

Model Checking Safety and Liveness via k -Induction and Witness Refinement with Constraint Generation

Nils Timm*, Stefan Gruner, Madoda Nxumalo, Josua Botha

Department of Computer Science, University of Pretoria, South Africa

Abstract

In this article, we revise our constraint-based abstraction refinement technique for checking temporal logic properties of concurrent software systems. Our technique employs predicate abstraction and SAT-based three-valued bounded model checking. In contrast to classical refinement techniques where a single state space model is iteratively explored and refined with predicates, our approach is as follows: We use a coarsely-abstracted global state space model where we check for abstract witness paths for the property of interest. For each detected abstract witness we construct a local model whose state space is restricted to refinements of the witness only. On the local models we check whether the witness is real or spurious. We eliminate spurious witnesses in the global model via *spurious segment constraints*, which do not increase the state space complexity. Our technique is complete and terminates when a real witness in a local model can be detected, or no more witnesses in the global model exist.

While our technique was originally restricted to the verification of *safety* properties, we extend it here to the verification of *liveness* properties. For this, we make use of the state recording translation of the input system, which reduces liveness model checking to safety checking. Another restriction of our original approach was its incompleteness due to the nature of *bounded* model checking. Here we show how abstraction refinement-based bounded model checking can be combined with the k -induction principle, which enables unbounded model checking. Our approach is iterative with regard to the bound. The extended approach also allows us to define enhanced concepts for *strengthening* the constraints that we use to rule out spurious behaviour and for *reusing* constraints between bound iterations. We demonstrate that our approach enables the complete verification of safety and liveness properties with a reduced state space complexity and a better solving time in comparison to classical abstraction refinement techniques.

Key words: Three-valued abstraction, SAT-based bounded model checking, Constraint generation, k -induction, Liveness

*Corresponding author

Email addresses: ntimm@cs.up.ac.za (Nils Timm), sg@cs.up.ac.za (Stefan Gruner)

1 Introduction

Three-valued abstraction (3VA) [1] is an established technique for reducing the complexity of software verification. It proceeds by generating an abstract representation of an input software system over predicates with the possible truth values *true*, *false* and *unknown*, where the latter value is used to express the loss of information due to abstraction. The state space of an abstract software system can then be represented as a three-valued model. The evaluation of temporal logic properties on such models is known as *three-valued model checking* (3MC) [2]. Under three-valued abstraction both *true* and *false* model checking results can be transferred to the modelled software system, whereas an *unknown* result indicates that the current model is too coarse for a definite outcome. In the latter case a so-called *unconfirmed witness* is produced, which is an execution path in the abstract state space with some *unknown* transitions or predicates that characterises a potential violation of the property of interest. *Witness-guided abstraction refinement* [3] then iteratively adds further predicates to the abstract model until a previously *unconfirmed* witness turns out to be *definite*, or no more witnesses exist. The described approach follows the classical *abstract-check-refine* paradigm [4] where a single model that represents the entire system is iteratively refined. Since each refinement iteration involves an exponential growth of the state space to be explored, this approach can easily suffer from state explosion.

In [5] we introduced a novel abstraction refinement technique that facilitates the verification of safety properties of concurrent software systems with an improved state space complexity. In this approach we make use of two kinds of state space models: We use a *global model* that considers all parts of the underlying system, and we use *local models* that are restricted to previously detected unconfirmed witness paths. Both global and local models are subject to three-valued abstraction. But only the local models are refined by adding predicates, whereas the global model is iteratively pruned via *spurious witness constraints* derived from local models. Our technique proceeds as follows: In the same manner as in a classical *abstract-check-refine* approach, we start with a coarsely-abstracted global model of the input system and we check whether the safety property of interest can be proven or refuted. If the check returns *unknown* along with an unconfirmed witness, then we derive new predicates for refinement. Now instead of refining the global model, we construct a new local model that is narrowed down to refinements of the unconfirmed witness only. Checking the local model either proves the previously unconfirmed witness to be *definite* or to be *spurious*. In the first case we are done. In the latter case we generate a constraint for ruling out the spurious witness. In the subsequent pruning iteration we return to the global model and prune its state space via the generated *spurious witness constraint*. The procedure terminates when either no more witnesses in the global model exist or a definite witness in a local model can be detected. In contrast to standard approaches to constraint reusing in bounded model checking [6] our technique does not only enable the reuse of constraints between bound iterations but also between different levels

1 of abstraction. A constraint generated based on a refined local model is also
2 admissible for the more abstract global model. Our approach reduces the state
3 space complexity in two ways. Refinement is only applied to *local* models whose
4 state space is already strongly limited by being restricted to refinements of a
5 certain unconfirmed witness. The state space of the global model is pruned
6 by spurious witness constraints derived from local models. But the refinement
7 predicates that were used in the local model in order to derive these constraints
8 do *not* have to be added to the global model. Hence, we gain precision in the
9 global model without enlarging its state space. The price that we pay is an in-
10 creased number of global pruning iterations and local refinement iterations until
11 a definite result can be obtained. The actual number depends on the *strength*
12 of the generated constraints in terms of ruling out spurious behaviour. The
13 spuriousness of a witness typically originates from a *spurious segment* of the
14 path that it represents. A constraint that rules out all paths that exhibit the
15 *spurious segment* is naturally stronger than a constraint that only rules out the
16 spurious witness itself.

17 As a background technique we use *satisfiability-based three-valued bounded*
18 *model checking* [7] in order to process the model checking problems to be solved
19 within our abstraction refinement approach. We have shown that the abstracted
20 input systems together with the safety property to be checked can be directly
21 encoded into propositional logic such that the construction of an explicit state
22 space model is avoided [7]. Encoded three-valued bounded model checking prob-
23 lems can be solved via two Boolean satisfiability checks. The first check con-
24 siders an over-approximating completion of the encoding where all *unknowns*
25 are assumed to be *true*. The second check considers an under-approximating
26 completion where all *unknowns* are assumed to be *false*. If both completions
27 are satisfiable, then the corresponding model checking result is *true*. If both
28 completions are unsatisfiable, then the corresponding model checking result is
29 *false*. Otherwise the result is *unknown*. In [5] we demonstrated that in the
30 satisfiability-based approach a spurious segment of a witness is characterised by
31 a part of an unsatisfiable core of the SAT-encoded problem. The constraint for
32 ruling out the spurious segment corresponds to the negation of this part. This
33 enables us to efficiently generate spurious segment constraints via *unsatisfiable*
34 *core extraction* [8]. Bounded model checking, as we used it in [3, 5, 7], is in-
35 herently incomplete. The bound $k \in \mathbb{N}$ restricts the length of execution paths
36 of the modelled system, which makes this technique only usable for detecting
37 property violations but not for proving their absence. Completeness for finite-
38 state systems can be theoretically established by iterating over the bound until
39 a *completeness threshold* is reached [9]. However, the determination of mini-
40 mal or tight (close to minimal) completeness thresholds is a computationally
41 hard problem by itself and even tight thresholds are mostly still impractical for
42 efficient verification. In [5] we added a bound iteration loop around our abstrac-
43 tion refinement technique. The loop is *incremental* [10] in the sense that clauses
44 learned by the solver for a particular model M in iteration k can be reused in
45 for pruning the search space of the *same* model in iteration $k + 1$. In addi-
46 tion, our novel approach allows for further pruning based on spurious segment

1 constraints. We proved that *iteration-independent* spurious segment constraints
2 derived from *any* local model can be used for restricting the search space of the
3 global model in all bound iterations. The admissibility of reusing constraints
4 between different models gives us pruning capabilities that are beyond what is
5 feasible with standard incremental SAT solving [10]. The approach allows us to
6 avoid computational overhead caused by repeating constraint generation that
7 has been already conducted based on local models in previous iterations

8 While our novel refinement approach [5] already revealed promising results
9 in terms of reducing the state space complexity of verification tasks in com-
10 parison to classical abstraction refinement [3], it was still subject to a number
11 of drawbacks and limitations. Our original definition of iteration-independence
12 was overly restrictive and resulted in a very limited reusability of constraints.
13 In this extended article we define seven different types of independence (resp.
14 dependence) of spurious segment constraints and we prove an admissible form
15 of reuse for each type. This also includes the additional strengthening of *fully-*
16 *independent* constraints and the adaptation of *initial state-independent* con-
17 straints for reuse in higher bound iterations. The type of independence of a
18 constraint follows immediately from unsatisfiable core extraction. Our enhanced
19 constraint reusing concept facilitates a more extensive pruning of spurious be-
20 haviour, and thus, a better verification performance.

21 A second limitation of [5] is that it is incomplete due to the nature of bounded
22 model checking. Although we use a bound iteration loop, this does not allow
23 us to prove the absence of property violations in large-scale state space mod-
24 els. Completeness of bounded model checking can be established via *k-induction*
25 [11]. This technique was originally introduced for the verification of safety prop-
26 erties of hardware systems. It proceeds as follows: Given a state space model of
27 the system to be analysed and a safety predicate *safe*, it is checked whether all
28 paths of length *k* that start in an initial state of the model are safe, i.e. whether
29 *safe* holds in each state along the paths. This is the *base case* of *k-induction*,
30 which is equivalent to standard bounded model checking. If the base case holds,
31 then the *inductive step* is checked: Assuming *k* consecutive states where *safe*
32 holds in each state, then *safe* also has to hold in every $(k + 1)$ -st successor state.
33 The inductive step does not restrict the *k* consecutive states to start in an ini-
34 tial state. If the inductive step holds as well, then it can be concluded that all
35 *unbounded* execution paths of the modelled system are also safe. Otherwise the
36 procedure needs to be repeated with an incremented *k*. For finite-state models
37 termination is guaranteed and the final bound is typically considerably smaller
38 than a precomputed approximation of a completeness threshold. In [12] we al-
39 ready demonstrated that *k-induction* is compatible with three-valued bounded
40 model checking. In this article we show that the *k-induction* technique can
41 be also combined with our novel abstraction refinement approach. For this,
42 we introduce a shared bound iteration loop and within this loop two separate
43 refinement loops, one for the base case and one for the inductive step. This
44 combination enables us to conduct *complete* verification of safety properties via
45 SAT-based model checking. The base case and the inductive step are two dis-
46 tinct problems to be solved. Although they exhibit certain similarity, constraint

1 reuse between the two is not admissible in general. However, we show that gen-
 2 erated *initial state-independent* constraints can be reused between the base case
 3 and the inductive step, which gives us further pruning capabilities.

4 k -induction is limited to the verification of *safety* properties and so is our
 5 original approach presented in [5]. For concurrent systems also *liveness* proper-
 6 ties are of great importance. Liveness model checking under fairness involves a
 7 considerably higher complexity than safety checking and many safety checking
 8 techniques are not compatible with liveness. In order to facilitate the verification
 9 of liveness properties with our approach, we adopt the *state recording transla-*
 10 *tion* [13]. This translation transfers an original state space model into a state
 11 recording model, which reduces an original liveness model checking problem to
 12 a safety problem. The translation comes at the cost of a quadratic increase
 13 of the number of states, but it gives us the benefit to utilise efficient safety
 14 checking techniques in order to solve liveness problems. We show that the state
 15 recording translation can be already applied to our abstracted systems *before*
 16 the corresponding state space models are encoded in propositional logic.

17 1.1. Contributions and Relation to Previous Work

18 The main contributions of this work are the establishment of considerably en-
 19 hanced constraint reusing capabilities for our automatic abstraction refinement
 20 approach proposed in [5], the completion of previously incomplete verification
 21 techniques [3, 5, 7], and the extension of a pure safety checking technique to a
 22 technique that also supports liveness checking under fairness. Table 1 provides
 23 an overview on how this article, denoted as SCP 2020, extends and combines
 24 our previous work on SAT-based three-valued bounded model checking.

Table 1: Comparison with previous work.

	[7]	[3]	[12]	[5]	SCP 2020
liveness checking	✓	✓	✗	✗	✓
automatic refinement	✗	✓	✓	✓	✓
multi-model approach	✗	✗	✗	✓	✓
reuse of path constraints	✗	✗	✗	✓	✓ (enhanced)
completeness	✗	✗	✓	✗	✓

25 [7] and [3] support liveness model checking based on an *explicit* encoding
 26 of the liveness property to be checked, which turned out to be inefficient in
 27 experiments. Thus, in subsequent papers we focussed on safety model checking
 28 only. SCP 2020 re-introduces liveness support based on a *reduction* to safety.
 29 In [3] we introduced a fully-automatic abstraction refinement technique. Pred-
 30 icates for refinement are derived from unsatisfied clauses. Refined models are
 31 generated with the help of the prover Z3 [14]. The originally single-model ap-
 32 proach to abstraction refinement has been extended to a multi-model approach

1 in [5]. This also involved the introduction of path constraint reusing, which we
2 enhance in SCP 2020. Based on thresholds, complete model checking is theo-
3 retically conceivable but practically infeasible in [7], [3], [5]. The k -induction
4 approach allows for the feasibility of complete verification in [12] and SCP 2020.

5 We have implemented our approach. In experiments we demonstrate that
6 our constraint-based refinement technique allows for significant performance
7 improvements in comparison to classical abstraction refinement. Moreover, we
8 show that our implemented tool can compete with the Spin model checker [15]
9 for certain verification tasks.

10 1.2. Outline

11 The remainder of this article is organised as follows. In Section 2 we in-
12 troduce the concurrent software systems that we consider in our approach and
13 the three-valued abstraction technique that we employ. Section 3 provides the
14 background on three-valued bounded model checking. In Section 4 we show
15 how the state recording translation can be applied to our abstracted systems in
16 order to reduce liveness to safety model checking. In Section 5 we show how k -
17 induction can be combined with three-valued bounded model checking in order
18 to establish completeness. Section 6 reviews basic three-valued abstraction re-
19 finement. In Section 7 we introduce our novel witness refinement technique with
20 constraint reuse. We show in this section how the novel refinement technique
21 can be combined with k -induction. Moreover, we define the different types of
22 constraints that occur in our approach and we prove an admissible form of reuse
23 and strengthening for each type. In Section 8 we introduce our propositional
24 logic encoding of three-valued bounded model checking problems. Furthermore,
25 we show how constraint types can be determined based on unsatisfiable core
26 extraction. In Section 9 we introduce the implementation of our approach and
27 we present experimental results. Section 10 discusses related work. We conclude
28 this paper in Section 11 and give an outlook on future work.

29 2. Abstracted Concurrent Software Systems

30 We start with a brief introduction to the systems that we want to verify
31 and the abstraction technique that we use in our work. Our approach supports
32 integer arithmetic-based concurrent systems with the data types *int*, *bool* and
33 *semaphore* (but no arrays and pointers). Moreover, almost all control structures
34 of the C language are supported, such as *if-then-else*, *while-do*, *for* and *goto*. A
35 *concurrent software system* Sys consists of a number of possibly non-uniform
36 processes P_1 to P_n composed in parallel: $Sys = \parallel_{i=1}^n P_i$. It is defined over a
37 set of variables $Var = Var_{Sys} \cup Var_{PC}$. Var_{Sys} is a set of typed system vari-
38 ables, whereas Var_{PC} is a special set that holds for each process P_i a dedicated
39 program counter variable pc_i ranging over control locations from a set Loc_i .
40 Locations of a process are labelled with conditional operations with regard to
41 system variables and with a reference to the subsequent location.

42 An example for a concurrent system implementing mutual exclusion is de-
43 picted in Figure 1.

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

Figure 1: Concurrent system Sys .

1 Here we have two processes operating on a shared semaphore variable y .
2 Processes P_i can be formally represented as *control flow graphs* (CFGs) $G_i =$
3 $(Loc_i, \delta_i, \tau_i)$ where $Loc_i = \{0, \dots, |Loc_i - 1\}$ is a set of control locations given
4 as numbers. We implicitly assume that 0 is the initial location of a control flow
5 graph. $\delta_i \subseteq Loc_i \times Loc_i$ is a transition relation and $\tau_i : \delta_i \rightarrow Op$ is a function
6 labelling transitions with operations from a set Op .

7 **Definition 1 (Operations).**

8 *Let $Var = \{v_1, \dots, v_m\}$ be a set of typed variables. The set of operations Op on*
9 *Var consists of all statements of the form $assume(e) : [v_1 := e_1] \circ \dots \circ [v_m := e_m]$*
10 *where e is a Boolean expression over Var that acts as a guard. Moreover, \circ*
11 *is the append operator and $[...] \circ \dots \circ [...]$ is a list of type-correct assignments*
12 *where e_1, \dots, e_m are expressions over Var .*

13 Hence, every operation consists of a guard and a list of assignments. For
14 convenience, we sometimes just write e instead of $assume(e)$, we omit the *as-*
15 *sume* part completely if the guard is *true*, and we write $v_1 := e_1, \dots, v_m := e_m$
16 for an assignment list $[v_1 := e_1] \circ \dots \circ [v_m := e_m]$. The control flow graphs G_1
17 and G_2 corresponding to the processes of our example system are depicted in
18 Figure 2. G_1 and G_2 also illustrate the semantics of the operations $acquire(y, 1)$
19 and $release(y, 1)$.




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

20 A concurrent system given by n individual control flow graphs G_1, \dots, G_n
21 can be modelled by one composite CFG $G = (Loc, \delta, \tau)$ where $Loc = \times_{i=1}^n Loc_i$.
22 G is the product graph of all individual CFGs. We assume that for any Sys
23 a deterministic initialisation of all its variables is given in terms of a predicate
24 expression $Init$, e.g., $Init = (y = 1) \wedge (pc_1 = 0) \wedge (pc_2 = 0)$ for our example. A
25 computation of a concurrent system corresponds to a sequence where in each

1 step one process is non-deterministically selected and an operation at its cur-
 2 rent location that is not blocked by the guard is executed, i.e. the variables
 3 are updated according to the assignment part and the process advances to the
 4 consequent control location. We assume that for each step and each control lo-
 5 cation of a process there will be always at least one non-blocked operation that
 6 can be executed. This might be an explicit idling step, such as the waiting for
 7 the semaphore to become available at location 0 in Figure 2. We assume that
 8 all assignments of an operation are executed simultaneously and that an oper-
 9 ation assigns to each variable at most once. The state space of a system over
 10 Var corresponds to the set S_{Var} of all type-correct valuations of the variables.
 11 Given a state $s \in S_{Var}$ and an expression e over Var , then $s(e)$ denotes the
 12 valuation of e in s . Thus, a computation can be likewise considered as sequence
 13 of states $s_0 s_1 s_2 \dots$ where the transition from s_i to s_{i+1} correctly characterises
 14 the execution of the associated operation. For verifying properties of concurrent
 15 software systems typically only *fair* computations are considered. Our notion
 16 of fairness for concurrent systems is as follows: In an infinite computation, each
 17 process executes an operation infinitely often.

18 Control flow graphs allow us to model concurrent systems formally. For
 19 an efficient verification it is additionally required to reduce the state space
 20 complexity. For this purpose, we use *three-valued predicate abstraction* [1].
 21 Such an abstraction is an approximation in the sense that all definite verifi-
 22 cation results (*true*, *false*) obtained for an abstract system can be transferred
 23 to the original system. Only *unknown* results necessitate abstraction refine-
 24 ment [16]. For a given set A_{Sys} of atomic predicates over system variables,
 25 the corresponding three-valued abstraction of the system can be automatically
 26 constructed via a theorem prover. In our approach we employ the prover Z3
 27 [14]. In abstract systems operations do not refer to concrete system variables
 28 from Var_{Sys} but to predicates from a set A_{Sys} with the three-valued domain
 29 $\mathbb{3} = \{true, false, unknown\}$ which we typically abbreviate by $\{t, f, u\}$. *Unknown*
 30 is used to represent the loss of information due to abstraction and is a valid truth
 31 value as we operate with the three-valued *Kleene logic* \mathcal{K}_3 [17] whose semantics
 32 is given by the truth tables in Figure 3.

\wedge	<i>true</i>	<i>u</i>	<i>false</i>	\vee	<i>true</i>	<i>u</i>	<i>false</i>	\neg	
<i>true</i>	<i>true</i>	<i>u</i>	<i>false</i>	<i>true</i>	<i>true</i>	<i>true</i>	<i>true</i>	<i>true</i>	<i>false</i>
<i>u</i>	<i>u</i>	<i>u</i>	<i>false</i>	<i>u</i>	<i>true</i>	<i>u</i>	<i>u</i>	<i>u</i>	<i>u</i>
<i>false</i>	<i>false</i>	<i>false</i>	<i>false</i>	<i>false</i>	<i>true</i>	<i>u</i>	<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 $u \leq_{\mathcal{K}_3} true$,
 $u \leq_{\mathcal{K}_3} false$, 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 $\{p_1, \dots, p_m\} = A_{Sys}$ is a set of predicates and $a, b, a_1, b_1, \dots, a_m, b_m$
 2 are logical expressions over A_{Sys} . $choice(a, b)$ -expressions have the following
 3 semantics:

Definition 2 (Choice Expressions).

Let s be a state over a set of three-valued predicates A_{Sys} . Moreover, let a and b be logical expressions over A_{Sys} that may evaluate to true false or unknown in s . Then

$$s(choice(a, b)) = \begin{cases} true & \text{if, and only if, } s(a) \text{ is true,} \\ false & \text{if, and only if, } s(a) \text{ is not true and } s(b) \text{ is true,} \\ u & \text{else.} \end{cases}$$

4 Thus, predicates in A_{Sys} may be set to *unknown* by an abstract operation.
 5 Given a concrete operation $assume(e) : v_1 := e_1, \dots, v_m := e_m$ and a set
 6 of predicates A_{Sys} , a corresponding abstract operation $assume(choice(a, b)) :$
 7 $p_1 := choice(a_1, b_1), \dots, p_{m'} := choice(a_{m'}, b_{m'})$ has to satisfy the following
 8 implications $a \models e, b \models \neg e, a_i \models wp_{op}(p_i)$, and $b_i \models \neg wp_{op}(p_i)$ with $1 \leq i \leq$
 9 m' where a, b, a_i and b_i are logical expressions over A_{Sys} and $wp_{op}(p_i)$ is the
 10 weakest precondition of p_i with regard to op . The abstraction of operations can
 11 be performed fully-automatically. In our approach we use the prover Z3 [14] for
 12 generating abstract operations.

13 While our abstraction technique reduces the complexity induced by system
 14 variables, it preserves the original control flow. For this, the set of 2-valued
 15 predicates $A_{PC} = \{pc_i = l_i \mid i \in [1, n], l_i \in Loc_i\}$ is used that covers all
 16 locations of the system. A state s with $s(pc_i = l_i) = true$ means that in s
 17 process P_i is at control location l_i . Since a process can be only at one location
 18 at a time, $s(pc_i = l'_i)$ must be *false* for all $l'_i \neq l_i$. The overall set of predicates
 19 is $A = A_{Sys} \cup A_{PC}$.

20 The application of three-valued predicate abstraction ensures that for any
 21 state s and for any expression $choice(a, b)$ in an abstract control flow graph the
 22 following holds: $s(a) = true \Rightarrow s(b) = false$ and $s(b) = true \Rightarrow s(a) = false$.
 23 Moreover, the following equivalences hold:

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

A three-valued expression $choice(a, b)$ over A_{Sys} approximates a Boolean expression e over Var , written $choice(a, b) \preceq e$, if, and only if, a implies e and

b implies $\neg e$:

$$\text{choice}(a, b) \preceq e \equiv (a \models e) \wedge (b \models \neg e).$$

1 The three-valued approximation relation can be extended to operations as
2 follows [1]:

Definition 3 (Approximation of Operations).

Let Var be a set of variables and let A_{Sys} be a set of predicates over Var .
Moreover, let

$$op = \text{assume}(e) : v_1 := e_1, \dots, v_m := e_m$$

be a concrete operation over Var and let

$$op' = \text{assume}(\text{choice}(a, b)) : p_1 := \text{choice}(a_1, b_1), \dots, p_{m'} := \text{choice}(a_{m'}, b_{m'})$$

be an abstract operation over A_{Sys} . Then op' approximates op , written $op' \preceq op$,
if and only if

$$\text{choice}(a, b) \preceq e \wedge \bigwedge_{i=1}^{m'} (\text{choice}(a_i, b_i) \preceq wp_{op}(p_i))$$

3 where $wp_{op}(p_i)$ is the weakest precondition of p_i with respect to op .

4 An abstract system Sys' over A_{Sys} approximates a concrete system Sys over
5 Var , written $Sys' \preceq Sys$, if the systems have isomorphic control flow graphs and
6 the operations in the abstract system approximate the corresponding ones in
7 the concrete system. In the same manner we can also define the approximation
8 relation for pairs of abstract systems Sys' over A' and Sys'' over A'' with $A' \subseteq$
9 A'' .


10 Figure 4 depicts different degrees of abstraction of the concrete system Sys
11 in Figure 2. We have that $Sys' \preceq Sys'' \preceq Sys''' \preceq Sys$. For illustration: the
12 abstract operation $(y > 0) := \text{choice}((y > 0), \text{false})$ in Sys'' sets the predicate
13 $(y > 0)$ to *true* if $(y > 0)$ was *true* before, and it never sets the predicate
14 to *false*. This is a sound three-valued approximation of the concrete operation
15 $y := y + 1$ over the predicate $(y > 0)$.

16 The state space of an abstract system is defined as $S = S_{A_{Sys}} \times S_{A_{PC}}$ where
17 $S_{A_{Sys}}$ is the set of all possible valuations of the three-valued predicates in A_{Sys}
18 and $S_{A_{PC}}$ is the set of all possible valuations of the two-valued predicates in
19 A_{PC} . So far we have seen how concurrent systems can be formally represented
20 and abstracted. Next we will take a look on how bounded model checking of
21 abstracted systems is defined.

22 **3. Three-Valued Bounded Model Checking**


23 CFGs allow us to model the *control flow* of a concurrent system. The veri-
24 fication of a system additionally requires to explore a corresponding *state space*

no predicates



(a) CFGs representing an abstract system Sys' over $A_{Sys'} = \emptyset$.


$(y > 0)$: predicate where $(y > 0) = true$;



(b) CFGs representing an abstract system Sys'' over $A_{Sys''} = \{(y > 0)\}$.

$(y > 0)$: predicate where $(y > 0) = true$;

$(y > 1)$: predicate where $(y > 1) = false$;



(c) CFGs representing an abstract system Sys''' over $A_{Sys'''} = \{(y > 0), (y > 1)\}$.

Figure 4: Control flow graphs representing different degrees of abstraction of the concrete system Sys .

1 model. Since we use three-valued abstraction, we need a model that incorpo-
 2 rates the truth values *true*, *false* and *unknown*. *Three-valued Kripke structures*
 3 are models with a three-valued domain for transitions and labellings of states:

4 **Definition 4 (Three-Valued Kripke Structure).**

5 A *three-valued Kripke structure* over a set of atomic predicates A is a tuple
 6 $M = (S, I, R, L, F)$ where

- 7 • S is a finite set of states,
- 8 • $I \subseteq S$ is a set of initial states,
- 9 • $R : S \times S \rightarrow \mathfrak{3}$ is a transition function such that $\forall s \in S : \exists s' \in S$ with
 10 $R(s, s') \in \{true, unknown\}$,
- 11 • $L : S \times A \rightarrow \mathfrak{3}$ is a labelling function that associates a truth value with each
 12 atomic predicate in each state,
- 13 • $F \subseteq \mathcal{P}(\{(s, s') \mid R(s, s') \in \{true, unknown\}\})$ is a set of fairness con-
 14 straints where each constraint is a set of *non-false* transitions.

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

18 **Definition 5 (Concurrent System as Three-Valued Kripke Structure).**

19 Let $Sys = \parallel_{i=1}^n P_i$ over Var be a concurrent system abstracted over a set
 20 of system predicates A_{Sys} and given by a composite control flow graph $G =$
 21 (Loc, δ, τ) and an initial state predicate $Init$. The corresponding three-valued
 22 Kripke structure is a tuple $M = (S, I, R, L, F)$ over the set of atomic predicates
 23 $A = A_{Sys} \cup A_{PC}$ where $A_{PC} = \{(pc_i = l_i) \mid i \in [1..n], l_i \in Loc_i\}$ with

- 24 • $S := S_{A_{Sys}} \times S_{A_{PC}}$,
- 25 • $I := \{s \in S \mid s(Init) = true\}$,
- 26 • $R(s, s') := \bigvee_{i=1}^n R_i(s, s') :=$
 27 $\bigvee_{i=1}^n (\delta_i(l_i, l'_i) \wedge \bigwedge_{i' \neq i} (l_{i'} = l'_{i'}) \wedge s(choice(a, b)) \wedge \bigwedge_{j=1}^m s'(p_j) = s(choice(a_j, b_j)))$
 28 *assuming that l_i is the location of P_i with $s(pc_i = l_i) = true$,*
 29 *l'_i is the location of P_i with $s'(pc_i = l'_i) = true$ and*
 30 $\tau_i(l_i, l'_i) = assume(choice(a, b)) : p_1 := choice(a_1, b_1), \dots, p_m := choice(a_m, b_m)$,
- 31 • $L(s, p) := s(p)$ for each $p \in A$,
- 32 • $F := \{\{(s, s') \mid R_i(s, s') \in \{true, unknown\}\} \mid i \in [1..n]\}$,
 33 *i.e. for each P_i one constraint that contains all transitions caused by P_i .*
 34

1 The abstraction technique that we use guarantees that all states of the re-
 2 sulting three-valued Kripke structure have a unique labelling. Hence, we can
 3 assume that a pair of states $s, s' \in S$ with $\forall p \in A : L(s, p) = L(s', p)$ implies
 4 that $s = s'$.

5 A three-valued Kripke structure representing the state space of our example
 6 system can be defined over the predicate set $A = \{(pc_1 = 0), (pc_1 = 1), (pc_2 =$
 7 $0), (pc_2 = 1)\}$. Here we only have predicates over the control flow, but so far no
 8 predicates over the semaphore variable y . The corresponding Kripke structure
 9 is depicted in Figure 5. For the sake of simplicity, only the predicates that
 10 evaluate to *true* are shown for each state. Each transition is labelled with its
 11 truth value and the identifier of the process that causes it. As we can see, the
 12 lack of predicates over y leads to several *unknown* transitions in the model.

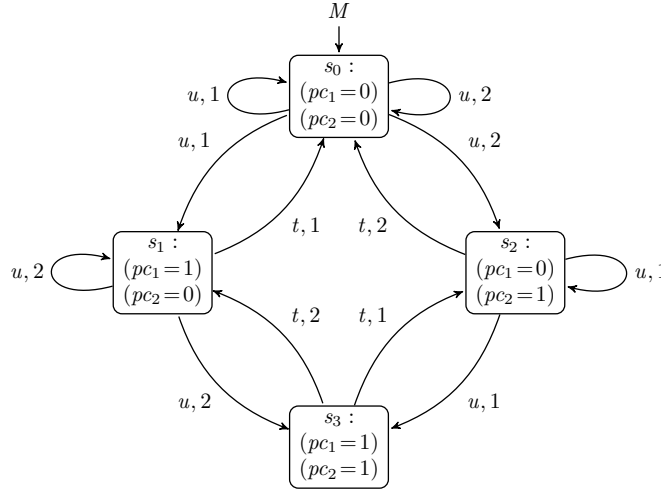


Figure 5: Three-valued Kripke structure corresponding to abstracted mutual exclusion system.

13 Computations of a modelled system correspond to paths of a Kripke struc-
 14 ture:

15 **Definition 6 (Path).**

16 Let $M = (S, I, R, L, F)$ be a three-valued Kripke structure. A *path* π of M is a
 17 sequence of states $s_0 s_1 s_2 \dots$ with $s_0 \in I$ and $\forall i : R(s_i, s_{i+1}) \in \{true, unknown\}$.
 18 π_i denotes the i -th state of π and π^i denotes the i -th suffix $s_i s_{i+1} \dots$ of π . Π_M
 19 denotes the set of all paths in M and Π_M^F denotes the set of all infinite paths
 20 in M that are fair with regard to F .

21 A fair path of our example Kripke structure must take infinitely many transi-
 22 tions associated with P_1 and infinitely many transitions associated with P_2 .
 23 While paths can be generally infinitely long, bounded model checking only looks
 24 at finite k -prefixes $\pi_0 \dots \pi_k$ of paths π where $k \in \mathbb{N}$ is the so-called bound. A
 25 finite prefix can still represent an infinite path if the prefix has a k -loop.

1 **Definition 7 (k -Loop).**

2 *Let π be a path of a three-valued Kripke structure M and let $l, k \in \mathbb{N}$ with $l \leq k$.*
 3 *Then π has a (k, l) -loop if $R(\pi_k, \pi_l) \in \{\text{true}, \text{unknown}\}$ and π is of the form*
 4 *$v \cdot w^\omega$ where $v = \pi_0 \dots \pi_{l-1}$ and $w = \pi_l \dots \pi_k$. π has a k -loop if there exists an*
 5 *$l \leq k$ such that π has a (k, l) -loop.*

6 On prefixes with or without a loop we can evaluate temporal logic properties.
 7 Here we use the *linear-time temporal logic* LTL.

Definition 8 (Syntax of LTL).

Let A be a set of atomic predicates. The syntax of LTL formulae ψ is inductively defined as

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

8 where p denotes arbitrary atomic predicates from A .

9 The temporal operator \mathbf{G} is read as *globally*, \mathbf{F} is read as *finally* (or *eventually*), and \mathbf{X} is read as *next*. For the sake of simplicity, we omit the temporal operator \mathbf{U} (*until*). Due to the extended domain of truth values in three-valued Kripke structures, the bounded evaluation of LTL formulae is based on the Kleene logic \mathcal{K}_3 (compare Section 2). For the bounded evaluation of LTL formulae on paths of three-valued Kripke structures we have to distinguish between prefixes *with* and *without* a k -loop.

16 **Definition 9 (Three-Valued Bounded Evaluation of LTL).**

17 *Let $M = (S, I, R, L, F)$ over A be a three-valued Kripke structure. Moreover,*
 18 *let $k \in \mathbb{N}$ and let π be a path of M with a k -loop. Then the bounded evaluation*
 19 *of an LTL formula ψ on the k -prefix of π , written $[\pi \models \psi]_k^i$ where $i \leq k$ denotes*
 20 *the current position along the path, is inductively defined as follows:*

$$\begin{aligned} [\pi \models p]_k^i &\equiv L(\pi_i, p) \\ [\pi \models \neg p]_k^i &\equiv \neg L(\pi_i, p) \\ [\pi \models \psi \vee \psi']_k^i &\equiv [\pi \models \psi]_k^i \vee [\pi \models \psi']_k^i \\ [\pi \models \psi \wedge \psi']_k^i &\equiv [\pi \models \psi]_k^i \wedge [\pi \models \psi']_k^i \\ [\pi \models \mathbf{G}\psi]_k^i &\equiv \bigwedge_{j \geq i} ([\pi \models \psi]_k^j \wedge R(\pi_j, \pi_{j+1})) \\ [\pi \models \mathbf{F}\psi]_k^i &\equiv \bigvee_{j \geq i} ([\pi \models \psi]_k^j \wedge \bigwedge_{j'=i}^{j-1} R(\pi_j, \pi_{j+1})) \\ [\pi \models \mathbf{X}\psi]_k^i &\equiv [\pi \models \psi]_k^{i+1} \wedge R(\pi_i, \pi_{i+1}) \end{aligned}$$

Let π be a path of M without a k -loop. Then the bounded evaluation of an LTL formula ψ on the k -prefix of π is defined as follows:

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

21 *The other cases are identical to the case where π has a k -loop.*

1 Note that in contrast to standard LTL, the *three-valued* evaluation of LTL for-
 2 mularae requires to take the transition relation into account: Transitions in three-
 3 valued models are either *true* or *unknown* and the value of a transition may
 4 affect the result of model checking. Moreover, note that in the *no k-loop* case of
 5 bounded model checking any formula of the form $\mathbf{G} \psi$ always evaluates to *false*
 6 because ψ might not hold at the successor position $k + 1$ of the considered path.

For convenience, we introduce the following abbreviation that allows for a
 shorter representation of certain temporal logic properties. Let ψ be an arbitrary
 LTL formula then:

$$\psi_i := \mathbf{X}^i \psi = \underbrace{\mathbf{X} \dots \mathbf{X}}_{i\text{-times}} \psi.$$

7 Hence, ψ_i states that ψ has to hold for the i -suffix of a considered path.

8 The bounded evaluation of temporal logic properties on entire three-valued
 9 Kripke structures is what is known as *three-valued bounded model checking* with
 10 the possible outcomes *true*, *false* and *unknown*. Here we distinguish between
 11 universal and existential model checking.

Definition 10 (Three-Valued Bounded Model Checking).

Let $M = (S, I, R, L, F)$ be a three-valued Kripke structure over A , let $k \in \mathbb{N}$ and
 let ψ be an LTL formula. The corresponding universal model checking problem
 is

$${}_A[M, I \models_{\forall}^F \psi]_k = \bigwedge_{\pi \in \Pi_M^F, \pi_0 \in I} [\pi \models \psi]_k^0$$

and the corresponding existential model checking problem is

$${}_A[M, I \models_{\exists}^F \psi]_k = \bigvee_{\pi \in \Pi_M^F, \pi_0 \in I} [\pi \models \psi]_k^0.$$

12 Note that model checking *without* fairness constraints is a special case of
 13 model checking under fairness with $F = \emptyset$. For certain verification tasks fairness
 14 is not relevant and thus can be ignored in order to avoid an unnecessary growth
 15 of complexity. If we neglect fairness, we will simply write $M = (S, I, R, L)$ for
 16 a Kripke structure and ${}_A[M, I \models_Q \psi]_k$ with $Q \in \{\forall, \exists\}$ for a model checking
 17 problem.

18 In practice, universal model checking is of major interest as it allows to show
 19 that *all* computations of a modelled system satisfy certain temporal logic prop-
 20 erties. Certain model checking techniques such as satisfiability-based bounded
 21 model checking are only defined for the existential case. However, universal
 22 model checking can always be transformed into existential model checking based
 23 on the equation ${}_A[M, I \models_{\forall}^F \psi]_k = \neg {}_A[M, I \models_{\exists}^F \neg \psi]_k$, which also makes exis-
 24 tential approaches generally applicable. In this work we will particularly focus
 25 on the existential case. In [3] Lemma 1 has been proven:

26 **Lemma 1**

27 *Let Sys over Var be a concurrent system and let A_a and A_r be sets of pred-*

1 indicates over Var with $A_a \subset A_r$. Let Sys_a over A_a and Sys_r over A_r be ab-
 2 stract systems with $Sys_a \preceq Sys$, $Sys_r \preceq Sys$ and $Sys_a \preceq Sys_r$. Moreover,
 3 let $M_a = (S_a, I_a, R_a, L_a, F_a)$ be the three-valued Kripke structure modelling the
 4 state space of Sys_a , let $M_r = (S_r, I_r, R_r, L_r, F_r)$ be the structure modelling the
 5 state space of Sys_r , let ψ be an LTL formula and $k \in \mathbb{N}$ be a bound. Then the
 6 following holds:

- 7 1. $A_a[M_a, I_a \models_{\exists}^F \psi]_k = true \Rightarrow A_r[M_r, I_r \models_{\exists}^F \psi]_k = true$
- 8 2. $A_a[M_a, I_a \models_{\exists}^F \psi]_k = false \Rightarrow A_r[M_r, I_r \models_{\exists}^F \psi]_k = false$

9 Hence, all definite (*true* and *false*) bounded model checking results obtained
 10 under three-valued abstraction of Sys over a predicate set A_a can be transferred
 11 to any refined abstraction over an extended predicate set A_r with $A_a \subset A_r$.
 12 Since the most refined abstraction represents the concrete state space of Sys ,
 13 definite results can be also transferred to the concrete system. An *unknown*
 14 result indicates that the current level of abstraction is too coarse and further
 15 predicates need to be added.

Bounded model checking is inherently incomplete as it only considers paths
 up to a length k . The existence of a k -prefix that satisfies a formula ψ allows
 us to conclude that ψ is also existentially satisfied in the unbounded case. But
 the non-existence of such a k -prefix does not allow us to conclude that ψ is not
 satisfied in the unbounded case. A straightforward but typically impracticable
 way of making bounded model checking complete is to iterate over all possible
 bounds up to a completeness threshold ct [9] with:

$$A[M, I \models_{\exists}^F \psi]_{ct} = A[M, I \models_{\exists}^F \psi]_{\infty}$$

16 Two types of temporal logic properties are of particular interest in model
 17 checking concurrent software systems: *safety* and *liveness*. Checking safety
 18 properties such as mutual exclusion can be done via reachability analysis for
 19 which many efficient techniques exist. In particular, only loop-free paths need
 20 to be considered and fairness can be neglected for safety model checking. How-
 21 ever, many requirements of concurrent systems cannot be formulated as safety
 22 properties. For instance, the requirement that the processes of a system make
 23 continuous progress refers to a liveness property. Checking liveness requires
 24 the consideration of fairness and the exploration of paths with loops, which is
 25 significantly more complex than reachability analysis.

In the following we exemplify one safety and one liveness model checking
 problem that we will consider throughout this work. For our running example,
 we can define the universal safety model checking problem $A[M, I \models_{\forall} \mathbf{G} safe]_k$
 without fairness constraints where $safe = \neg(pc_1 = 1) \vee \neg(pc_2 = 1)$. This charac-
 terises the mutual exclusion requirement that globally at most one process shall
 be at the critical location 1 at the same time. A *counterexample* to this require-
 ment can be always given by a finite loop-free prefix that reaches a state where
 both processes are at the critical location. The complementary existential model
 checking problem is $A[M, I \models_{\exists} \mathbf{F} \neg safe]_k$ where we check for the existence of a

path prefix that violates mutual exclusion. In existential model checking we call a path that satisfies the temporal logic property a *witness*. A counterexample for a universal problem always corresponds to a witness for the complementary existential problem and vice versa. A completeness threshold for model checking safety properties is the *diameter* of the underlying Kripke structure, i.e. the longest distance between any two states. In incremental bounded model checking that iterates over k from 0 up to a completeness threshold, we can assume that in all previous iterations $k' < k$ no witness has been detected (otherwise model checking would have terminated). This allows us to slightly re-formulate the property to be checked. The formula

$$\mathbf{O}^k \neg safe \equiv safe_0 \wedge safe_1 \wedge \dots \wedge safe_{k-1} \wedge \neg safe_k$$

1 where $safe_i = \mathbf{X}^i safe$, characterises that $safe$ is *only* (\mathbf{O}) violated in the k -th
 2 state of a k -prefix. Model checking such a conjunctive formula is generally more
 3 efficient than checking a disjunctive formula, since the conjunctive form involves
 4 a stronger restriction of the search space.

5 In the three-valued scenario, we have to distinguish between two kinds of
 6 witnesses, *definite* ones when the model checking result is *true* and *unconfirmed*
 7 ones when the result is *unknown*.

Definition 11 (Definite Witness for Safety).

Let ${}_A[M, I \models_{\exists} \mathbf{O}^k \neg safe]_k$ be a three-valued bounded model checking problem with regard to safety where $safe$ is a propositional logic expression over A . A *definite witness* for $\mathbf{O}^k \neg safe$ is a prefix $\omega = \pi_0 \dots \pi_k$ of a path $\pi \in \Pi_M$, $\pi_0 \in I$ with

$$\pi_k(safe) = false \text{ and } \forall 0 \leq i < k : R(\pi_i, \pi_{i+1}) = true$$

8 A definite witness implies that a safety violation has been detected, and thus,
 9 no further model checking runs are required. In our example Kripke structure
 10 from Figure 5 there exists no definite witness for $\mathbf{O}^k \neg safe$.

11 An *unknown* result in three-valued bounded model checking indicates that
 12 there exists a path $\pi \in \Pi_M$ such that its k -prefix $\omega = \pi_0 \dots \pi_k$ is an *unconfirmed*
 13 *witness* for the safety formula $\mathbf{O}^k \neg safe$.

Definition 12 (Unconfirmed Witness for Safety).

Let ${}_A[M, I \models_{\exists} \mathbf{O}^k \neg safe]_k$ be a three-valued bounded model checking problem with regard to safety where $safe$ is a propositional logic expression over A . An *unconfirmed witness* for $\mathbf{O}^k \neg safe$ is a prefix $\omega = \pi_0 \dots \pi_k$ of a path $\pi \in \Pi_M$, $\pi_0 \in I$ with either

$$\begin{aligned} \pi_k(safe) &= unknown, \text{ or} \\ \pi_k(safe) &= false \text{ and } \exists 0 \leq i < k \text{ with } R(\pi_i, \pi_{i+1}) = unknown \end{aligned}$$

14 For our Kripke structure from Figure 5 the path prefix $s_0 s_1 s_3$ is an uncon-
 15 firmed witness for $\mathbf{O}^k \neg safe$.

1 A fair universal liveness model checking problem for our running example is
 2 $_A[M, I \models_{\forall}^F \mathbf{GF} \textit{progress}]_k$ where $\textit{progress} = (pc_1 = 1) \vee (pc_2 = 1)$. This char-
 3 acterises the requirement that in a fair computation always eventually some
 4 process will reach the critical location. A counterexample would be a path
 5 that reaches an infinite loop where each process continues to execute operations
 6 but no process will be in the critical location ever again. The complementary
 7 existential model checking problem is $_A[M, I \models_{\exists}^F \mathbf{FG} \neg \textit{progress}]_k$. A witness
 8 for the existential problem corresponds to a counterexample to the comple-
 9 mentary universal problem. Again, we can distinguish between definite and
 10 unconfirmed witnesses for liveness. In our example, the prefix $s_0 \xrightarrow{u,1} s_0 \xrightarrow{u,2} s_0$
 11 which has a fair (2,0)-loop is an unconfirmed witness for the liveness formula
 12 $\mathbf{FG} \neg \textit{progress}$. Thus, the corresponding three-valued bounded model checking
 13 result is *unknown*.

14 So far, our three-valued bounded model checking approach fails for our run-
 15 ning example because of the following two reasons. Firstly, we are only able to
 16 obtain *unknown* results. Hence, three-valued abstraction refinement is necessary
 17 for which we will present a technique in Section 6. Secondly, bounded model
 18 checking is incomplete and completeness thresholds for liveness properties are
 19 hard to compute and typically impracticable for efficient verification. We ap-
 20 proach this limitation in two ways. We adopt the *state recording* technique of
 21 [13] that allows to translate liveness model checking problems into safety prob-
 22 lems, for which smaller completeness thresholds exist. Moreover, we employ the
 23 *k-induction* technique [11], which allows to make safety bounded model check-
 24 ing complete without the need of a predetermined threshold. We start with the
 25 reduction of liveness to safety in the subsequent section.

26 4. From Liveness to Safety via State Recording

27 In this section, we review the state recording technique originally introduced
 28 in [13] which allows to reduce model checking liveness to model checking safety.
 29 In particular, we show that state recording can be implemented based on a
 30 transformation of our concurrent systems to be verified. For the purpose of an
 31 easier understanding, we divide the transformation into two steps: a first step
 32 that enables *loop detection* and a second step that allows for model checking
 33 *liveness under fairness*. The first step translates the input system into a state
 34 recording system for the purpose of loop detection. A looped path corresponds
 35 to a finite prefix where the last state is identical to an arbitrary previous state.
 36 Hence, the search for a loop can be performed by "recording" an arbitrary
 37 state along an explored path and checking whether this state will be reached
 38 again. The recording creates a copy of the current state. Thus, a state recording
 39 system introduces a copy p^c of each original predicate p in order to represent
 40 the recorded state.

41 Definition 13 (State Recording System).

42 Let $Sys = \parallel_{i=1}^n P_i$ be a concurrent system abstracted over a set of system
 43 predicates A_{Sys} and given by n control flow graphs $G_i = (Loc_i, \delta_i, \tau_i)$. Moreover,

1 let $A = A_{Sys} \cup A_{PC}$ with $A_{PC} = \{(pc_i = l_i) \mid i \in [1..n], l_i \in Loc_i\}$ be the
 2 overall set of predicates and let $Init$ be the initial state predicate of Sys . Then
 3 the corresponding *state recoding system* Sys^{sr} is defined over the predicate set
 4 A^{sr} with initial state predicate $Init^{sr}$ and given by the control flow graphs
 5 $G_i^{sr} = (Loc_i, \delta_i, \tau_i^{sr})$ where

- 6 • $A^{sr} := A \cup A^c \cup Aux$
 7 where $A^c = \{p^c \mid p \in A\}$ is a set of copies of the original predicates in A
 8 and $Aux = \{record, recorded\}$ is a set of Boolean auxiliary predicates,
- 9 • $Init^{sr} := Init \wedge \bigwedge_{p^c \in A^c} (p^c \leftrightarrow p) \wedge \neg recorded,$
- $\forall (l_i, l'_i) \in \delta_i : \tau_i^{sr}(l_i, l'_i) :=$
 - $\tau_i(l_i, l'_i) \circ [record := *]$ (1)
 - $\circ [recorded := record \vee recorded]$ (2)
 - $\circ [\forall p^c \in A^c : p^c := (firstRecord \rightarrow p) \wedge (\neg firstRecord \rightarrow p^c)]$ (3)

10 where $*$ denotes the non-deterministic choice between *true* and *false*
 11 and $firstRecord = record \wedge \neg recorded$.

12

13 Hence, the state recoding translation adds a copy p^c of each original pred-
 14 icate $p \in A$ and two Boolean auxiliary predicates *record* and *recorded* to the
 15 system. The copies receive the same initialisation as the original predicates.
 16 While there is no restriction with regard to the initialisation of *record*, the pred-
 17 icate *recorded* is initially set to *false*. The control flow and the operations of
 18 the original system are preserved, but each original operation gets extended by
 19 assignments to the new predicates. Remember that the list of assignments asso-
 20 ciated with an operation is executed in an *atomic* manner. Thus, if an operation
 21 changes the value of *record* from *false* to *true*, then on the right-hand side of the
 22 assignment $recorded := record \vee recorded$ the predicate *record* is still evaluated
 23 with *false*.

24 The extension allows to reduce the detection of a looped path to a reach-
 25 ability problem on the level of the corresponding Kripke structure. For now
 26 we can ignore fairness. Let $M = (S, I, R, L)$ over A be the Kripke struc-
 27 ture modelling the original system Sys and let $M^{sr} = (S^{sr}, I^{sr}, R^{sr}, L^{sr})$ over
 28 A^{sr} be the Kripke structure modelling Sys^{sr} . The set of states of M^{sr} is
 29 $S^{sr} = S_A \times S_{A^c} \times S_{Aux}$ and each state is a triple $s^{sr} = \langle s, s^c, s^{aux} \rangle$ where
 30 s refers to the state of Sys . M^{sr} still comprises all the behaviour of M with
 31 regard to original predicate set A , i.e. for each path π in M there exists a path
 32 π^{sr} in M^{sr} with $\forall i \geq 0 : \forall p \in A : L(\pi_i, p) = L(\pi_i^{sr}, p)$ and vice versa. The
 33 new predicates and extended operations add a state recoding mechanism to the
 34 modelled system, which works as follows. Each execution of an extended opera-
 35 tion along a computational path now sets the value of the new predicate *record*

1 non-deterministically to either *true* or *false* (1). When *record* evaluates to *true*
2 for the first time, this indicates that the current state of the original system is
3 selected to be recorded. In order to ensure that a recording only happens once,
4 the predicate *recorded* is used. While initially evaluated with *false*, *recorded* is
5 set to *true* after a *record* has been triggered and then remains *true* in all further
6 computational steps (2). The actual state recording now has the condition that
7 $firstRecord = record \wedge \neg recorded$ holds. The recording happens by assigning the
8 predicate copies p^c to the current values of the original predicates p (3).

9 While classical state space exploration is memoryless, state recording allows
10 to check whether a previously recorded state can be reached again, which reduces
11 loop detection to a reachability problem. Let $looped = recorded \wedge \bigwedge_{p \in A} (p \leftrightarrow p^c)$,
12 then the model checking problem $A^{sr} [M^{sr}, I^{sr} \models_{\exists} \mathbf{F} looped]_k$ is equivalent to
13 the question of whether there exists a looped path of length $k - 1$ in the original
14 system. Note that if the state recording model checking problem has the bound
15 k , then only paths of length $k - 1$ will be considered. This is due to the fact that
16 under state recording a detected loop indicates that the state s_k is identical to
17 some previously reached state s_l . Hence, the associated loop is actually from
18 position $k - 1$ to l .

19 The prefix of a path π^{sr} depicted in Figure 6 illustrates reachability-based
20 loop detection in a state recording Kripke structure M^{sr} corresponding to some
21 original Kripke structure M .




Figure 6: Witness for a looped path.

22 As we can see, π^{sr} comprises the looped path $\pi = s_0 s_1 \cdot (s_2 s_3 s_4)^\omega$ of some
23 original model. The execution of the operation associated with the transition
24 from π_1^{sr} to π_2^{sr} sets *record* to *true* for the first time. Hence, the original state
25 s_2 is the one chosen to be recorded. Note that this transition only indicates
26 the *choice* of the state to be recorded. The *actual recording* happens by the
27 execution of the operation associated with the subsequent transition, as only in
28 π_2^{sr} the condition *firstRecord* for state recording evaluates *true*. In the overall
29 state π_3^{sr} the state copy s^c has been set to s_2 , while the state of the original
30 system has already changed to s_3 . In π_4^{sr} the predicate *record* is *true* again, but
31 since *firstRecord* does not hold any more the state copy will not be overwritten
32 in the subsequent state. The path finally reaches the state π_5^{sr} where $s = s^c$ as
33 well as *recorded* holds. Thus, we can conclude that the prefix of π^{sr} characterises
34 a looped path of the original model. Since we use *existential* model checking,
35 *any* reachable original state will be considered for recording. Hence, if the state
36 recording model checking problem $A^{sr} [M^{sr}, I^{sr} \models_{\exists} \mathbf{F} looped]_k$ returns *false*, then
37 no looped path of length $k - 1$ exists in the original model.

1 So far, the state recording approach only allows to detect looped paths but
 2 not to perform liveness model checking. In Definition 14, we further extend
 3 state recording systems in order to enable model checking liveness properties
 4 under fairness assumptions.

5 **Definition 14 (State Recording System with Liveness Extension).**

6 Let Sys^{sr} over $A^{sr} = A \cup A^c \cup Aux$ be a state recording system given by n
 7 control flow graphs $G_i^{sr} = (Loc_i, \delta_i, \tau_i^{sr})$. Moreover, let $\mathbf{FG} \neg progress$ be a
 8 temporal logic formula characterising a liveness violation where $progress$ is a
 9 predicate expression over A . Then the corresponding *state recoding system with*
 10 *liveness extension* Sys^{le} is defined over the predicate set A^{le} with initial state
 11 predicate $Init^{le}$ and given by the control flow graphs $G_i^{le} = (Loc_i, \delta_i, \tau_i^{le})$ where

- 12 • $A^{le} := A^{sr} \cup \mathbb{F} \cup \{live\}$
 13 where $\mathbb{F} = \{fair_i \mid i \in [1..n]\}$ is a set of Boolean fairness predicates
 14 and $live$ is a Boolean auxiliary predicate,
- 15 • $Init^{le} := Init^{sr} \wedge \bigwedge_{i \in [1..n]} \neg fair_i \wedge \neg live$,
- $\forall (l_i, l'_i) \in \delta_i : \tau_i^{le}(l_i, l'_i) :=$

$$\tau_i^{sr}(l_i, l'_i) \circ [fair_i := loopStarted] \quad (1)$$

$$\circ [live := live \vee (loopStarted \wedge progress)] \quad (2)$$

16 where $loopStarted = record \vee recorded$.

17

18 For each process P_i the liveness extension adds a predicate $fair_i$ to the
 19 system. Moreover, a predicate $live$ is added. All new predicates are initially
 20 assigned to *false*. The predicate $fair_i$ is used to indicate whether a computation
 21 is fair with regard to P_i in the sense that the process infinitely often executes
 22 an operation. On looped execution paths this is the case if P_i executes an
 23 operation after the loop has started. Hence, an operation $\tau_i^{le}(l_i, l'_i)$ by P_i sets
 24 $fair_i$ to *true* if the condition $loopStarted = record \vee recorded$ is satisfied (1). The
 25 predicate expression $loopStarted$ is used to indicate whether the starting point of
 26 a *potential* loop has already been reached and recorded. (If the recorded starting
 27 point belongs to an *actual* loop is determined by solving the model checking
 28 problem defined in Corollary 1.) Once a computation of a state recording system
 29 leads to a state where $loopStarted$ holds, this expression will also evaluate to
 30 *true* in all future states resulting from the computation (Definition 13). Thus, a
 31 predicate $fair_i$ will also evaluate to *true* forever once it has been *true* for the first
 32 time along a computational path. The predicate $live$ is used to indicate whether
 33 $progress$ holds infinitely often, i.e. at some point after the loop has started (2).
 34 By combining the results of [13] with our definitions, we get Corollary 1:

Corollary 1

Let Sys be a concurrent system and M be the Kripke structure modelling the

state space of Sys abstracted over A . Let $progress$ be a predicate expression over A . Moreover, let Sys^{le} be the state recording system with liveness extension corresponding to Sys and M^{le} be the Kripke structure modelling the state space of Sys^{le} over A^{le} . Then the following holds:

$$A[M, I \models_{\exists}^F \mathbf{FG} \neg progress]_k \Leftrightarrow A^{le}[M^{le}, I^{le} \models_{\exists} \mathbf{F}(looped \wedge fair \wedge \neg live)]_{k+1}$$

1 where $looped = recorded \wedge \bigwedge_{p \in A}(p \leftrightarrow p^c)$ and $fair = \bigwedge_{i=1}^n fair_i$.

2 Hence, liveness model checking under fairness can be reduced to simple safety
3 model checking.

4 The prefix of a path π^{le} depicted in Figure 7 illustrates reachability-based
5 detection of liveness violations in a state recording Kripke structure with liveness
6 extension M^{le} corresponding to some original Kripke structure M .

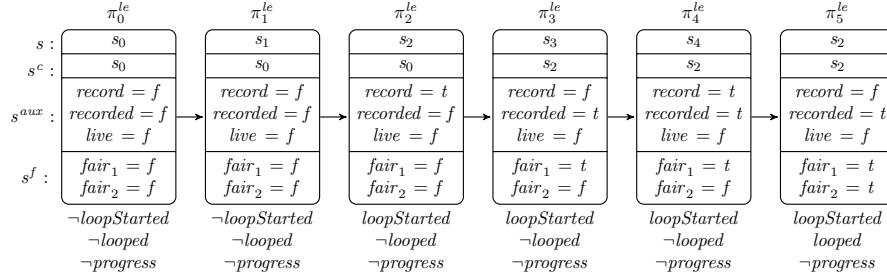


Figure 7: Witness for liveness violation.

7 As we can see, π^{le} comprises the looped path $\pi = s_0 s_1 \cdot (s_2 s_3 s_4)^\omega$ of some
8 original model M . The reachability of the state π_5^{le} allows us to conclude that
9 there exists a looped path satisfying fairness but violating liveness. Similar as
10 proposed in Section 2, an incremental approach to bounded model checking
11 allows us to re-formulate the property to be checked to $\mathbf{O}^k(looped \wedge fair \wedge \neg live)$
12 which gives us a more restricted search space.

13 With the state recording translation we are able to detect liveness violations
14 in our abstracted system via safety model checking. The price to be paid for
15 the reduction from liveness to safety is a quadratic increase of the number of
16 states in the state recording Kripke structure, i.e. $|S^{le}| = \mathcal{O}(|S|^2)$ [13]. How-
17 ever, for safety model checking many efficient algorithms and typically smaller
18 completeness thresholds exist. While the determination of tight thresholds for
19 checking safety is also challenging, an alternative approach to the completeness
20 of bounded model checking is k -induction [11]. We review k -induction in Section
21 5 and show how we have integrated it into our abstraction-based three-valued
22 bounded model checking technique in Section 6.

1 5. Unbounded Model Checking via k -Induction

In the previous section we have shown that liveness model checking problems can be translated into safety model checking problems. Hence, we can assume that all our verification problems to be solved are of the form

$$A[M, I \models_{\forall} \mathbf{G} \text{ safe}]_{\infty}$$

where *safe* is an arbitrary predicate expression over the set of atomic predicates A . This unbounded model checking problem requires the consideration of all infinite paths of the model in order to prove that no property violations exist. The k -induction approach [11] allows to reduce an unbounded safety model checking problem into two bounded model checking problems: In the *base case* it is checked whether all k -prefixes of paths starting in an initial state $s \in I$ of M are *safe*, i.e.

$$A[M, I \models_{\forall} \bigwedge_{i=0}^k \text{ safe}_i]_k.$$

In the *inductive step* it is checked whether, assuming a path prefix of k *safe* states, also any successor state is *safe*, i.e.

$$A[M, S \models_{\forall} (\bigwedge_{i=0}^k \text{ safe}_i) \rightarrow \text{ safe}_{k+1}]_{k+1}.$$

2 Note that in the inductive step the considered prefixes can start in an arbitrary state $s \in S$. As proven in [11], if there exists a k for which the base case
 3 fails, then the property $\mathbf{G} \text{ safe}$ is violated. And, if there exists a k for which
 4 both the base case and the inductive step hold, then the property $\mathbf{G} \text{ safe}$ holds
 5 for the model. k -induction is typically performed *incrementally* with regard to
 6 the bound. Thus, if the base case holds but the inductive step fails then k gets
 7 incremented and the model checking problems corresponding to the increased
 8 bound are solved.
 9

The universal problems above refer to the safety of *all* paths. However, model checking via satisfiability solving is based on the existential case. We have already seen that each universal model checking problem can be transformed into a complementary existential problem referring to the existence of an unsafe path:

$$A[M, I \models_{\exists} \mathbf{F} \neg \text{ safe}]_{\infty}$$

The corresponding existential base case is

$$A[M, I \models_{\exists} \bigvee_{i=0}^k \neg \text{ safe}_i]_k$$

and the existential inductive step is

$$A[M, S \models_{\exists} (\bigwedge_{i=0}^k \text{ safe}_i) \wedge \neg \text{ safe}_{k+1}]_{k+1}.$$

Now an unsafe path exists if for some k the existential base case holds, whereas

all paths are safe if both the existential base case and the existential inductive step fail. We want to follow an incremental approach with regard to the bound. Thus, when checking the base case for some k we can assume that all shorter base cases have already been proven to be safe, and we can add these facts as constraints to the problem to be solved:

$$A[base]_k := A[M, I \models \exists \underbrace{\left(\bigwedge_{i=0}^{k-1} safe_i \right) \wedge \neg safe_k}_k]_k = \mathbf{O}^k \neg safe$$

1 This strengthening of the temporal logic formula to be checked involves a re-
 2 striction of the state space to be explored, and thus, allows for model checking
 3 with an improved efficiency.

In order to make the k -induction approach *complete*, i.e. terminating for finite-state systems, it is necessary to restrict the inductive step to *loop-free* computations [11]. This gives us a slightly revised inductive step:

$$A[step]_{k+1} := A[M, S \models \exists \underbrace{\left(\bigwedge_{i=0}^k safe_i \right) \wedge \neg safe_{k+1} \wedge loopFree_{0,k+1}}_{k+1}]_{k+1} = \mathbf{O}^{k+1} \neg safe$$

4 where $loopFree_{0,k+1} = \bigwedge_{0 \leq i < j \leq k+1} \left(\bigvee_{p \in AP} ((p_i \wedge \neg p_j) \vee (\neg p_i \wedge p_j)) \right)$.

5 Adding the loop-free constraint to the inductive step is an implicit way
 6 of determining whether the current bound is a completeness threshold of the
 7 model checking problem to be solved. Hence, a threshold does not have to be
 8 explicitly computed in the k -induction approach. Algorithm 1 illustrates the
 9 basic principle of incremental k -induction:

Algorithm 1: Incremental k -induction.

```

1 for  $k = 0$  to  $\infty$  do
2   if  $(A[base]_k = true)$  then
3     return "safety property violated"
4   if  $(A[base]_k = false$  and  $A[step]_{k+1} = false)$  then
5     return "safety property holds"

```

10 Hence, incremental k -induction iterates over the bound until either a prop-
 11 erty violation can be detected or until the bound is sufficiently large to conclude
 12 that the property holds. In [11] it has been proven that the k -induction approach
 13 yields the correct unbounded model checking result.

14 However, Algorithm 1 does not take into account that our model checking
 15 problems have a *three-valued* domain, i.e. solving the base case or the inductive
 16 step might yield *unknown*. In order to handle this case, we combine k -induction
 17 with three-valued abstraction refinement, which we discuss in the next section.

1 6. Basic Three-Valued Abstraction Refinement

2 Solving a three-valued bounded model checking problem $_A[M, I \models_{\exists} \psi]_k$,
 3 where the temporal logic formula ψ characterises the violation of a safety prop-
 4 erty, has the possible outcomes *false*, *true* and *unknown*. A *false* result indicates
 5 that there exists no k -prefix in M that is a witness for ψ . A *true* result indicates
 6 that there exists a k -prefix ω in M that is a *definite witness* for ψ (Definition
 7 11). An *unknown* result indicates that there exists a k -prefix ω in M that is
 8 an *unconfirmed witness* for ψ (Definition 12). Model checking tools are typi-
 9 cally not only capable of returning the result of the input problem, but also the
 10 corresponding witness in case the result is *true* or *unknown* [3].

11 Our example Kripke structure in Figure 5 abstracts the mutual exclusion
 12 system over the predicate set $A = \{(pc_1 = 0), (pc_1 = 1), (pc_2 = 0), (pc_2 = 1)\}$.
 13 Assuming a bound of $k = 2$ and the temporal logic formula $\psi = \mathbf{O}^k \neg safe$ with
 14 $safe = \neg((pc_1 = 1) \wedge (pc_2 = 1))$, three-valued bounded model checking will return
 15 *unknown* along with the corresponding unconfirmed witness shown in Figure 8.

$$\omega = s_0 s_1 s_3 = (0, 0) \xrightarrow{u} (1, 0) \xrightarrow{u} (1, 1)$$

Figure 8: Unconfirmed witness.

16
 17 In this representation, a tuple (l_1, l_2) denotes a state where process P_1 is
 18 at location l_1 and P_2 is at location l_2 . Moreover, \xrightarrow{u} denotes an *unknown*
 19 transition between states. The witness ω is *unconfirmed* because it reaches the
 20 abstract state $(1, 1)$ where safety is definitely violated, but *unknown* transi-
 21 tions are taken in order to reach this state. An unconfirmed witness implies that the
 22 current level of abstraction, characterised by the predicate set A , is too coarse
 23 for a definite model checking result. In this case, refinement in the sense of
 24 extending the set A is required. In [3] we defined the function *analyseWitness*.
 25 It takes an unconfirmed witness $\omega = \pi_0 \dots \pi_k$ for a model checking problem
 26 $_A[M, I \models_{\exists} \mathbf{O}^k safe]_k$ as an input where M models the state space of a system
 27 Sys abstracted over A . Based on the concrete operations of Sys and the weakest
 28 precondition calculus *analyseWitness* derives suitable predicates for refinement,
 29 which works as follows:

- 30 1. Let $p = safe$. One possible reason for $\omega = \pi_0 \dots \pi_k$ being an unconfirmed
 31 witness is that $\pi_k(p) = unknown$.
 - 32 • Determine the largest index i , $0 \leq i < k$, with $\pi_i(p) \neq unknown$.
 - 33 • Determine the concrete operation op of Sys that is associated with
 34 the transition (π_i, π_{i+1}) .
 - 35 • Let $p' = wp_{op}(p)$ be the weakest precondition of p with respect to
 36 op . If $p' \notin A$ then return p' and \bar{p}' as predicates for refinement. Else
 37 repeat Step 1 for p' .

- 1 2. The other possible reason for $\omega = \pi_0 \dots \pi_k$ being unconfirmed is a transi-
2 tion (π_i, π_{i+1}) with $R(\pi_i, \pi_{i+1}) = \textit{unknown}$ for some index $0 \leq i < k$.
- 3 • Determine the concrete operation op of Sys that is associated with
4 the transition (π_i, π_{i+1}) .
 - 5 • Let $\textit{assume}(e)$ be the guard of op . If $e \notin A$ then return e and \bar{e} as
6 predicates for refinement. Else repeat Step 1 for e .

7 For our current example and the unconfirmed witness $\omega = s_0 s_1 s_3$ we get
8 $\textit{analyseWitness}(s_0 s_1 s_3) = \{(y > 0), (\overline{y > 0})\}$. This is the set containing the guard
9 $(y > 0)$, associated with the *acquire* operation in the mutual exclusion system, as
10 well as its complement $(\overline{y > 0})$ ¹. While a detailed description of *analyseWitness*
11 can be found in [3], we subsequently focus on the integration of this function into
12 a fully-automatic abstraction refinement algorithm. Algorithm 2 defines a basic
13 three-valued abstraction refinement algorithm *AR* that utilises *analyseWitness*.

Algorithm 2: $AR(A[M, I \models_{\exists} \psi]_k)$

```

1 loop forever do /*refinement loop*/
2   if  $A[M, I \models_{\exists} \psi]_k = \textit{false}$  then
3     return false, ‘no witness for safety violation of length  $k$  exists’
4   if  $A[M, I \models_{\exists} \psi]_k = \textit{true}$  and definite witness  $\omega$  then
5     return true, ‘ $\omega$  is a definite witness for safety violation’
6   if  $A[M, I \models_{\exists} \psi]_k = \textit{unknown}$  and unconfirmed witness  $\omega$  then
7      $A := A \cup \textit{analyseWitness}(\omega)$ 

```

14 *AR* takes a three-valued bounded model checking problem $A[M, I \models_{\exists} \psi]_k$
15 as an input. We assume that M is the state space model corresponding to the
16 system to be verified abstracted over the initial predicate set A . Within the
17 algorithm the model checking problem is solved. If the outcome is *true* or *false*
18 the algorithm terminates and returns the definite result. In case of an *unknown*
19 result, an unconfirmed witness ω is generated and new predicates are derived
20 via *analyseWitness*. The predicate set of the next iteration is defined by the
21 predicates of the current iteration joined with the newly derived predicates.
22 Now the steps of model checking the Kripke structure corresponding to the
23 extended predicate set and deriving new refinement predicates are repeated until
24 a definite result can be obtained. The termination of *AR* is guaranteed for finite-
25 state systems. In [3] we showed that the result of *AR* correctly characterises
26 the computational behaviour of the system modelled by M . For our running

¹In contrast to the Boolean predicates over the control flow, predicates over system variables have a *three-valued-valued* domain as they may evaluate to *unknown* due to abstraction. In order to enable the later reduction of three-valued bounded model checking to Boolean satisfiability, there must be a complementary predicate \bar{p} with $\bar{p} \equiv \neg p$ for each predicate p over system variables [18].

1 example with the fixed bound $k = 2$ *AR* terminates with a *false* result after one
 2 refinement iteration that adds the predicates $(y > 0)$ and $(\overline{y} > \overline{0})$. Thus, there
 3 does not exist a witness of length 2 that violates mutual exclusion.

4 *AR* can be easily integrated into the algorithm for incremental k -induction.
 5 For the integration it is advisable to use distinct predicate sets for the base case
 6 and for the inductive step. In this way both problems can be independently
 7 refined, which is typically more efficient than always forcing the same level of
 8 abstraction, i.e. identical predicate sets for both problems. An approach that
 9 combines incremental k -induction with abstraction refinement will process our
 10 running example as follows: In bound iteration $k = 0$ the base case yields *false*
 11 and the inductive step yields *true*. Hence, we move to bound iteration $k = 1$
 12 where the base case yields *false* and the step yields *unknown*. Consequently,
 13 the inductive step gets refined by adding the predicates $(y > 0)$ and $(\overline{y} > \overline{0})$ to
 14 its associated predicate set. After refinement, the inductive step yields *true*.
 15 Thus, we move to bound iteration $k = 2$ where we receive an *unknown* result
 16 for the base case. We refine the base case by adding $(y > 0)$ and $(\overline{y} > \overline{0})$ to
 17 its associated predicate set. After refinement, the base case yields *false*. The
 18 algorithm continues with the inductive step which yields also *false* for the current
 19 bound iteration. Hence, we can conclude that mutual exclusion is generally not
 20 violated for the system under consideration. Note that in each bound iteration
 21 k , the bound of the inductive step is always $k + 1$. In the next section we
 22 introduce our enhanced abstraction refinement technique.

23 7. Witness Refinement and Constraint Generation

24 Abstraction refinement-based model checking is still challenged by the state
 25 explosion problem. Each additional predicate involves an exponential growth
 26 of the state space to be explored. In the following, we introduce an enhanced
 27 abstraction refinement algorithm that allows to reduce the number of predicates
 28 that are actually considered during model checking. Our enhancement is based
 29 on restricting the search space of the model checking problem by *path constraints*
 30 that can be formulated as temporal logic formulae:

Definition 15 (Path Constraints).

Let $M = (S, I, R, L)$ be a three-valued Kripke structure defined over a set of
 atomic predicates A . Moreover, let $\pi_0 \dots \pi_k$ be the k -prefix of a path π in M .
 Then the corresponding *focussing path constraint* is

$$\sigma(\pi_0 \dots \pi_k) := \bigwedge_{i=0}^k \left(\left(\bigwedge_{p \in \mathcal{T}(\pi_i)} p_i \right) \wedge \left(\bigwedge_{p \in \mathcal{F}(\pi_i)} \neg p_i \right) \right)$$

and the corresponding *excluding path constraint* is

$$\bar{\sigma}(\pi_0 \dots \pi_k) := \bigvee_{i=0}^k \left(\left(\bigvee_{p \in \mathcal{T}(\pi_i)} \neg p_i \right) \vee \left(\bigvee_{p \in \mathcal{F}(\pi_i)} p_i \right) \right)$$

31 where $\mathcal{T}(\pi_i) = \{p \in A \mid L(\pi_i, p) = \text{true}\}$ and $\mathcal{F}(\pi_i) = \{p \in A \mid L(\pi_i, p) = \text{false}\}$.

1 Hence, a focussing path constraint is of the form $\phi_0 \wedge \dots \wedge \phi_k$ and an excluding
 2 path constraint is of the form $\phi_0 \vee \dots \vee \phi_k$ where each ϕ_i is a propositional logic
 3 expression over i -indexed predicates ($\phi_i = \mathbf{X}^i \phi$).

Given a prefix $\pi_0 \dots \pi_k$, the corresponding focussing constraint $\sigma(\pi_0 \dots \pi_k)$
 is an LTL formula that is only satisfied for paths π' that have the same definite
 properties as $\pi_0 \dots \pi_k$. Such a constraint can be especially useful in the
 context of three-valued abstraction where refinement always preserves definite
 properties but may also make previously *unknown* properties definite. For the
 unconfirmed witness ω depicted in Figure 8 we can construct the corresponding
 path constraint $\sigma(\omega)$.

$$\begin{aligned} \sigma(\omega) = & (pc_1=0)_0 \wedge \neg(pc_1=1)_0 \wedge (pc_2=0)_0 \wedge \neg(pc_2=1)_0 \\ & \wedge \neg(pc_1=0)_1 \wedge (pc_1=1)_1 \wedge (pc_2=0)_1 \wedge \neg(pc_2=1)_1 \\ & \wedge \neg(pc_1=0)_2 \wedge (pc_1=1)_2 \wedge \neg(pc_2=0)_2 \wedge (pc_2=1)_2 \end{aligned}$$

Now after a refinement step that adds the predicate $(y > 0)$, the prefix

$$\omega' = (0, 0, (y > 0) = t) \longrightarrow (1, 0, (y > 0) = u) \xrightarrow{u} (1, 1, (y > 0) = f)$$

satisfies the constraint $\sigma(\omega)$, since ω' follows the same control flow as ω , whereas
 the prefix

$$\omega'' = (0, 0, \cdot) \longrightarrow (0, 1, \cdot) \longrightarrow (1, 1, \cdot)$$

4 where the \cdot denotes an arbitrary valuation of the predicate $(y > 0)$, does not
 5 satisfy $\sigma(\omega)$ because the prefixes ω and ω'' differ in their control flow. We will
 6 show that focussing path constraints can be used to focus on a particular un-
 7 confirmed witness in order to confirm it or to show that it is spurious after
 8 refinement. Conversely, we will use excluding path constraints to rule out un-
 9 confirmed witnesses that turn out to be spurious. A constraint can be added to
 10 a model checking problem by simply conjugating it with the safety property to
 11 be checked. We will see that, in particular in the context of satisfiability-based
 12 model checking, path constraints can substantially narrow down the search space
 13 of the model checking problem to be solved.

While the algorithm *AR* follows the classical *abstract-check-refine* loop [4],
 we now introduce an enhanced abstraction refinement algorithm that makes use
 of path constraints. The new algorithm is based on a loop

$$\mathit{abstract-check}-(\mathit{refineWitness-checkWitness})^*-\mathit{generateConstraint}$$

14 where the $*$ denotes that the steps in brackets belong to an internal loop with
 15 potentially multiple iterations. The idea of the new approach is as follows. If
 16 abstraction-based model checking on a model that covers the global state space
 17 returns an unconfirmed witness ω , then we start an internal refinement loop with
 18 a local model that is restricted to refinements of the witness ω only. The local
 19 model can be straightforwardly obtained by using the focussing path constraint
 20 $\sigma(\omega)$, which masks out all prefixes that differ from ω . The *witness refinement*

1 *loop* either results in a definite witness, which means we are done, or it tells
 2 us that ω is spurious. In the latter case, we generate the constraint $\bar{\sigma}(\omega)$ that
 3 excludes the unconfirmed witness ω from further consideration. In the next
 4 overall loop, we return to the global model and we use the constraint $\bar{\sigma}(\omega)$ in
 5 order to restrict the state space exploration. But we do not need to add the
 6 refinement predicates that we used in the local model in order to generate the
 7 constraint. Hence, we have two forms of refinement respectively concretisation
 8 here: predicate refinement along unconfirmed witnesses in a local model and
 9 the pruning of infeasible paths via constraints in the global model. The latter
 10 does not involve any increase of the state space. Algorithm 3 shows our new
 11 abstraction refinement algorithm *WRC*.

Algorithm 3: $WRC(A[M, I \models_{\exists} \psi]_k)$

```

1   $\Sigma_k := true$  /*cumulative excluding path constraint*/
2  loop forever do /*global constraint loop*/
3  |   if  $A[M, I \models_{\exists} \Sigma_k \wedge \psi]_k = false$  then
4  |   |   return false, ‘no witness for safety violation of length  $k$  exists’
5  |   if  $A[M, I \models_{\exists} \Sigma_k \wedge \psi]_k = true$  and definite witness  $\omega$  then
6  |   |   return true, ‘ $\omega$  is a definite witness for safety violation’
7  |   if  $A[M, I \models_{\exists} \Sigma_k \wedge \psi]_k = unknown$  and unconfirmed witness  $\omega$  then
8  |   |    $A^\omega := A \cup analyseWitness(\omega)$ 
9  |   |   loop forever do /*refinement loop local to  $\omega^*$ */
10 |   |   |   if  $A^\omega[M, I \models_{\exists} \sigma(\omega) \wedge \psi]_k = false$  then
11 |   |   |   |   /*  $\omega$  is spurious */
12 |   |   |   |    $\Sigma_k := \Sigma_k \wedge \bar{\sigma}(\omega)$ 
13 |   |   |   |   goto 3
14 |   |   |   if  $A^\omega[M, I \models_{\exists} \sigma(\omega) \wedge \psi]_k = true$  and witness  $v$  then
15 |   |   |   |   return true, ‘ $v$  is a definite witness for safety violation’
16 |   |   |   if  $A^\omega[M, I \models_{\exists} \sigma(\omega) \wedge \psi]_k = unknown$  and witness  $v$  then
17 |   |   |   |    $A^\omega := A^\omega \cup analyseWitness(v)$ 

```

12 *WRC* consists of an outer constraint loop where we operate on a global state
 13 space model defined over a global predicate set A and cumulative excluding path
 14 constraint Σ_k that is initially *true*, i.e. no paths are excluded. While A remains
 15 constant throughout the execution of the algorithm, Σ_k will be gradually ex-
 16 tended with constraints that rule out *spurious witnesses*. The cases where a
 17 definite result is obtained in the outer loop are identically handled as in *AR*. If
 18 an *unknown* result together with an unconfirmed witness ω is obtained in the
 19 outer loop, then the algorithm enters an inner refinement loop *local* to ω . In the
 20 inner loop, we use a model defined over the predicate set A^ω . A^ω is initialised
 21 as the union of A and the refinement predicates derived from the unconfirmed
 22 witness ω . Moreover, the temporal logic formula to be checked is conjugated

1 with the focussing path constraint $\sigma(\omega)$, which restricts the feasible paths to
 2 those whose prefix is a refinement of ω . Hence, the model checking problem
 3 in the inner loop has a refined state space defined over A^ω , but the employed
 4 model is local in the sense that the state space exploration is narrowed down to
 5 refinements of ω . In case of a *false* result in the inner loop, we have that ω is
 6 a spurious witness. We then extend the cumulative excluding path constraint
 7 Σ_k by the constraint $\bar{\sigma}(\omega)$, which excludes ω from further consideration, and we
 8 return to the outer loop where we operate again with the global model and the
 9 original predicate set A . In case of a *true* result in the inner loop, we obtain
 10 a definite witness v that is a refinement of ω . Thus, *WRC* can terminate. In
 11 case of an *unknown* result in the inner loop, we obtain an unconfirmed witness
 12 v that is a refinement of ω . We then derive new predicates from v and continue
 13 with a further refinement iteration local to ω . We get the following theorem
 14 with regard to the return values of *AR* and *WRC*:

15 **Theorem 1**

16 Let ${}_A[M, I \models_{\exists} \psi]_k$ be a three-valued bounded model checking problem where M
 17 is a state space model of a system *Sys* abstracted over A and ψ is an LTL safety
 18 formula defined over A . Then the following holds:

- 19 1. $AR({}_A[M, I \models_{\exists} \psi]_k) = true$ iff $WRC({}_A[M, I \models_{\exists} \psi]_k) = true$
 20 2. $AR({}_A[M, I \models_{\exists} \psi]_k) = false$ iff $WRC({}_A[M, I \models_{\exists} \psi]_k) = false$

21 *Proof.* See <http://github.com/ssfm-up/TVMC/raw/master/SCIC0Proofs.pdf>

22 Hence, both algorithms return the same result for the same input model check-
 23 ing problem with regard to a system *Sys*. We have already shown in [3] that
 24 the result of *AR* correctly characterises the computational behaviour of *Sys*.
 25 Thus, we can conclude that the result of *WRC* also correctly characterises the
 26 behaviour of *Sys*.

27 We now illustrate how *WRC* processes verification tasks based on our run-
 28 ning example with regard to the mutual exclusion system and a bound of $k = 2$.
 29 The initial predicate set is again: $A = \{(pc_1 = 0), (pc_1 = 1), (pc_2 = 0), (pc_2 = 1)\}$.
 30 The set is de facto reducible to just $\{(pc_1 = 0), (pc_2 = 0)\}$ by assuming the equiv-
 31 alences $(pc_1 = 1) \equiv \neg(pc_1 = 0)$ and $(pc_2 = 1) \equiv \neg(pc_2 = 0)$, which we effectively
 32 do in the implementation of our approach. However, for illustrative purposes
 33 we use the expanded set A here. In the first global constraint iteration, *WRC*
 34 detects the unconfirmed witness $\omega = (0, 0) \xrightarrow{u} (1, 0) \xrightarrow{u} (1, 1)$. Similar as in
 35 our illustration of the basic algorithm, *WRC* now derives the refinement pred-
 36 icates $(y > 0)$ and $(\overline{y > 0})$. But the predicates are added to the *local* predicate
 37 set A^ω . Moreover, the focussing path constraint $\sigma(\omega)$ is added to the model
 38 checking problem, which gives us a local state space model. Hence, when we are
 39 solving the refined problem in the inner loop, the valuation of the control flow
 40 predicates in each state along a prefix is now fixed by $\sigma(\omega)$. This means that the
 41 complexity of the state space to be explored is solely induced by the predicates
 42 over y . Model checking yields that ω is spurious. Consequently, $\bar{\sigma}(\omega)$ is added
 43 to the cumulative excluding path constraint Σ_k via conjunction. This excludes

1 any further consideration of ω and its possible refinements. The next iteration
 2 detects another unconfirmed witness $\omega' = (0, 0) \xrightarrow{u} (0, 1) \xrightarrow{u} (1, 1)$. *WRC*
 3 now enters a refinement loop local to ω' . It detects that ω' is also spurious,
 4 and thus, can be ruled out via the excluding path constraint $\bar{\sigma}(\omega')$. In the final
 5 constraint iteration *WRC* terminates with the definite result that no witness
 6 for safety violation of length $k = 2$ exists.

7 With *WRC* we are able to reduce the number of predicates that actually
 8 contribute to the size of the state space to be explored. In our simple example,
 9 *WRC* had to solve model checking problems on global and local models with
 10 a maximum number of two predicates, whereas *AR* had to solve a problems
 11 with maximum four predicates. The price that we pay is an increased number
 12 of model checking runs. In our experiments we will show that the savings due
 13 to the reduced number of predicates typically outweigh the extra costs due to
 14 additional model checking runs. Similar to *AR*, the enhanced algorithm can be
 15 straightforwardly combined with a bound iteration loop ranging over k . The
 16 corresponding algorithm is depicted below.

Algorithm 4: $k\text{-IND}_{(A_b[base]_k, A_s[step]_k)}$

```

1  for  $k = 0$  to  $\infty$  do
2  |   if  $WRC_{(A_b[base]_k)} = true$  then
3  |   |   return "safety property violated"
4  |   else if  $WRC_{(A_s[step]_{k+1})} = false$  then
5  |   |   return "safety property holds"

```

17 Here the input $A_b[base]_k$ is a three-valued bounded model checking problem
 18 corresponding to the base case of the k -induction approach where A_b is the ini-
 19 tial set of predicates and k is the bound parameter which gets initialised within
 20 the algorithm. Likewise, $A_s[step]_k$ is a three-valued bounded model checking
 21 problem corresponding to the inductive step with initial predicate set A_s .

22 *7.1. Constraint Strengthening and Reuse Between Bound Iterations*

23 In general, it is not admissible to reuse generated constraints $\bar{\sigma}(\omega)$ for ruling
 24 out spurious witnesses between bound iterations and between the base case and
 25 the inductive step. For instance, a witness $\omega = \pi_0 \dots \pi_k$ might be spurious
 26 in iteration k , but it might be still the prefix of a definite witness in some
 27 later iteration $k + j$. However, we have identified different *types* of spurious
 28 witness constraints. Depending on the type, a spurious witness constraint may
 29 be straightforwardly reused or adapted for reuse in a higher bound iteration.

30 The algorithm *WRC* identifies an unconfirmed witness ω to be spurious if
 31 $A^\omega[M, I \models \exists \sigma(\omega) \wedge \psi]_k$ yields *false* (line 10). We can re-formulate this model
 32 checking problem as $A^\omega[M, S \models \exists I_0 \wedge \sigma(\omega) \wedge \psi]_k$ with $I_0 = \bigvee_{s \in I} \sigma(s)$. This
 33 gives us an equivalent model checking problem where the initial state constraint
 34 is part of the temporal logic formula to be checked. A sufficient condition for
 35 $A^\omega[M, S \models \exists I_0 \wedge \sigma(\omega) \wedge \psi]_k = false$ is that any conjunct or combination of

1 conjuncts of the overall formula $I_0 \wedge \sigma(\omega) \wedge \psi$ is not satisfied by M . We denote
 2 such a conjunct (combination) as a *cause of violation* of the model checking
 3 problem:

4 **Definition 16 (Cause of Violation).**

5 Let ${}_A[M, S \models_{\exists} \phi \wedge \phi']_k$ be a bounded model checking problem that yields *false*.
 6 Then ϕ is a *cause of violation* of the problem if ${}_A[M, S \models_{\exists} \phi]_k = \text{false}$.

Evidently, for every model checking problem that yields *false*, the overall LTL formula is always a cause of violation itself. However, smaller causes may exist. For our abstraction refinement approach and the model checking problem ${}_{A^\omega}[M, S \models_{\exists} I_0 \wedge \sigma(\omega) \wedge \psi]_k$ local to an unconfirmed witness ω a cause of violation is also a cause of spuriousness of the witness. While the cumulative constraint Σ_k of the *current* bound iteration is simply extended by the spurious witness constraint $\bar{\sigma}(\omega)$ (line 12), $\bar{\sigma}(\omega)$ may have to be adapted for reuse in a *higher* bound iteration $k + j$. In our context, adaptation means to increment the index values that occur in $\bar{\sigma}(\omega)$ according to the higher bound. We define a j -increment of the indices occurring in a constraint as

$$\bar{\sigma}(\omega)_j := \bar{\sigma}(\omega)[i \leftarrow i + j \mid 0 \leq i \leq k].$$

7 Thus, a j -increment substitutes i -indexed predicates p_i in $\bar{\sigma}(\omega)$ by $(i + j)$ -
 8 indexed predicates p_{i+j} . While the incrementation of indices may be *necessary*
 9 for reusing a constraint in a higher iteration, it may be additionally *feasible*
 10 (but not necessary) to strengthen a generated constraint. Note that since the
 11 focussing path constraint $\sigma(\omega)$ is a pure conjunction, a cause of violation may
 12 only contain a sub-formula $\sigma(\omega)^{sub}$ of $\sigma(\omega)$. We will see that in this case the
 13 stronger formula $\bar{\sigma}(\omega)^{sub}$ is a feasible constraint for excluding spuriousness. We
 14 denote such a $\bar{\sigma}(\omega)^{sub}$ as a *spurious segment constraint*.

15 We have proven Theorem 2 that assigns to the possible causes of violation of
 16 ${}_{A^\omega}[M, S \models_{\exists} I_0 \wedge \sigma(\omega) \wedge \psi]_k$ admissible extensions of the constraints Σ_{k+j} with
 17 $j \in \mathbb{N}$.

18 **Theorem 2**

19 *Let ϕ_x be a cause of violation of the three-valued bounded model checking problem*
 20 ${}_{A^\omega}[M, I \models_{\exists} \sigma(\omega) \wedge \psi]_k$ *local to an unconfirmed witness ω . Then in bound*
 21 *iteration $k + j$ with $j \in \mathbb{J}$ it is admissible to extend the cumulative path constraint*
 22 Σ_{k+j} *of the corresponding global model checking problem ${}_A[M, I \models_{\exists} \Sigma_{k+j} \wedge \psi]_{k+j}$*
 23 *as follows: $\Sigma_{k+j} := \Sigma_{k+j} \wedge \phi_x$ where*

$\phi_1 = I_0 \wedge \sigma(\omega)^{sub} \wedge \psi$	$\varphi_1 = \bar{\sigma}(\omega)^{sub}$	$\mathbb{J} = \{0\}$
$\phi_2 = I_0 \wedge \sigma(\omega)^{sub}$	$\varphi_2 = \bar{\sigma}(\omega)^{sub}$	$\mathbb{J} = \mathbb{N}$
$\phi_3 = \sigma(\omega)^{sub} \wedge \psi$	$\varphi_3 = \bar{\sigma}(\omega)_j^{sub}$	$\mathbb{J} = \mathbb{N}$
$\phi_4 = I_0 \wedge \psi$	$\varphi_4 = \text{false}$	$\mathbb{J} = \{0\}$
$\phi_5 = \sigma(\omega)^{sub}$	$\varphi_5 = \bigwedge_{l=0}^j \bar{\sigma}(\omega)_l^{sub}$	$\mathbb{J} = \mathbb{N}$
$\phi_6 = \psi$	$\varphi_6 = \text{false}$	$\mathbb{J} = \{0\}$
$\phi_7 = I_0$	$\varphi_7 = \text{false}$	$\mathbb{J} = \mathbb{N}$

1 *Proof.* See <http://github.com/ssfm-up/TVMC/raw/master/SCICOProofs.pdf>

2
 3 Here the set \mathbb{J} denotes the range of bound iterations for which the constraint φ_x
 4 is admissible. If $\mathbb{J} = \{0\}$ then φ_x is an admissible constraint in iteration $k + 0$
 5 only. If $\mathbb{J} = \mathbb{N}$ then φ_x is admissible in all iterations $k + 0, k + 1, k + 2, \dots$

6 In our context, *admissible* means that the added spurious segment constraint
 7 $\bar{\sigma}(\omega)^{sub}$ will only rule out paths that would anyway turn out to be spurious after
 8 the refinement of the global model checking problem $A[M, I \models \exists \Sigma_{k+j} \wedge \psi]_{k+j}$
 9 over some extended predicate set $A' \supset A$. However, with our approach we are
 10 able to exclude all paths that exhibit the spurious segment *without* refining the
 11 global problem. For each distinct cause of violation we defined a descriptive
 12 *type* of the associated constraint:

cause of violation	constraint type
ϕ_1	<i>full-dependent</i>
ϕ_2	<i>property-independent</i>
ϕ_3	<i>initial state-independent</i>
ϕ_4	<i>redundant I</i>
ϕ_5	<i>fully-independent</i>
ϕ_6	<i>redundant II</i>
ϕ_7	<i>redundant III</i>

13
 14 We now illustrate certain interesting cases of causes of violation resp. con-
 15 straint types and the consequent reuse of constraints. Remember that ψ is of
 16 the form $\mathbf{O}^k \neg safe$. Hence, in bound iteration k we check for safety violation at
 17 position k whereas in a higher bound iteration $k + j$ we check for safety violation
 18 at position $k + j$.

- 19 • If the cause of violation is of the form ϕ_1 , this tells us that there exists no
 20 k -prefix $s_0 \dots s_k$ satisfying $\sigma(\omega)^{sub}$ that starts in an initial state $s_0 \in I$ and
 21 ends in a state s_k in which *safe* is violated. Hence, the cause of violation is
 22 linked to an initial state at position 0 *and* to an error state at position k .
 23 We denote the corresponding constraint as *fully-dependent* on the current
 24 bound iteration. The excluding path constraint $\bar{\sigma}(\omega)^{sub}$ is admissible in
 25 the current iteration k . However, in any bound iteration $k + j$ with $j > 0$
 26 we search for a path prefix from an initial state to an error state at position
 27 $k + j$ rather than at position k . Such a prefix may still satisfy $\sigma(\omega)^{sub}$,
 28 and thus, $\bar{\sigma}(\omega)^{sub}$ is not an admissible constraint in iteration $k + j$.
- 29 • If the cause of violation is of the form ϕ_2 , this tells us that there exists no
 30 k -prefix $s_0 \dots s_k$ satisfying $\sigma(\omega)^{sub}$ that starts in an initial state $s_0 \in I$.
 31 Hence, the cause of violation is linked to an initial state at position 0
 32 but *not* to any error state. We denote the corresponding constraint as
 33 (error) *property-independent*. If there exists no k -prefix satisfying $\sigma(\omega)^{sub}$
 34 in iteration k then there is also not such a prefix in any iteration $k + j$.
 35 Thus, it is admissible to reuse the excluding path constraint $\bar{\sigma}(\omega)^{sub}$ in
 36 higher bound iterations.

- 1 • If the cause of violation is of the form ϕ_3 , this tells us that there exists
2 no k -prefix $s_0 \dots s_k$ satisfying $\sigma(\omega)^{sub}$ ends in a state s_k in which *safe* is
3 violated. Since the cause of violation is *not* linked to an initial state such a
4 k -prefix may start in an *arbitrary* state. The cause is however linked to an
5 error state at position k . We denote the corresponding constraint as *initial*
6 *state-independent*. In bound iteration $k+j$ with $j > 0$ we search for a path
7 prefix $s_0 \dots s_{k+j}$ from an initial state to an error state at position $k+j$.
8 Note that the j -suffix $s_j \dots s_{k+j}$ of such a prefix is a k -prefix leading to
9 an error state. Hence, $\sigma(\omega)^{sub}$ must be violated for this j -suffix. In order
10 to exclude corresponding paths it is admissible to use the j -incremented
11 constraint $\bar{\sigma}(\omega)_j^{sub}$ in iteration $k+j$.
- 12 • If the cause of violation is of the form ϕ_5 , this tells us that there exists
13 no k -prefix $s_0 \dots s_k$ satisfying $\sigma(\omega)^{sub}$. Hence, the cause of violation is
14 neither linked to an initial state nor to an error state. We denote the
15 corresponding constraint as *fully-independent*. We have that $\sigma(\omega)^{sub}$ is
16 violated for arbitrary sequences of states $s_0 \dots s_k$ in the model. Thus,
17 in any bound iteration $k+j$ all l -incremented constraints $\bar{\sigma}(\omega)_l^{sub}$ with
18 $0 \leq l \leq j$ are admissible.

19 As we can see in Theorem 2, there are also causes of violation where we can
20 immediately conclude that the model checking problem yields *false*, regardless
21 of the constraint. In these cases we denote the corresponding constraint as *re-*
22 *dundant*. In case of a *redundant I* or *redundant II* constraint the *false* result
23 only holds for the current bound iteration. In the latter case the *false* result is
24 additionally transferable from the base case to the corresponding inductive step.
25 In case of a *redundant III* constraint the *false* result can be even transferred to
26 higher bound iterations.

27 Based on a similar argumentation we are also able to characterise admissible
28 constraint reuse between the base case and the inductive step: If the cause of
29 violation of the base case does not contain the initial state condition I_0 then
30 the corresponding constraint is initial state-independent and can be reused for
31 the inductive step (because there is no initial state condition in the step). And,
32 if the cause of violation of the inductive step does not contain the loop-free
33 condition $loopFree_{0,k+1}$ then the corresponding constraint can be reused for the
34 base case (because there is no loop-free condition in the base case). Based on
35 the results of Theorem 2 constraint reusing can be easily integrated into the
36 algorithm *k-IND* that invokes the algorithm *WRC* for ascending bounds.

Remember our running example where we detected in bound iteration $k = 2$
of the base case that the unconfirmed witness $\omega = (0, 0) \xrightarrow{u} (1, 0) \xrightarrow{u} (1, 1)$ is
spurious. A corresponding cause of violation is

$$\sigma(\omega)^{sub} = (pc_1 = 0)_0 \wedge (pc_2 = 0)_0 \wedge (pc_1 = 1)_2 \wedge (pc_2 = 1)_2$$

37 The cause is of the form ϕ_5 as introduced in Theorem 2. Hence, the correspond-
38 ing spurious witness constraint $\bar{\sigma}(\omega)^{sub}$ is full-independent, i.e. neither linked
39 to an initial state nor to an error state, and thus, reusable in all future bound

1 iterations $k + j$ of the base case and the inductive step. Moreover, all index
 2 increments $\bar{\sigma}(\omega)_l^{sub}$ with $0 \leq l \leq j$ are also admissible constraints in iterations
 3 $k + j$. Note that the cause of violation $\sigma(\omega)^{sub}$ is a real sub-formula of $\sigma(\omega)$,
 4 which positively affects the strength of the spurious segment constraint $\bar{\sigma}(\omega)^{sub}$
 5 and its index increments. The constraint reveals that there does not exist any
 6 sequence of two transitions leading from $(0, 0)$ to $(1, 1)$ – independent of which
 7 state will be reached between $(0, 0)$ and $(1, 1)$. After reusing this constraint for
 8 the inductive step, the algorithm k -IND will immediately terminate with the
 9 result that the safety properties holds.

10 We have outlined the benefits of cause-based constraint reusing between
 11 bound iterations in terms of ruling out spurious behaviour. However, so far we
 12 have not discussed how causes of violation of a model checking problem can
 13 be automatically and efficiently computed. In the next section we introduce
 14 the propositional logic encoding of three-valued bounded model checking. The
 15 encoding enables us to solved model checking problems via Boolean satisfiability
 16 (SAT) solving. Moreover, we will see that the SAT-based approach reduces the
 17 detection of causes of violation to the extraction of *unsatisfiable cores* of a
 18 propositional logic formula, for which efficient tools exist.

19 8. Reduction to Propositional Logic Satisfiability

In our previous work [7] we showed how a three-valued bounded model check-
 ing problem ${}_A[M, I \models_{\exists} \psi]_k$ can be encoded as a propositional logic formula
 ${}_A[[M, I, \psi, k]]$. The encoding corresponds to an implicit problem representation
 such that the construction of an explicit Kripke structure is avoided. The for-
 mula ${}_A[[M, I, \psi, k]]$ is defined over a set of Boolean atoms $Atoms$, the constants
true, *false*, and a special atom \perp that is used to represent the *unknowns* due to
 abstraction. The atom \perp occurs solely non-negated in ${}_A[[M, I, \psi, k]]$. Based on
 the encoding, three-valued bounded model checking can be performed via two
 satisfiability checks. The first check considers an *under-approximating comple-*
tion, marked with ‘-’, where all \perp ’s are assumed to be *false*:

$${}_A[[M, I, \psi, k]]^- := {}_A[[M, I, \psi, k]][\perp \mapsto \textit{false}]$$

and the second check considers an *over-approximating completion*, marked with
 ‘+’, where all \perp ’s are assumed to be *true*:

$${}_A[[M, I, \psi, k]]^+ := {}_A[[M, I, \psi, k]][\perp \mapsto \textit{true}].$$

20 Here $[\perp \mapsto z]$, $z \in \{\textit{true}, \textit{false}\}$ denotes the assumption that the special atom
 21 \perp is assigned to z . This gives us the notion of *3-valued satisfiability sat₃*:

Definition 17 (sat₃).

Let ${}_A[[M, I, \psi, k]]$ over $Atoms$ be the propositional logic encoding of a three-
 valued bounded model checking problem ${}_A[M, I \models_{\exists} \psi]_k$. Then **sat₃** is defined

as:

$$\mathbf{sat}_3({}_A\llbracket M, I, \psi, k \rrbracket) = \begin{cases} true & \text{if } \mathbf{sat}({}_A\llbracket M, I, \psi, k \rrbracket^-) = true \\ false & \text{if } \mathbf{sat}({}_A\llbracket M, I, \psi, k \rrbracket^+) = false \\ unknown & \text{else} \end{cases}$$

- 1 Thus, a \mathbf{sat}_3 problem is reducible to two Boolean satisfiability problems. In [7]
 2 the following lemma has been proven:

Lemma 2 *Let ${}_A\llbracket M, I, \psi, k \rrbracket$ over Atoms be the propositional logic encoding of a three-valued bounded model checking problem ${}_A[M, I \models_{\exists} \psi]_k$. Then:*

$${}_A[M, I \models_{\exists} \psi]_k = \mathbf{sat}_3({}_A\llbracket M, I, \psi, k \rrbracket)$$

Hence, by solving \mathbf{sat}_3 we obtain the result of the encoded three-valued bounded model checking problem. If the results of the two Boolean satisfiability checks are $\mathbf{sat}({}_A\llbracket M, \psi, k \rrbracket^-) = false$ and $\mathbf{sat}({}_A\llbracket M, \psi, k \rrbracket^+) = true$, then we can conclude that the result of the encoded problem is *unknown*. In this case, a truth assignment $\mathcal{A} : \text{Atoms} \rightarrow \{true, false\}$ that satisfies ${}_A\llbracket M, \psi, k \rrbracket^+$ characterises an unconfirmed witness ω . Thus, witness generation in the SAT-based approach is straightforward. The details on how the formula ${}_A\llbracket M, I, \psi, k \rrbracket$ is built and on how witnesses ω can be derived from satisfying assignments \mathcal{A} can be found in [7]. ${}_A\llbracket M, I, \psi, k \rrbracket$ is in *conjunctive normal form* (CNF) and we assume a representation of the CNF formula as a set of sets of literals $\{\{l, \dots, l'\}, \dots, \{l'', \dots, l'''\}\}$. The construction of ${}_A\llbracket M, I, \psi, k \rrbracket$ is divided into the encoding of initial states I , the encoding of k unrollings of the transition relation of M and the encoding of the temporal logic formula ψ for bound k :

$${}_A\llbracket M, I, \psi, k \rrbracket = \llbracket M, k \rrbracket \cup \llbracket I \rrbracket \cup \llbracket \psi, k \rrbracket$$

In the expanded representation, we omit the reference to the associated predicate set A , as this is clear from the context. Since path constraints Σ_k and $\sigma(\omega)$ are also temporal logic formulae, the propositional logic encoding of model checking problems with constraints, as used in the refinement algorithm *WRC*, is straightforward. For a global model checking problem ${}_A[M, I \models_{\exists} \Sigma_k \wedge \psi]_k$ with a cumulative path constraint Σ_k we get

$${}_A\llbracket M, I, \Sigma_k, \psi, k \rrbracket = \llbracket M, k \rrbracket \cup \llbracket I \rrbracket \cup \llbracket \Sigma_k \rrbracket \cup \llbracket \psi, k \rrbracket$$

and for a model checking problem ${}_{A\omega}[M, I \models_{\exists} \sigma(\omega) \wedge \psi]_k$ local to an unconfirmed witness ω we get

$${}_{A\omega}\llbracket M, I, \sigma(\omega), \psi, k \rrbracket = \llbracket M, k \rrbracket \cup \llbracket I \rrbracket \cup \llbracket \sigma(\omega) \rrbracket \cup \llbracket \psi, k \rrbracket.$$

- 3 This allows us to redefine the algorithm *WRC* as a satisfiability-based version
 4 *SATWRC*.
 5 As we can see, each three-valued bounded model checking problem to be

Algorithm 5: $SATWRC(A[M, I \models_{\exists} \psi]_k)$

```
1  $\Sigma_k := true$  /*cumulative excluding path constraint*/
2 loop forever do /*global constraint loop*/
3   if  $\text{sat}_3(A[M, I, \Sigma_k, \psi, k]) = false$  then
4     return false, ‘no witness for safety violation of length  $k$  exists’
5   if  $\text{sat}_3(A[M, I, \Sigma_k, \psi, k]) = true$  and definite witness  $\omega$  then
6     return true, ‘ $\omega$  is a definite witness for safety violation’
7   if  $\text{sat}_3(A[M, I, \Sigma_k, \psi, k]) = unknown$  and unconfirmed witness  $\omega$ 
   then
8      $A^\omega := A \cup \text{analyseWitness}(\omega)$ 
9     loop forever do /*refinement loop local to  $\omega^*$ */
10    if  $\text{sat}_3(A^\omega[M, I, \sigma(\omega), \psi, k]) = false$  then
11      /*  $\omega$  is spurious */
12       $\Sigma_k := \Sigma_k \wedge \bar{\sigma}(\omega)$ 
13      goto 3
14    if  $\text{sat}_3(A^\omega[M, I, \sigma(\omega), \psi, k]) = true$  and witness  $v$  then
15      return true, ‘ $v$  is a definite witness for safety violation’
16    if  $\text{sat}_3(A^\omega[M, I, \sigma(\omega), \psi, k]) = unknown$  and witness  $v$  then
17       $A^\omega := A^\omega \cup \text{analyseWitness}(v)$ 
```

1 solved is substituted by a **sat**₃ problem. Thus, the soundness of *SATWRC*
2 in terms of returning the correct model checking result follows from Theorem
3 1 and Lemma 2. The algorithm extends the cumulative path constraint Σ_k
4 in the same manner as *WRC*. The extended constraint then gets encoded as
5 part of $A[M, I, \Sigma_k, \psi, k]$. So far, none of the improvements that we introduced
6 in the previous section are integrated into the algorithm. Remember that our
7 improvements with regard to constraint strengthening and constraint reusing
8 were based on causes of violation of a model checking problem. We will now
9 show that there is a strong correspondence between a *cause of violation* of an
10 explicit model checking problem and an *unsatisfiable core* of the propositional
11 logic encoding of a model checking problem.

12 **Definition 18 (Unsatisfiable Core).**

13 Let α be a propositional logic formula in conjunctive normal form with $\text{sat}_3(\alpha) =$
14 *false*. An *unsatisfiable core* is a subset $\alpha_{uc} \subseteq \alpha$ of clauses of α with $\text{sat}_3(\alpha_{uc}) =$
15 *false*.

State-of-the-art SAT solvers support the efficient extraction of small or
even minimal unsatisfiable cores [19]. Thus, when *SATWRC* detects that
 $\text{sat}_3(A^\omega[M, I, \sigma(\omega), \psi, k])$ yields *false*, i.e. that ω is a spurious witness, then

we can extract an unsatisfiable core of the encoded model checking problem

$${}_{A^\omega} \llbracket M, I, \sigma(\omega), \psi, k \rrbracket = \llbracket M, k \rrbracket \cup \llbracket I \rrbracket \cup \llbracket \sigma(\omega) \rrbracket \cup \llbracket \psi, k \rrbracket$$

which is a sub-formula

$${}_{A^\omega} \llbracket M, I, \sigma(\omega), \psi, k \rrbracket_{uc} = \llbracket M, k \rrbracket_{uc} \cup \llbracket I \rrbracket_{uc} \cup \llbracket \sigma(\omega) \rrbracket_{uc} \cup \llbracket \psi, k \rrbracket_{uc}$$

with

$$\llbracket M, k \rrbracket_{uc} \subseteq \llbracket M, k \rrbracket, \llbracket I \rrbracket_{uc} \subseteq \llbracket I \rrbracket, \llbracket \sigma(\omega) \rrbracket_{uc} \subseteq \llbracket \sigma(\omega) \rrbracket \text{ and } \llbracket \psi, k \rrbracket_{uc} \subseteq \llbracket \psi, k \rrbracket.$$

1 Hence, such an unsatisfiable core consists of a model-related part $\llbracket M, k \rrbracket_{uc}$,
 2 an initial state-related part $\llbracket I \rrbracket_{uc}$, a constraint-related part $\llbracket \sigma(\omega) \rrbracket_{uc}$ and a
 3 property-related part $\llbracket \psi, k \rrbracket_{uc}$. For our approach, the constraint related part
 4 is of major interest. We have proven the following lemma:

5 **Lemma 3**

6 Let $\text{sat}_3({}_{A^\omega} \llbracket M, I, \sigma(\omega), \psi, k \rrbracket) = \text{false}$ and let $\llbracket \sigma(\omega) \rrbracket_{uc}$ be the constraint-related
 7 part of the unsatisfiable core of ${}_{A^\omega} \llbracket M, I, \sigma(\omega), \psi, k \rrbracket$. Then there exists a unique
 8 sub-formula $\sigma(\omega)^{uc}$ of the constraint $\sigma(\omega)$ such that $\llbracket \sigma(\omega) \rrbracket_{uc} \subseteq \llbracket \sigma(\omega)^{uc} \rrbracket$ and
 9 $\text{sat}_3({}_{A^\omega} \llbracket M, I, \sigma(\omega)^{uc}, \psi, k \rrbracket) = \text{false}$.

10 *Proof.* See <http://github.com/ssfm-up/TVMC/raw/master/SCICOProofs.pdf>
 11
 12

13 Hence, if SAT-based model checking yields that the unconfirmed witness ω is
 14 spurious, then the constraint-related part $\llbracket \sigma(\omega) \rrbracket_{uc}$ of the unsatisfiable core en-
 15 codes a sub-formula $\sigma(\omega)^{uc}$ of $\sigma(\omega)$ which already contributes to the violation of
 16 the encoded model checking problem. While $\sigma(\omega)$ characterises the *entire* spu-
 17 rious witness ω , $\sigma(\omega)^{uc}$ characterises a *segment* of ω that is already spurious by
 18 itself. Both complements $\bar{\sigma}(\omega)$ and $\bar{\sigma}(\omega)^{uc}$ are admissible constraints for ruling
 19 out spurious behaviour. However, while the spurious witness constraint $\bar{\sigma}(\omega)$
 20 excludes ω only, the spurious segment constraint $\bar{\sigma}(\omega)^{uc}$ excludes *all* paths that
 21 exhibit the spurious segment. We can immediately obtain $\sigma(\omega)^{uc}$ from $\llbracket \sigma(\omega) \rrbracket_{uc}$
 22 by applying the inverse $\llbracket \cdot \rrbracket^{-1}$ of the encoding $\llbracket \cdot \rrbracket$: $\sigma(\omega)^{uc} := \llbracket \llbracket \sigma(\omega) \rrbracket_{uc} \rrbracket^{-1}$.

23 Beside extracting the stronger spurious fragment constraints, unsatisfiable
 24 cores allow for further improvements of *SATWRC*. Note that it is possible that
 25 certain parts of an unsatisfiable core $\llbracket M, k \rrbracket_{uc} \cup \llbracket I \rrbracket_{uc} \cup \llbracket \sigma(\omega) \rrbracket_{uc} \cup \llbracket \psi, k \rrbracket_{uc}$ are
 26 the empty set, e.g. $\llbracket I \rrbracket_{uc} = \emptyset$. In the following, we will omit the empty parts of
 27 an unsatisfiable core of an encoded model checking problem and we assume that
 28 all shown parts are *non-empty*. In Theorem 3 we consider the different cases
 29 of unsatisfiable cores of ${}_{A^\omega} \llbracket M, I, \sigma(\omega), \psi, k \rrbracket$ with empty parts and we assign an
 30 admissible extension of the cumulative constraint Σ_{k+j} with $j \in \mathbb{N}$ to each case.

31 **Theorem 3**

32 *Let ϕ_x be an unsatisfiable core of the encoding ${}_{A^\omega} \llbracket M, I, \sigma(\omega), \psi, k \rrbracket$ of a three-*
 33 *valued bounded model checking problem local to an unconfirmed witness ω . Then*

1 in bound iteration $k + j$ with $j \in \mathbb{J}$ it is admissible to extend the cumulative path
 2 constraint Σ_{k+j} of the encoding $A \llbracket M, I, \Sigma_{k+j}, \psi, k + j \rrbracket$ of the corresponding
 3 global model checking problem as follows: $\Sigma_{k+j} := \Sigma_{k+j} \wedge \varphi_x$ with

4	$\phi_1 = \llbracket M, k \rrbracket_{uc} \cup \llbracket I \rrbracket_{uc} \cup \llbracket \sigma(\omega) \rrbracket_{uc} \cup \llbracket \psi, k \rrbracket_{uc}$	$\varphi_1 = \bar{\sigma}(\omega)^{uc}$	$\mathbb{J} = \{0\}$
	$\phi_2 = \llbracket M, k \rrbracket_{uc} \cup \llbracket I \rrbracket_{uc} \cup \llbracket \sigma(\omega) \rrbracket_{uc}$	$\varphi_2 = \bar{\sigma}(\omega)^{uc}$	$\mathbb{J} = \mathbb{N}$
	$\phi_3 = \llbracket M, k \rrbracket_{uc} \cup \llbracket \sigma(\omega) \rrbracket_{uc} \cup \llbracket \psi, k \rrbracket_{uc}$	$\varphi_3 = \bar{\sigma}(\omega)_j^{uc}$	$\mathbb{J} = \mathbb{N}$
5	$\phi_4 = \llbracket M, k \rrbracket_{uc} \cup \llbracket I \rrbracket_{uc} \cup \llbracket \psi, k \rrbracket_{uc}$	$\varphi_4 = false$	$\mathbb{J} = \{0\}$
	$\phi_5 = \llbracket M, k \rrbracket_{uc} \cup \llbracket \sigma(\omega) \rrbracket_{uc}$	$\varphi_5 = \bigwedge_{l=0}^j \bar{\sigma}(\omega)_l^{uc}$	$\mathbb{J} = \mathbb{N}$
	$\phi_6 = \llbracket M, k \rrbracket_{uc} \cup \llbracket \psi, k \rrbracket_{uc}$	$\varphi_6 = false$	$\mathbb{J} = \{0\}$
	$\phi_7 = \llbracket M, k \rrbracket_{uc} \cup \llbracket I \rrbracket_{uc}$	$\varphi_7 = false$	$\mathbb{J} = \mathbb{N}$

6 where $\bar{\sigma}(\omega)^{uc} = \neg \llbracket \sigma(\omega) \rrbracket_{uc}^{-1}$.

7 *Proof.* See <http://github.com/ssfm-up/TVMC/raw/master/SCICOProofs.pdf>

8

9

10 As we can see, Theorem 3 is the SAT-based version of Theorem 2. The
 11 new theorem allows us to implicitly determine causes of violation of a model
 12 checking problem and to derive corresponding constraints via unsatisfiable core
 13 extraction. Subsequently, we outline some of the cases that are considered in
 14 Theorem 3. If the unsatisfiable core is *initial state-independent* (ϕ_2 with $\llbracket I \rrbracket_{uc} =$
 15 \emptyset), then the extracted spurious segment constraint $\bar{\sigma}(\omega)^{uc}$ is an admissible
 16 constraint in all bound iterations $k + j$. If the unsatisfiable core is *property-*
 17 *independent* (ϕ_4 with $\llbracket \psi \rrbracket_{uc} = \emptyset$), then the j -increment of the extracted spurious
 18 segment constraint $\bar{\sigma}(\omega)^{uc}$ is an admissible constraint in all bound iterations $k +$
 19 j . If the unsatisfiable core is *fully-independent*, i.e. *both* initial state-independent
 20 and property-independent (ϕ_5), then all l -increments of the extracted spurious
 21 segment constraint $\bar{\sigma}(\omega)^{uc}$ with $0 \leq l \leq j$ are admissible constraints in all bound
 22 iterations $k + j$.

23 The admissible use, reuse and incrementation of spurious segment con-
 24 straints can be straightforwardly integrated into a combination of the algorithms
 25 *k-IND* and *SATWRC* by storing all reusable constraints in a global set Σ_{global}
 26 and by initialising Σ_k with Σ_{global} in each bound iteration. Note that the index
 27 increments of constraints in Σ_{global} may be parameterised with regard to the
 28 current bound. Thus, a constraint $\bar{\sigma}(\omega)^{uc}$ in iteration k may be adapted to
 29 $\bar{\sigma}(\omega)_{j+1}^{uc}$ in iteration $k + 1$. In the subsequent section we introduce the imple-
 30 mentation of our approach and we present experimental results.

1 9. Implementation and Experiments

2 We have prototypically implemented our novel refinement approach on top of
3 the SAT-based three-valued bounded model checker TVAMCUS². Our tool takes
4 a concurrent system Sys within integer arithmetic as a first input. It supports
5 almost all control structures of the C language as well as *int*, *bool* and *semaphore*
6 as data types. The second input is a temporal logic formula that is either of the
7 form \mathbf{G} *safe* (safety) or \mathbf{GF} *progress* (liveness) where *safe* and *progress* are pred-
8 icate expressions over the control flow or the variables of the system. The tool
9 employs a three-valued abstractor [1] that automatically builds abstract control
10 flow graphs corresponding to the system and an initial predicate set A that
11 covers the predicates referenced in the temporal logic formula to be checked.
12 In case the input is a liveness formula, we apply the state recording translation
13 to the abstract control flow graphs, which reduces liveness checking to safety
14 checking. Hence, the formula to be evaluated will be always of the form \mathbf{G} *safe*.
15 TVAMCUS iterates over the bound starting with $k = 0$. In each bound iteration,
16 the three-valued bounded model checking problems $A[base]_k$ and $A[step]_{k+1}$ cor-
17 responding to the inputs are encoded into propositional logic. For each, the base
18 case and the inductive step we run our refinement algorithm until both invo-
19 cations yield a definite result. Within TVAMCUS we use an implementation of
20 MINISAT [20] for processing the satisfiability problems and for extracting unsat-
21 isfiability cores. If the base case yields *true*, we have detected a property violation.
22 If both the base case and the inductive step yield *false*, we have proven that
23 no property violation exists. Otherwise our tool proceeds to the next bound
24 iteration and repeats the computations for the incremented bound. Generated
25 spurious segment constraints are reused between bound iterations and between
26 the base case and the inductive step according to our results in terms of the
27 admissible reuse of constraints.

28 In our experimental evaluation, we considered implementations of mutual
29 exclusion algorithms for multiple processes, namely Dijkstra’s mutual exclusion
30 algorithm and Lamport’s bakery algorithm. Both algorithms ensure *mutual*
31 *exclusion* and *progress* in the sense that some process will be always able to
32 enter the critical section. While Dijkstra’s algorithm is prone to *starvation*,
33 Lamport’s algorithm ensures starvation-freedom. Moreover, we considered a
34 *deadlock*-prone instantiation of a semaphore-based dining philosopher system
35 with ten philosophers and forks. We checked safety and liveness properties of
36 the systems under the assumption of weak fairness in the case of liveness. In
37 experiments we compared our model checker TVAMCUS with SPIN [15]. SPIN
38 is an established model checking tool for verifying temporal logic properties of
39 concurrent systems under fairness. Thus, TVAMCUS and SPIN focus on similar
40 kinds of verification tasks. We employed our TVAMCUS model checker in two
41 different modes. In TVAMCUS-AR we used standard abstraction refinement,
42 whereas in TVAMCUS-WRC we used our novel refinement approach with mul-

²available at <https://github.com/siocnarff/tvamacus>

1 tuple models and constraint generation. The experiments were conducted on a
 2 2.6 GHz Intel Core i5 system with 8 GB. The experimental results are depicted
 3 in Table 2. A “✓” indicates that the property holds for the system, whereas a
 4 “✗” indicates that the property does not hold. Moreover, a “-” means that the
 5 model checker did not solve the task within 2 hours.

Table 2: Experimental results

case study	processes	TVAMCUS-AR		TVAMCUS-WRC		SPIN
		maximum predicates	time	maximum predicates	time	time
DIJKSTRA	2	10	2.0s	8	0.8s	0.1s
MUTUAL EXCLUSION (✓)	3	15	9.2s	12	3.2s	0.2s
	4	20	19.9s	16	4.8s	3.9s
DIJKSTRA	2	16	14.2s	10	1.3s	0.3s
PROGRESS (✓)	3	24	227s	15	9.3s	4.8s
	4	32	-	20	178s	166s
DIJKSTRA	2	16	14.7s	10	2.8s	0.4s
STARVATION-FREEDOM (✗)	3	24	806s	15	20.5s	9.1s
	4	32	-	20	343s	307s
LAMPORT	2	18	248s	10	12.6s	17.7s
PROGRESS (✓)	3	30	-	18	169s	-
PHILOSOPHERS						
DEADLOCK-FREEDOM (✗)	10	30	98.1s	20	54.6s	93.6s

6 If we compare the results for TVAMCUS-AR and TVAMCUS-WRC we can ob-
 7 serve that the multi-model approach allows to significantly reduce the maxi-
 8 mum number of predicates in the problems to be solved. The reduced amount
 9 of predicates results from the fact that TVAMCUS-WRC makes use of a global
 10 model that is solely defined over the set of initial predicates, and of local mod-
 11 els where the initial predicates are fixed by an unconfirmed witness and only
 12 the refinement predicates contribute to the state space complexity. Constraints
 13 derived from local models are added to the global model without increasing its
 14 state space complexity. Operating on smaller models under TVAMCUS-WRC
 15 results in considerably faster model checking in all experiments in comparison
 16 to TVAMCUS-AR. If we compare the results for TVAMCUS-WRC and SPIN for
 17 the DIJKSTRA case studies we can see that SPIN clearly outperforms TVAMCUS-
 18 WRC for smaller instances with two and three processes. In the case of four pro-
 19 cesses SPIN is still faster than TVAMCUS-WRC but the differences are less sig-
 20 nificant. The LAMPORT and PHILOSOPHERS case studies show that TVAMCUS-
 21 WRC is even capable of outperforming SPIN for the specified instances. These
 22 promising results demonstrate that our multi-model approach with constraint
 23 reuse can compete with existing tools. We have ideas for further enhancing our
 24 approach that we discuss in the outlook of this article.

1 10. Related Work

2 Some of the earliest work on the application of three-valued reasoning about
3 software specifications and their properties is [21]. There, however, no reason-
4 ing algorithm was provided to deal with the *unknown* in a constructive manner.
5 More recent work on three-valued reasoning about system specifications can be
6 found in [22], whereby techniques of *theorem-proving* are used additionally to
7 deal with *unknown* model checking results. Other interesting applications of
8 three-valued model checking can be found in the domain of *multi-agent systems*
9 [23], whereas model checking of temporal logic properties with *multi-valued* log-
10 ics is described in [24].

11 Our verification technique is related to existing approaches for improving the
12 classical *abstract-check-refine* paradigm [4, 25]. *Lazy abstraction* [26, 27, 28, 29]
13 is a concept that builds and refines a single abstract model where different
14 parts of the model exhibit different degrees of precision. This is achieved by
15 adding refinement predicates only at parts where they are required for proving
16 the spuriousness of witnesses. The major difference to our approach is that we
17 work with one full and multiple partial models. Only the partial models are
18 refined in order to *prove* whether a particular witness is spurious or not. In
19 the full model we take proven spuriousness as a *fact* in order to prune the state
20 space. The separation of proving and eliminating spuriousness enables us to
21 conduct verification on smaller models in comparison to lazy abstraction where
22 only a single model is used. Another refinement strategy for SAT-based model
23 checking, which also uses a combination of over- and under-approximations in
24 order to solve reachability problems can be found in [30, 31].

25 The propositional logic encoding of bounded model checking problems has
26 been initially introduced in [32]. An improved encoding that is linear in both
27 the size of the formula to be checked and the length of the bound has been
28 presented in [33]. Both techniques are incomplete. Liveness model checking
29 is supported by explicitly evaluating liveness formulae on lasso-shaped paths.
30 In our approach we enabled complete liveness checking by reducing liveness to
31 reachability via state recording [13] and by employing induction. Earliest ap-
32 proaches to complete bounded model checking via *k*-induction can be found
33 in [11, 34] where induction is used for the verification of safety properties of
34 finite-state systems. *k*-induction for infinite-state systems such as communi-
35 cation protocols and timed automata has been proposed in [35]. In [36] we
36 can find the use of counterexample-guided *k*-induction for bug detection, where
37 counterexamples produced from over-approximating the loops are exploited to
38 shorten the number of steps that are required to find bugs. Other upper-bound
39 considerations, also with the purpose of making bounded model checking com-
40 plete, can be found in [37].

41 Another related approach is *local abstraction refinement* [38] which extends
42 the lazy abstraction idea. The technique also adds predicates only to relevant
43 parts of the model. While a new predicate typically splits an abstract state in
44 two refined states, local abstraction refinement uses heuristics for determining
45 whether a single refined state is sufficient for the underlying verification task.

1 This enables smaller state spaces. The approach is still based on a single model,
2 and thus, does not have the same state space reduction capabilities as our multi-
3 model approach.

4 Our work also is related to *conditional model checking* (CMC) [39], which
5 reformulates model checking as follows: If model checking fails (due to state
6 explosion) to fully prove or disprove the property of interest, then it at least
7 returns a condition under which the property holds. This allows for a sequential
8 combination of model checking runs where a first run generates a condition and a
9 second run checks whether the condition holds. Our approach can be regarded as
10 an application and generalisation of the CMC idea in the context of abstraction
11 refinement. We take unconfirmed witnesses as conditions for our partial models
12 and we use conditions for excluding spurious witnesses in the full model.

13 Whereas our approach exploits control flow knowledge for the purpose of
14 abstraction refinement in the context of SAT solving, the exploitation of control
15 flow knowledge to obtain better results in SMT solving is described in [40],
16 whereas the approach of [41] is designed to verify program properties by way of
17 CHC-solving with constrained Horn clauses. An automatic on-the-fly decompo-
18 sition of large specifications into their most interesting or most relevant parts
19 can be found in [42].

20 The practical applicability of SAT-based model checking for various applica-
21 tion purposes is well known, for instance [43] in the railway domain. While we
22 are using SAT solvers to reason about temporal logic properties of concurrent
23 systems, it is in principle possible to use SAT solving also for reasoning about
24 other modal logic system properties that are not temporal e.g. [44].

25 In a wider context one can observe that model checking via SAT solving,
26 with potentially *unknown* results for any α -interpretation of the given logical
27 atoms, is somewhat similar to the path coverage problem in software testing
28 where different value interpretations of the input variables of a program lead
29 to different execution paths being covered, whereby the problem arises which
30 execution paths remain *unknown* for a given set of input values. In [45] we can
31 find an incremental method that generates input data from atomic predicates in
32 software specifications with the aim of achieving full path coverage. Somewhat
33 similar to our approach, distance-based reasoning is applied in [45] to incremen-
34 tally refine the generation of input data leading to better path coverage in the
35 software model under analysis.

36 11. Conclusion and Outlook

37 We presented an iterative abstraction refinement technique for the verifica-
38 tion of temporal logic properties of concurrent software systems. The novelty
39 of our approach is that we use separate models for producing abstract witness
40 paths and for checking whether witnesses are definite or spurious. Our local
41 models are restricted to refinements of particular witnesses only. The abstract
42 state space of our global model is pruned via constraints derived from local
43 models. We hereby gain precision in the global model without increasing its
44 state space. Our multi-model approach allows for a significant reduction of the

1 state space complexity in comparison to single-model approaches. It comes at
2 the cost of an increased number of constraint generation iterations. Our *con-*
3 *straint strengthening* concept enables us to diminish this number, which gives
4 us a space- and time-efficient verification technique. Our approach employs
5 *satisfiability-based bounded model checking* [46], and thus, profits from the ca-
6 pabilities of today’s SAT solvers. Since bounded model checking is inherently
7 incomplete, we integrated the *k-induction principle* [47] into our verification
8 technique. This iterative approach reduces an unbounded model checking prob-
9 lem to two bounded model checking problems, which enables us to perform
10 *complete* verification via satisfiability solving. We showed that constraints for
11 ruling out spurious behaviour can be generated via *unsatisfiable core extraction*
12 [8] and we introduced a concept for the admissible *reuse of constraints* between
13 bound iterations. In general, *k-induction* is limited to model checking *safety*
14 properties. We additionally enabled *liveness* checking by applying the *state*
15 *recording translation* [13] to the systems to be verified, which reduces liveness
16 to safety checking. In experiments we received promising performance results
17 with our new approach.

18 As future work we intend to enhance our constraint strengthening concept
19 based on ideas adopted from *symmetry reduction* [48] and *partial-order reduc-*
20 *tion* [49]. Since we focus on the verification of concurrent *systems*, these systems
21 typically exhibit a considerable amount of symmetry. If a constraint has been
22 generated that reveals that a *particular* pair of processes will never be at a
23 certain location at the same time, this observation may be transferable to *ar-*
24 *bitrary* pairs based on symmetry arguments. Similarly, certain behaviour may
25 be spurious, independent of the order in which the processes execute their op-
26 erations. Thus, a constraint that does not only exclude one but all orders that
27 lead to spuriousness may be admissible as well. We are working on a concept
28 for detecting symmetric and order-independent constraints that will allow us
29 to rule out spurious behaviour on a larger scale. Moreover, we plan to extend
30 our verification technique to systems beyond linear integer arithmetic and to
31 multi-agent systems.

32 References

- 33 [1] J. Schriebl, H. Wehrheim, D. Wonisch, Three-valued spotlight abstractions,
34 in: A. Cavalcanti, D. R. Dams (Eds.), FM 2009: Formal Methods, Vol.
35 5850 of LNCS, Springer Berlin Heidelberg, Berlin, Heidelberg, 2009, pp.
36 106–122.
- 37 [2] G. Bruns, P. Godefroid, Model checking partial state spaces with 3-valued
38 temporal logics, in: N. Halbwachs, D. Peled (Eds.), Computer Aided Verifi-
39 cation, Vol. 1633 of LNCS, Springer Berlin Heidelberg, Berlin, Heidelberg,
40 1999, pp. 274–287.
- 41 [3] N. Timm, S. Gruner, Three-valued bounded model checking with cause-
42 guided abstraction refinement, Science of Computer Programming 175

- 1 (2019) 37 – 62. doi:<https://doi.org/10.1016/j.scico.2019.02.002>.
2 URL [http://www.sciencedirect.com/science/article/pii/
3 S0167642319300206](http://www.sciencedirect.com/science/article/pii/S0167642319300206)
- 4 [4] E. Clarke, O. Grumberg, S. Jha, Y. Lu, H. Veith, Counterexample-guided
5 abstraction refinement, in: E. A. Emerson, A. P. Sistla (Eds.), Computer
6 Aided Verification, Vol. 1855 of LNCS, Springer Berlin Heidelberg, 2000,
7 pp. 154–169.
- 8 [5] N. Timm, S. Gruner, Abstraction refinement with path constraints for 3-
9 valued bounded model checking, in: C. Artho, P. C. Ölveczky (Eds.), For-
10 mal Techniques for Safety-Critical Systems, Springer International Pub-
11 lishing, Cham, 2019, pp. 139–157.
- 12 [6] O. Strichman, Accelerating bounded model checking of safety properties,
13 Formal Methods in System Design 24 (1) (2004) 5–24. doi:10.1023/B:
14 FORM.0000004785.67232.f8.
15 URL <https://doi.org/10.1023/B:FORM.0000004785.67232.f8>
- 16 [7] N. Timm, S. Gruner, M. Harvey, A bounded model checker for three-valued
17 abstractions of concurrent software systems, in: L. Ribeiro, T. Lecomte
18 (Eds.), Formal Methods: Foundations and Applications, Vol. 10090 of
19 LNCS, Springer International Publishing, Cham, 2016, pp. 199–216.
- 20 [8] A. Nadel, Boosting minimal unsatisfiable core extraction, in: Proceedings
21 of the 2010 Conference on Formal Methods in Computer-Aided Design,
22 FMCAD '10, FMCAD Inc, Austin, TX, 2010, pp. 221–229.
- 23 [9] D. Kroening, J. Ouaknine, O. Strichman, T. Wahl, J. Worrell, Linear com-
24 pleteness thresholds for bounded model checking, in: G. Gopalakrishnan,
25 S. Qadeer (Eds.), Computer Aided Verification, Springer Berlin Heidelberg,
26 Berlin, Heidelberg, 2011, pp. 557–572.
- 27 [10] A. Nadel, V. Ryvchin, O. Strichman, Ultimately incremental sat, in:
28 C. Sinz, U. Egly (Eds.), SAT 2014, Springer International Publishing, 2014,
29 pp. 206–218.
- 30 [11] N. Een, N. Sörensson, Temporal induction by incremental sat solving, Elec-
31 tronic Notes in Theoretical Computer Science 89 (4) (2003) 543 – 560, bMC
32 2003. doi:[https://doi.org/10.1016/S1571-0661\(05\)82542-3](https://doi.org/10.1016/S1571-0661(05)82542-3).
- 33 [12] N. Timm, S. Gruner, M. Harvey, Constraint reusing and k-induction for
34 three-valued bounded model checking, in: T. Massoni, M. R. Mousavi
35 (Eds.), Formal Methods: Foundations and Applications, Springer Inter-
36 national Publishing, Cham, 2018, pp. 126–143.
- 37 [13] A. Biere, C. Artho, V. Schuppan, Liveness checking as safety checking,
38 Electronic Notes in Theoretical Computer Science 66 (2) (2002) 160 –
39 177, fMICS'02, 7th International ERCIM Workshop in Formal Meth-
40 ods for Industrial Critical Systems (ICALP 2002 Satellite Workshop).

- 1 doi:[https://doi.org/10.1016/S1571-0661\(04\)80410-9](https://doi.org/10.1016/S1571-0661(04)80410-9).
2 URL [http://www.sciencedirect.com/science/article/pii/
3 S1571066104804109](http://www.sciencedirect.com/science/article/pii/S1571066104804109)
- 4 [14] L. Moura, N. Bjørner, Z3: An efficient SMT solver, in: C. R. Ramakrishnan,
5 J. Rehof (Eds.), Tools and Algorithms for the Construction and Analysis of
6 Systems, Vol. 4963 of Lecture Notes in Computer Science, Springer-Verlag
7 Berlin Heidelberg, 2008, pp. 337–340. doi:10.1007/978-3-540-78800-3_
8 24.
9 URL http://dx.doi.org/10.1007/978-3-540-78800-3_24
- 10 [15] G. J. Holzmann, The model checker spin, IEEE Transactions on software
11 engineering 23 (5) (1997) 279–295.
- 12 [16] N. Timm, H. Wehrheim, M. Czech, Heuristic-guided abstraction re-
13 finement for concurrent systems, in: T. Aoki, K. Taguchi (Eds.),
14 ICFEM, Vol. 7635 of LNCS, Springer, 2012, pp. 348–363. doi:10.1007/
15 978-3-642-34281-3_25.
- 16 [17] M. Fitting, Kleene’s three valued logics and their children, Fundamenta
17 Informaticae 20 (1-3) (1994) 113–131.
- 18 [18] H. Wehrheim, Bounded model checking for partial Kripke structures, in:
19 J. S. Fitzgerald, A. E. Haxthausen, H. Yenigun (Eds.), Theoretical As-
20 pects of Computing — ICTAC 2008, Vol. 5160 of LNCS, Springer Berlin
21 Heidelberg, 2008, pp. 380–394.
- 22 [19] A. Nadel, Boosting minimal unsatisfiable core extraction, in: Proceedings
23 of the 2010 Conference on Formal Methods in Computer-Aided Design,
24 FMCAD ’10, FMCAD Inc, Austin, TX, 2010, pp. 221–229.
25 URL <http://dl.acm.org/citation.cfm?id=1998496.1998537>
- 26 [20] N. Sorensson, N. Een, Minisat v1. 13-a sat solver with conflict-clause min-
27 imization, SAT 2005 (53) (2005) 1–2.
- 28 [21] D. Kourie, An approach to defining abstractions, refinements and enrich-
29 ments.
- 30 [22] A. Bernasconi, C. Menghi, P. Spoletini, L. D. Zuck, C. Ghezzi, From model
31 checking to a temporal proof for partial models, in: A. Cimatti, M. Sirjani
32 (Eds.), Software Engineering and Formal Methods, Springer International
33 Publishing, Cham, 2017, pp. 54–69.
- 34 [23] A. Lomuscio, J. Michaliszyn, Verifying multi-agent systems by model check-
35 ing three-valued abstractions, in: Proceedings of the 2015 International
36 Conference on Autonomous Agents and Multiagent Systems, AAMAS ’15,
37 International Foundation for Autonomous Agents and Multiagent Systems,
38 Richland, SC, 2015, pp. 189–198.
39 URL <http://dl.acm.org/citation.cfm?id=2772879.2772907>

- 1 [24] Y. Li, M. Droste, L. Lei, Model checking of linear-time properties
2 in multi-valued systems, *Information Sciences* 377 (2017) 51 – 74.
3 doi:<https://doi.org/10.1016/j.ins.2016.10.030>.
4 URL [http://www.sciencedirect.com/science/article/pii/
5 S0020025516313548](http://www.sciencedirect.com/science/article/pii/S0020025516313548)
- 6 [25] S. Shoham, O. Grumberg, 3-valued abstraction: More precision at less cost,
7 *Information and Computation* 206 (11) (2008) 1313–1333.
- 8 [26] T. A. Henzinger, R. Jhala, R. Majumdar, G. Sutre, Lazy abstraction, in:
9 *Proceedings of the 29th ACM SIGPLAN-SIGACT Symposium on Princi-
10 ples of Programming Languages, POPL '02*, ACM, New York, NY, USA,
11 2002, pp. 58–70.
- 12 [27] K. Madhukar, M. Srivas, B. Wachter, D. Kroening, R. Metta, Verifying
13 synchronous reactive systems using lazy abstraction, in: *2015 Design, Au-
14 tomation Test in Europe Conference Exhibition (DATE)*, 2015, pp. 1571–
15 1574.
- 16 [28] R. Degiovanni, P. Ponzio, N. Aguirre, M. Frias, Improving lazy abstraction
17 for SCR specifications through constraint relaxation, *Software Testing, Ver-
18 ification and Reliability* 28 (2) (2018) e1657.
- 19 [29] K. Hsu, R. Majumdar, K. Mallik, A. Schmuck, Lazy abstraction-based
20 control for reachability, *CoRR* abs/1804.02722. [arXiv:1804.02722](https://arxiv.org/abs/1804.02722).
- 21 [30] J. Li, S. Zhu, Y. Zhang, G. Pu, M. Y. Vardi, Approximate reachability,
22 *arXiv preprint arXiv:1611.04946*.
- 23 [31] J. Li, S. Zhu, Y. Zhang, G. Pu, M. Y. Vardi, Safety model checking with
24 complementary approximations, in: *Proceedings of the 36th International
25 Conference on Computer-Aided Design, ICCAD '17*, IEEE Press, Piscat-
26 away, NJ, USA, 2017, pp. 95–100.
27 URL <http://dl.acm.org/citation.cfm?id=3199700.3199713>
- 28 [32] A. Biere, A. Cimatti, E. Clarke, Y. Zhu, Symbolic model checking without
29 bdds, in: W. R. Cleaveland (Ed.), *Tools and Algorithms for the Construc-
30 tion and Analysis of Systems*, Springer Berlin Heidelberg, Berlin, Heidel-
31 berg, 1999, pp. 193–207.
- 32 [33] T. Latvala, A. Biere, K. Heljanko, T. Junttila, Simple bounded ltl model
33 checking, in: A. J. Hu, A. K. Martin (Eds.), *Formal Methods in Computer-
34 Aided Design*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2004, pp.
35 186–200.
- 36 [34] M. Sheeran, S. Singh, G. Stålmarck, Checking safety properties using in-
37 duction and a sat-solver, in: W. A. Hunt, S. D. Johnson (Eds.), *Formal
38 Methods in Computer-Aided Design*, Springer Berlin Heidelberg, Berlin,
39 Heidelberg, 2000, pp. 127–144.

- 1 [35] L. de Moura, H. Rueß, M. Sorea, Bounded model checking and induction:
2 From refutation to verification, in: W. A. Hunt, F. Somenzi (Eds.), Com-
3 puter Aided Verification, Springer Berlin Heidelberg, Berlin, Heidelberg,
4 2003, pp. 14–26.
- 5 [36] M. Y. R. Gadelha, L. C. Cordeiro, D. A. Nicole, Counterexample-guided k-
6 induction verification for fast bug detection, CoRR abs/1706.02136. arXiv:
7 1706.02136.
8 URL <http://arxiv.org/abs/1706.02136>
- 9 [37] I. Konnov, H. Veith, J. Widder, On the completeness of bounded
10 model checking for threshold-based distributed algorithms: Reach-
11 ability, Information and Computation 252 (2017) 95 – 109.
12 doi:<https://doi.org/10.1016/j.ic.2016.03.006>.
13 URL [http://www.sciencedirect.com/science/article/pii/
14 S0890540116000432](http://www.sciencedirect.com/science/article/pii/S0890540116000432)
- 15 [38] H. Fecher, S. Shoham, Local abstraction-refinement for the mu-calculus,
16 in: D. Bošnački, S. Edelkamp (Eds.), Model Checking Software, Vol. 4595
17 of LNCS, Springer Berlin Heidelberg, Berlin, Heidelberg, 2007, pp. 4–23.
- 18 [39] D. Beyer, T. A. Henzinger, M. E. Keremoglu, P. Wendler, Conditional
19 model checking: A technique to pass information between verifiers, in:
20 Proceedings of the ACM SIGSOFT 20th International Symposium on the
21 Foundations of Software Engineering, FSE '12, ACM, New York, NY, USA,
22 2012, pp. 57:1–57:11.
- 23 [40] J. Chen, F. He, Control flow-guided smt solving for program verification,
24 in: Proceedings of the 33rd ACM/IEEE International Conference on Auto-
25 mated Software Engineering, ASE 2018, ACM, New York, NY, USA, 2018,
26 pp. 351–361. doi:10.1145/3238147.3238218.
27 URL <http://doi.acm.org/10.1145/3238147.3238218>
- 28 [41] E. DE ANGELIS, F. FIORAVANTI, A. PETTOROSSO, M. PROIETTI,
29 Predicate pairing for program verification, Theory and Practice of Logic
30 Programming 18 (2) (2018) 126?166. doi:10.1017/S1471068417000497.
- 31 [42] S. Apel, D. Beyer, V. Mordan, V. Mutilin, A. Stahlbauer, On-the-fly de-
32 composition of specifications in software model checking, in: Proceedings
33 of the 2016 24th ACM SIGSOFT International Symposium on Foundations
34 of Software Engineering, FSE 2016, ACM, New York, NY, USA, 2016, pp.
35 349–361. doi:10.1145/2950290.2950349.
36 URL <http://doi.acm.org/10.1145/2950290.2950349>
- 37 [43] P. James, Sat-based model checking and its applications to train control
38 systems, Swansea University.
- 39 [44] T. Caridroit, J.-M. Lagniez, D. Le Berre, T. de Lima, V. Montmirail, A
40 sat-based approach for solving the modal logic s5-satisfiability problem, in:
41 Thirty-First AAAI Conference on Artificial Intelligence, 2017.

- 1 [45] P. Zhao, S. Liu, A software tool to support the “vibration” method, in:
2 C. Tian, F. Nagoya, S. Liu, Z. Duan (Eds.), *Structured Object-Oriented*
3 *Formal Language and Method*, Springer International Publishing, Cham,
4 2018, pp. 171–186.
- 5 [46] A. Biere, A. Cimatti, E. M. Clarke, O. Strichman, Y. Zhu, Bounded model
6 checking., *Handbook of Satisfiability* 185 (2009) 457–481.
- 7 [47] T. Wahl, *The k-induction principle*, Northeastern University, College of
8 Computer and Information Science (2013) 1–2.
- 9 [48] A. Miller, A. Donaldson, M. Calder, Symmetry in temporal logic model
10 checking, *ACM Computing Surveys (CSUR)* 38 (3) (2006) 8.
- 11 [49] D. Peled, *Partial-Order Reduction*, Springer International Publishing,
12 Cham, 2018, pp. 173–190. doi:10.1007/978-3-319-10575-8_6.
13 URL https://doi.org/10.1007/978-3-319-10575-8_6