardly derived from clauses of the encoding that are unsatisfied under  
9 assignments characterising potential witness/error paths. Due to the efficiency  
10 of modern SAT solvers we achieve promising performance results with our overall  
11 approach.

12 As future work we plan to optimise our technique by integrating *incremental*  
13 SAT solving [33] into the tool and by developing a concept for reusing parts of  
14 the encoding between the consecutive refinement iterations. Finally, we want to  
15 develop SAT solving *heuristics* tailored to the structure of our encodings [34] in  
16 order to further accelerate our approach.

## 17 Acknowledgements

18 We thank Dewald de Jager and Matthias Harvey for help with implementing  
19 our approach.

## 20 References

- 21 [1] S. Shoham, O. Grumberg, 3-valued abstraction: More precision at less cost,  
22 Information and Computation 206 (11) (2008) 1313–1333.
- 23 [2] J. Schriebl, H. Wehrheim, D. Wonisch, Three-valued spotlight abstractions,  
24 in: A. Cavalcanti, D. Dams (Eds.), FM, Vol. 5850 of LNCS, Springer, 2009,  
25 pp. 106–122.
- 26 [3] G. Bruns, P. Godefroid, Model checking partial state spaces with 3-valued  
27 temporal logics, in: CAV 1999, LNCS, Springer Berlin Heidelberg, 1999,  
28 pp. 274–287.
- 29 [4] A. Cimatti, E. Clarke, F. Giunchiglia, M. Roveri, NuSMV: a new symbolic  
30 model checker, Int. Jour. on Softw. Tools for Techn. Transfer 2 (4) (2000)  
31 410–425.
- 32 [5] A. Biere, A. Cimatti, E. M. Clarke, O. Strichman, Y. Zhu, Bounded model  
33 checking., Handbook of Satisfiability 185 (2009) 457–481.
- 34 [6] O. Grumberg, 3-Valued Abstraction for (Bounded) Model Checking,  
35 Springer Berlin Heidelberg, Berlin, Heidelberg, 2009, pp. 21–21. doi:  
36 10.1007/978-3-642-04761-9\_2.

- 1 [7] H. Wehrheim, Bounded model checking for partial Kripke structures, in:  
2 J. Fitzgerald, A. Haxthausen (Eds.), ICTAC, Vol. 5160 of LNCS, Springer,  
3 2008, pp. 380–394. doi:10.1007/978-3-540-85762-4\_26.
- 4 [8] N. Timm, S. Gruner, M. Harvey, A Bounded Model Checker for  
5 Three-Valued Abstractions of Concurrent Software Systems, Springer  
6 International Publishing, Cham, 2016, pp. 199–216. doi:10.1007/  
7 978-3-319-49815-7\_12.  
8 URL [http://dx.doi.org/10.1007/978-3-319-49815-7\\_12](http://dx.doi.org/10.1007/978-3-319-49815-7_12)
- 9 [9] E. Clarke, O. Grumberg, S. Jha, Y. Lu, H. Veith, Counterexample-guided  
10 abstraction refinement, in: CAV, Vol. 1855 of LNCS, Springer, 2000, pp.  
11 154–169. doi:10.1007/10722167\_15.
- 12 [10] N. En, N. Srensson, Temporal induction by incremental sat solving,  
13 Electronic Notes in Theoretical Computer Science 89 (4) (2003) 543 – 560,  
14 bMC’2003, First International Workshop on Bounded Model Checking.  
15 doi:[https://doi.org/10.1016/S1571-0661\(05\)82542-3](https://doi.org/10.1016/S1571-0661(05)82542-3).  
16 URL [http://www.sciencedirect.com/science/article/pii/  
17 S1571066105825423](http://www.sciencedirect.com/science/article/pii/S1571066105825423)
- 18 [11] N. Timm, H. Wehrheim, On symmetries and spotlights - verifying param-  
19 eterised systems, in: J. Dong, H. Zhu (Eds.), ICFEM, Vol. 6447 of LNCS,  
20 Springer, 2010, pp. 534–548.
- 21 [12] N. Timm, H. Wehrheim, M. Czech, Heuristic-guided abstraction re-  
22 finement for concurrent systems, in: T. Aoki, K. Taguchi (Eds.),  
23 ICFEM, Vol. 7635 of LNCS, Springer, 2012, pp. 348–363. doi:10.1007/  
24 978-3-642-34281-3\_25.
- 25 [13] M. Fitting, Kleene’s 3-valued logics and their children, Fund. Inf. 20 (1-3)  
26 (1994) 113–131.
- 27 [14] D. Kroening, J. Ouaknine, O. Strichman, T. Wahl, J. Worrell, Linear com-  
28 pleteness thresholds for bounded model checking, in: CAV, Springer, 2011,  
29 pp. 557–572.
- 30 [15] N. Timm, Bounded model checking für partielle Systeme, Masters thesis,  
31 University of Paderborn.
- 32 [16] N. Timm, S. Gruner, Three-valued bounded model checking with cause-  
33 guided abstraction refinement – proofs, Tech. rep., Department of Com-  
34 puter Science, University of Pretoria (August 2018).  
35 URL <http://hdl.handle.net/2263/66136>
- 36 [17] L. Moura, N. Bjørner, Z3: An efficient SMT solver, in: C. R. Ramakrishnan,  
37 J. Rehof (Eds.), Tools and Algorithms for the Construction and Analysis of  
38 Systems, Vol. 4963 of Lecture Notes in Computer Science, Springer-Verlag  
39 Berlin Heidelberg, 2008, pp. 337–340. doi:10.1007/978-3-540-78800-3\_

- 1 24.  
2 URL [http://dx.doi.org/10.1007/978-3-540-78800-3\\_24](http://dx.doi.org/10.1007/978-3-540-78800-3_24)
- 3 [18] D. Le Berre, A. Parrain, The Sat4j library, release 2.2, Journal on Satisfi-  
4 ability, Boolean Modeling and Computation 7 (2010) 59–64.
- 5 [19] J. Burch, E. Clarke, K. McMillan, D. Dill, L. Hwang, Symbolic model  
6 checking: 1020 states and beyond, Information and Computation 98 (2)  
7 (1992) 142 – 170. doi:[https://doi.org/10.1016/0890-5401\(92\)](https://doi.org/10.1016/0890-5401(92)90017-A)  
8 90017-A.  
9 URL [http://www.sciencedirect.com/science/article/pii/](http://www.sciencedirect.com/science/article/pii/S089054019290017A)  
10 089054019290017A
- 11 [20] E. W. Dijkstra, Solution of a problem in concurrent programming control,  
12 Commun. ACM 8 (9) (1965) 569–. doi:[10.1145/365559.365617](https://doi.org/10.1145/365559.365617).  
13 URL <http://doi.acm.org/10.1145/365559.365617>
- 14 [21] D. Kroening, M. Tautschnig, CBMC–C bounded model checker, in:  
15 TACAS, Springer, 2014, pp. 389–391.
- 16 [22] F. Ivančić, Z. Yang, M. K. Ganai, A. Gupta, I. Shlyakhter, P. Ashar, F-Soft:  
17 Software verification platform, in: International Conference on Computer  
18 Aided Verification, Springer, 2005, pp. 301–306.
- 19 [23] I. Rabinovitz, O. Grumberg, Bounded model checking of concurrent pro-  
20 grams, in: CAV, Springer, 2005, pp. 82–97.
- 21 [24] A. Methni, M. Lemerre, B. Ben Hedia, S. Haddad, K. Barkaoui, Specifying  
22 and Verifying Concurrent C Programs with TLA+, Springer International  
23 Publishing, Cham, 2015, pp. 206–222. doi:[10.1007/978-3-319-17581-2\\_](https://doi.org/10.1007/978-3-319-17581-2_14)  
24 14.  
25 URL [http://dx.doi.org/10.1007/978-3-319-17581-2\\_14](http://dx.doi.org/10.1007/978-3-319-17581-2_14)
- 26 [25] S. Merz, The Specification Language TLA+, Springer Berlin Hei-  
27 delberg, Berlin, Heidelberg, 2008, pp. 401–451. doi:[10.1007/](https://doi.org/10.1007/978-3-540-74107-7_8)  
28 978-3-540-74107-7\_8.  
29 URL [http://dx.doi.org/10.1007/978-3-540-74107-7\\_8](http://dx.doi.org/10.1007/978-3-540-74107-7_8)
- 30 [26] A. Gupta, O. Strichman, Abstraction Refinement for Bounded Model  
31 Checking, Springer Berlin Heidelberg, Berlin, Heidelberg, 2005, pp. 112–  
32 124. doi:[10.1007/11513988\\_11](https://doi.org/10.1007/11513988_11).  
33 URL [http://dx.doi.org/10.1007/11513988\\_11](http://dx.doi.org/10.1007/11513988_11)
- 34 [27] D. Kroening, Computing over-approximations with bounded model check-  
35 ing, Electron. Notes Theor. Comput. Sci. 144 (1) (2006) 79–92. doi:  
36 [10.1016/j.entcs.2005.07.021](https://doi.org/10.1016/j.entcs.2005.07.021).  
37 URL <http://dx.doi.org/10.1016/j.entcs.2005.07.021>

- 1 [28] C. Y. Wu, C. A. Wu, C.-Y. Lai, C. Y. Huang, A counterexample-guided  
2 interpolant generation algorithm for sat-based model checking, in: 2013  
3 50th ACM/EDAC/IEEE Design Automation Conference (DAC), 2013, pp.  
4 1–6.
- 5 [29] A. Beer, S. Heidinger, U. Kühne, F. Leitner-Fischer, S. Leue, Symbolic  
6 causality checking using bounded model checking, in: B. Fischer, J. Gelden-  
7 huys (Eds.), Model Checking Software, Springer International Publishing,  
8 Cham, 2015, pp. 203–221.
- 9 [30] J. Y. Halpern, J. Pearl, Causes and explanations: A structural-model ap-  
10 proach. part i: Causes, The British journal for the philosophy of science  
11 56 (4) (2005) 843–887.
- 12 [31] R. Ramos, A. Sampaio, A. Mota, Systematic development of trustworthy  
13 component systems, in: A. Cavalcanti, D. R. Dams (Eds.), FM 2009: For-  
14 mal Methods, Springer Berlin Heidelberg, Berlin, Heidelberg, 2009, pp.  
15 140–156.
- 16 [32] M. S. C. Filho, M. V. M. Oliveira, A. Sampaio, A. Cavalcanti, Local livelock  
17 analysis of component-based models, in: K. Ogata, M. Lawford, S. Liu  
18 (Eds.), Formal Methods and Software Engineering, Springer International  
19 Publishing, Cham, 2016, pp. 279–295.
- 20 [33] A. Nadel, V. Ryvchin, O. Strichman, Ultimately incremental sat, in:  
21 C. Sinz, U. Egly (Eds.), Theory and Applications of Satisfiability Testing  
22 – SAT 2014, Springer International Publishing, Cham, 2014, pp. 206–218.
- 23 [34] N. Timm, S. Gruner, P. Sibanda, Model checking of concurrent software  
24 systems via heuristic-guided sat solving, in: M. Dastani, M. Sirjani (Eds.),  
25 Fundamentals of Software Engineering, Springer International Publishing,  
26 Cham, 2017, pp. 244–259.