

Model Checking as Program Verification by Abstract Interpretation

Paolo Baldan ✉ 

Department of Mathematics, University of Padua, Italy

Roberto Bruni ✉ 

Department of Computer Science, University of Pisa, Italy

Francesco Ranzato ✉ 

Department of Mathematics, University of Padua, Italy

Diletta Rigo ✉ 

Department of Mathematics, University of Padua, Italy

Abstract

Abstract interpretation offers a powerful toolset for static analysis, tackling precision, complexity and state-explosion issues. In the literature, state partitioning abstractions based on (bi)simulation and property-preserving state relations have been successfully applied to abstract model checking. Here, we pursue a different track in which model checking is seen as an instance of program verification. To this purpose, we introduce a suitable language – called **MOKA** (for **MO**del checking as abstract interpretation of Kleene Algebras) – which is used to encode temporal formulae as programs. In particular, we show that (universal fragments of) temporal logics, such as ACTL or, more generally, universal μ -calculus can be transformed into MOKA programs. Such programs return all and only the initial states which violate the formula. By applying abstract interpretation to MOKA programs, we pave the way for reusing more general abstractions than partitions as well as for tuning the precision of the abstraction to remove or avoid false alarms. We show how to perform model checking via a program logic that combines under-approximation and abstract interpretation analysis to avoid false alarms. The notion of locally complete abstraction is used to dynamically improve the analysis precision via counterexample-guided domain refinement.

2012 ACM Subject Classification Theory of computation \rightarrow Modal and temporal logics; Theory of computation \rightarrow Verification by model checking; Theory of computation \rightarrow Program verification; Theory of computation \rightarrow Programming logic; Theory of computation \rightarrow Abstraction

Keywords and phrases ACTL, μ -calculus, model checking, abstract interpretation, program analysis, local completeness, abstract interpretation repair, domain refinement, Kleene algebra with tests

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2025.8

Related Version *Extended Version*: <https://arxiv.org/abs/2506.05525> [1]

Funding This research was partially funded by the *Italian MUR*, under the PRIN 2022 PNRR project no. P2022HXNSC on “Resource Awareness in Programming: Algebra, Rewriting, and Analysis”, and the PNRR project *SEcurity and RIghts In the CyberSpace* (SERICS, PE00000014 – CUP H73C2200089001), by the INdAM-GNCS Projects RISICO (CUP E53C22001930001) and MARQ (CUP E53C24001950001), by a *WhatsApp Research Award* and by an *Amazon Research Award* for *AWS Automated Reasoning*.

1 Introduction

Abstraction is a fundamental craft for mastering complexity. In model checking, abstraction-guided space reduction allows to mitigate the well-known state explosion problem [15]. Abstract interpretation [16, 17] is the de facto standard framework for designing sound analyses. The idea of applying abstract interpretation to model checking has been widely



© Paolo Baldan, Roberto Bruni, Francesco Ranzato, and Diletta Rigo;
licensed under Creative Commons License CC-BY 4.0

36th International Conference on Concurrency Theory (CONCUR 2025).

Editors: Patricia Bouyer and Jaco van de Pol; Article No. 8; pp. 8:1–8:20

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

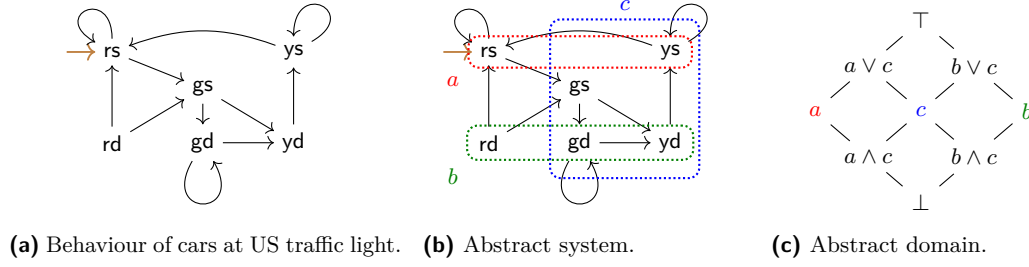
investigated, e.g., [2–5, 10–12, 14, 19–23, 25, 26, 28–30, 33, 35, 37, 38, 40, 44–47, 49, 50, 55] (§ 7 accounts for more closely related work). These approaches often rely on over-approximations that preserve some logical properties, like state partitioning abstractions or simulation-preserving relations.

Objective. In this work, we pursue a different approach, which consists in recasting model checking of temporal formulae directly in terms of program verification, for which the whole tool-suite of abstract interpretation is readily available. To this aim, we exploit an instance of a Kleene algebra with tests and fixpoints (KAF) [36], whose primitives allow to map each temporal formula to a program that can single out all and only the counterexamples to the formula. This program can then be analysed through a sound abstract interpretation that over-approximates the concrete behaviour, so that all possible counterexamples are exposed. However, this over-approximation might not faithfully model the program behaviour in the abstract domain, thus possibly mixing true and false alarms. This lack of precision cannot happen when the abstract interpretation is complete [18]. However, completeness happens quite seldom, and even if, in principle, it can be achieved by refining the abstract domain [27], the most abstract domain refinement ensuring completeness is often way too concrete (it might well coincide with the concrete domain itself).

To remove false alarms, the validity of temporal formulae can be analysed by deriving certain judgements for the corresponding program in a suitable program logic. This enables the use of techniques borrowed from locally complete abstract interpretation [7, 9], where the over-approximation provided by the abstract domain is paired with an under-approximating specification, in the style of O’Hearn’s incorrectness logic [43]. This way, any alarm raised by an incorrectness logic proof corresponds to a true counterexample and local completeness guarantees that derivable judgements either exposes some true counterexamples (if any) or proves that there are none. Moreover, the failure of a proof obligation during inference can point out how to dynamically refine the underlying abstraction to enhance the precision and expressiveness of the analysis, a technique called abstract interpretation repair [8].

Methodology (and a toy example). Our main insight is to design a meta-programming language, called MOKA (MODEL checking as abstract interpretation of Kleene Algebras), where suitable temporal formulae can be mapped to. We show how this can be done for ACTL or, more generally, for the single variable μ -calculus without the existential diamond modality. MOKA is a language of regular commands, based on Kleene algebra with tests (KAT) [34] augmented with fixpoint operators (KAF) [36]. MOKA leverages a small set of primitives to extend and filter computation paths. They can be combined with the usual KAF operators of sequential composition, join (i.e., nondeterministic choice), Kleene iteration, and least fixpoint computation. The corresponding programs are then analysed by abstract interpretation.

The language MOKA operates on computation paths $\langle \sigma, \Delta \rangle$, where $\sigma \in \Sigma$ represents the current system state, and $\Delta \in \mathcal{P}(\Sigma)$ represents the set of traversed states. In the following, we informally overload the symbol σ to denote $\langle \sigma, \emptyset \rangle$. Several computation paths can be stacked and unstacked through the MOKA primitives **push** and **pop** for dealing with nested temporal formulae. A generic stack is denoted $\langle \sigma, \Delta \rangle :: S$ and we call σ its *current* state. Each temporal formula φ is mapped to a MOKA program $\llbracket \overline{\varphi} \rrbracket$ that intuitively computes counterexamples to φ , whence the bar over φ in the application of the encoding $\llbracket \cdot \rrbracket$. Letting $\llbracket \cdot \rrbracket$ denote the usual (collecting) denotational semantics, a key result is, therefore, that a state σ satisfies the formula φ iff the execution of $\llbracket \overline{\varphi} \rrbracket$ on σ returns the empty set. Of course, the results



■ **Figure 1** Variations on the traffic light example from [12].

can be generalised additively to sets of stacks, for which $\llbracket \bar{\varphi} \rrbracket$ filters out exactly those stacks whose current state satisfies φ . We can thus summarise the semantic correspondence between formulae satisfaction and execution of their MOKA programs by $\llbracket \llbracket \bar{\varphi} \rrbracket \rrbracket P = \{\sigma \in P \mid \sigma \not\models \varphi\}$. As a special case, it follows that $\sigma \models \varphi$ iff $\llbracket \llbracket \bar{\varphi} \rrbracket \rrbracket \{\sigma\} = \emptyset$. For example, for an atomic proposition p , we let $\llbracket \bar{p} \rrbracket \triangleq \neg p?$, where $\neg p?$ is a MOKA primitive that filters out those states where p is valid. As a further example, the MOKA program for the formula $\text{AG } \varphi$ (stating that φ along every possible path always holds) is $\llbracket \overline{\text{AG } \varphi} \rrbracket = \text{push}; \text{next}^*; \llbracket \bar{\varphi} \rrbracket; \text{pop}$.

The intuition is quite simple: next^* generates all the successors by iterating the next operation, and they are filtered by $\llbracket \bar{\varphi} \rrbracket$, to expose those which fail to satisfy φ . The operations push and pop manage the stack structure. In the example, to attain the semantic correspondence above, whenever a state δ such that $\delta \not\models \varphi$ is reachable from some initial state σ , we want $\llbracket \llbracket \overline{\text{AG } \varphi} \rrbracket \rrbracket \{\sigma\}$ to return σ and not δ : the operation push serves to stash σ when looking for δ , and the operation pop to discard δ and unstash σ , to conclude $\sigma \not\models \text{AG } \varphi$. To make the next introductory examples easier to follow, we suggest that the reader ignore the push and pop symbols, hence their faded rendering in MOKA programs.

To get a flavour of the proposed approach, we sketch an easy example, which is a variation of [12, Examples 3.4, 3.7]. Consider the transition system in Figure 1a, which models the behaviour of cars at a US traffic light. State names consist of two letters, denoting, respectively, the traffic light status, i.e. red, yellow, green, and the car behaviour, i.e. stop, drive. We model check the validity in the initial state rs of the safety property $\varphi = \text{AG } (\neg rd)$, stating that it will never happen that the light is red and the car is driving. This is obviously true for the system, since rd is an unreachable state. We therefore consider the MOKA program $\llbracket \bar{\varphi} \rrbracket = \text{push}; \text{next}^*; rd?; \text{pop}$, and compute its semantics $\llbracket \llbracket \bar{\varphi} \rrbracket \rrbracket \{rs\}$, which turns out to be the empty set, thus allowing us to conclude that $rs \models \varphi$. In fact, after a few concrete steps of iteration, $\llbracket rd? \rrbracket \llbracket \text{next}^* \rrbracket \{rs\} = \llbracket rd? \rrbracket \{rs, gs, gd, yd, ys\} = \emptyset$.

How the abstraction works. We show how any sound abstraction of state properties in $\wp(\Sigma)$ of the original system can be lifted to a range of sound abstractions on stacks with varying degrees of precision, in a way that MOKA programs can be analysed by an abstract interpreter, denoted by $\llbracket \cdot \rrbracket^\sharp$. The analysis output is always sound, meaning that the set of counterexamples is over-approximated. In particular, letting α denote the abstraction map to an abstract lattice, and \perp the bottom element of the corresponding abstract computational domain, if $\llbracket \llbracket \bar{\varphi} \rrbracket \rrbracket^\sharp \alpha(\{\sigma\}) = \perp$, then $\sigma \models \varphi$ holds. Vice versa, if $\llbracket \llbracket \bar{\varphi} \rrbracket \rrbracket^\sharp \alpha(\{\sigma\}) \neq \perp$, then the abstract analysis might raise a false alarm because, in general, abstractions are not complete.

Back to the previous example, let us consider the (non-partitioning) abstract domain $A \triangleq \{\perp_A, a \wedge c, b \wedge c, a, b, c, a \vee c, b \vee c, \top_A\}$ in Figure 1c, induced by the abstract properties a , b and c , whose concretizations are represented as dotted boxes in Figure 1b. The abstract interpreter computes $\llbracket \llbracket \overline{\text{AG } (\neg rd)} \rrbracket \rrbracket^\sharp \alpha(\{rs\}) = \llbracket \text{push}; \text{next}^*; rd?; \text{pop} \rrbracket^\sharp a = \perp$, thus allowing us

to conclude that $rs \models AG(\neg rd)$. Since rs has a transition to gs , we have $\llbracket rd? \rrbracket^\# \llbracket next^* \rrbracket^\# a = \llbracket rd? \rrbracket^\# (a \vee c) = \perp$. It is worth remarking that in the abstract domain A , computing $\llbracket next^* \rrbracket^\# a$ requires two iterations instead of four, as it happens in the concrete computation. Consider now the formula $\psi = AG(g \rightarrow AX d)$, namely, “whenever the semaphore is green, then at the next state the car is driving”. Spelling out the MOKA encoding, we have: $\llbracket \psi \rrbracket = \text{push}; \text{next}^*; g?; \text{push}; \text{next}; \neg d?; \text{pop}; \text{pop}$. Again, ψ holds for the concrete system in Figure 1a, but here the abstract interpretation is imprecise, because $\llbracket \text{push}; \text{next}^*; g? \rrbracket^\# \alpha(\{rs\}) = c::a$ and $\llbracket \text{push}; \text{next}; \neg d? \rrbracket^\# (c::a) = (a \vee c)::c::a$, so that $\llbracket \llbracket \psi \rrbracket \rrbracket^\# \alpha(\{rs\}) = a \neq \perp$, thus raising a false alarm. As explained above, we exploit **push/pop** operators to manage the stack structure of domain elements and to recover the (abstract) states from which the computation started when property violations are found. For instance, in the case illustrated above, once the stack $(a \vee c)::c::a$ is obtained, representing an abstract trace where the property fails, its initial abstract state, a , is recovered by applying two successive **pop** operations.

To eliminate false alarms, we leverage the concept of local completeness in abstract interpretation. Roughly, the idea consists in focusing on the (abstract) computation path produced by some input of interest, and then refining the abstraction only when needed to make it complete *locally* to such computation. More precisely, we apply a variation of local completeness logic (LCL) [7, 9], which is here extended to deal with fixpoint operators. The LCL proof system, parametrised by a generic state abstraction A , works with O’Hearn-like judgements $\vdash_A [P] \text{ r } [Q]$, where r is a program and P, Q denote state properties or, equivalently, the underlying sets of states satisfying those properties. The triple $\vdash_A [P] \text{ r } [Q]$ is valid when Q is an under-approximation of the states reachable by r from P while the abstraction of Q over-approximates such reachable states, and the abstract computation of r on the precondition P is locally complete. Roughly, this can be expressed as $Q \subseteq \llbracket \text{r} \rrbracket P \subseteq \gamma \circ \alpha(Q)$. In our setting, where the program $\llbracket \varphi \rrbracket$ associated with a formula φ “returns” all the counterexamples to the validity of φ , an inference of $\vdash_A [P] \llbracket \varphi \rrbracket [Q]$ shows that $Q \subseteq \{\sigma \in P \mid \sigma \not\models \varphi\} \subseteq \gamma \circ \alpha(Q)$, thus bounding the set of counterexamples to φ in P . Hence, if $Q \neq \emptyset$ then φ does not hold for each $\sigma \in Q \subseteq P$, while φ holds for all states $\sigma \in P \setminus (\gamma \circ \alpha(Q))$. If, instead, $Q = \emptyset$, then $\alpha(Q) = \perp$, and φ holds in the whole P . As LCL derivations cannot succeed with locally incomplete abstractions, if some proof obligation fails, the abstract domain needs to be “fixed” to achieve local completeness. This can be accomplished, e.g., by applying the abstraction repair techniques defined in [8].

For the toy traffic light example, we can derive in LCL the program triple $\vdash_A [\{rs\}] \llbracket \varphi \rrbracket [\emptyset]$, that confirms the validity of the formula φ . Vice versa, an attempt to derive the triple $\vdash_A [\{rs\}] \llbracket \psi \rrbracket [\emptyset]$ fails because of local incompleteness. Roughly, we can successfully derive $\vdash_A [\{rs\}] \text{next}^*; g?; [\{gs, gd\}]$, but then the execution of **next** on $\{gs, gd\}$ is not locally complete, because $\alpha(\llbracket \text{next} \rrbracket \{gs, gd\}) = \alpha(\{gd, yd\}) = b \wedge c$, while $\llbracket \text{next} \rrbracket^\# \alpha(\{gs, gd\}) = \llbracket \text{next} \rrbracket^\# c = a \vee c$. The abstraction repair procedure of [8] would then lead to refine the abstract domain by adding a new abstract element c_1 to represent the concrete set $\gamma(c_1) = \{gs, gd\}$. Since abstract domains must be closed under meets, the addition of c_1 will also entail the addition of $c_1 \wedge b$ such that $\gamma(c_1 \wedge b) = \{gd\}$. Letting $A_1 \triangleq A \cup \{c_1 \wedge b, c_1\}$, we can still derive $\vdash_{A_1} [\{rs\}] \text{next}^*; g?; [\{gs, gd\}]$, then $\vdash_{A_1} [\{gs, gd\}] \text{next} [\{gd, yd\}]$, and finally $\vdash_{A_1} [\{gd, yd\}] \neg d? [\emptyset]$, which can be composed together to conclude $\vdash_{A_1} [\{rs\}] \llbracket \psi \rrbracket [\emptyset]$. In fact, in A_1 we just have $\llbracket \llbracket \psi \rrbracket \rrbracket^\# \alpha_1(\{rs\}) = \perp$.

Original contribution. We set up a theoretical framework for the systematic reduction of model checking of temporal logics to program verification, thereby enabling the re-use of abstract interpretation techniques and verifiers either directly or with minimal effort. This framework consists of:

- MOKA, a meta-programming language in which different temporal logics can be encoded so that the MOKA code for a formula φ computes exactly the counterexamples to φ ;
- a systematic technique for lifting abstractions over state properties to abstractions suitable for the static analysis of MOKA programs;
- an extension of the LCL program logic to handle least fixpoints, introducing a form of “false alarm-guided” abstraction refinement loop in the analysis of MOKA programs.

Synopsis. In § 2 we provide some basics about abstract interpretation and the (fragments) of temporal logics considered in the paper. In § 3 we introduce the language MOKA, and then prove in § 4 the key results that relate formula satisfaction with program execution. In § 5 we define a general technique for deriving abstract domains for the static analysis of MOKA programs. In § 6 we showcase how local completeness logic reasoning can be exploited in our framework. In § 7 we discuss related work. Finally, in § 8 we draw some conclusions and sketch future avenues of research. Full proofs of our results and additional material are available in the extended version of this paper [1].

2 Background

A complete lattice is a poset (L, \leq_L) where every subset $X \subseteq L$ has both least upper bound (lub) and greatest lower bound (glb), denoted by $\bigvee_L X$ and $\bigwedge_L X$, respectively, with $\perp_L \triangleq \bigvee_L \emptyset$ and $\top_L \triangleq \bigwedge_L \emptyset$. When no ambiguities can arise, a lattice will be denoted as L and subscripts will be omitted.

Given two complete lattices L_1 and L_2 , a function $f : L_1 \rightarrow L_2$, is monotone if $x \leq_1 y$ implies $f(x) \leq_2 f(y)$, and additive (resp., co-additive) if it preserves arbitrary lub (resp., glb). Any monotone function $f : L \rightarrow L$ on a complete lattice has both a least and a greatest fixpoint, denoted by $\text{lfp}(f)$ and $\text{gfp}(f)$, respectively. The set of functions $f : S \rightarrow L$ from a set S to a complete lattice L , denoted L^S , forms a complete lattice when endowed with the pointwise order s.t. $f \leq g$ if for all $s \in S$, $f(s) \leq_L g(s)$. If L_1, L_2 are complete lattices, then $L_1 \times L_2$ is their product lattice endowed with the componentwise order s.t. $(x_1, x_2) \leq (y_1, y_2)$ if $x_1 \leq_1 y_1$ and $x_2 \leq_2 y_2$. We write L^n for the n -ary product of L with itself and $L^+ \triangleq \{\perp, \top\} \cup \bigcup_{n \geq 1} L^n$ for the complete lattice of non-empty finite sequences in L , ordered by $x_1 \dots x_n \leq y_1 \dots y_m$ if $n = m$ and $x_i \leq_L y_i$ for all $i \in [1, n]$, with top \top and bottom \perp .

2.1 Abstract Interpretation

Let us recall the basics of abstract interpretation [17] (see [16] for a thorough account).

Given two complete lattices C and A , called the *concrete* and the *abstract* domain, respectively, a *Galois connection* (GC) $\langle \alpha, \gamma \rangle : C \rightleftarrows A$ is a pair of functions $\alpha : C \rightarrow A$ and $\gamma : A \rightarrow C$ s.t. $\alpha(c) \leq_A a \Leftrightarrow c \leq_C \gamma(a)$ for any $c \in C$ and $a \in A$.

The function α is referred to as abstraction map and turns out to be additive, while γ is the concretization map which is always co-additive. Intuitively, any abstract element $a \in A$ such that $c \leq \gamma(a)$ is a sound over-approximation for the concrete value c , while the abstraction $\alpha(c)$ is the most precise over-approximation of c in the abstract domain A , i.e., $\alpha(c) = \bigwedge_C \{a \mid c \leq_C \gamma(a)\}$ holds. The notation $A_{\alpha, \gamma}$ denotes an abstract domain endowed with its underlying GC, and we will omit subscripts when α and γ are clear from the context.


(a) Program c .(b) Control flow graph of c .(c) Predicate abstraction \mathbb{P} .

Figure 2 Program from [4, Figure 1], with control flow graph and predicate abstraction domain.

► **Example 2.1** (Image adjunction). Given any function $f : X \rightarrow Y$, let $f^>$ and $f^<$ denote the direct and inverse image of f , respectively. The pair $\langle f^>, f^< \rangle : \mathcal{P}(X) \rightarrow \mathcal{P}(Y)$ is a GC that we refer to as the *image adjunction* (this is an instance of [24, Exercise 7.18]). \dashv

The class of abstract domains on C , denoted by $\text{Abs}(C) \triangleq \{A_{\alpha, \gamma} \mid \langle \alpha, \gamma \rangle : C \rightleftharpoons A\}$, can be preordered by the domain refinement relation: $A' \sqsubseteq A$ when $\gamma_A(A) \subseteq \gamma_{A'}(A')$.

► **Example 2.2** (Control flow graphs and predicate abstraction). Any program can be represented by its control flow graph (CFG). Let Var be a set of variables valued in \mathbb{V} and denote by $\text{Env} = \mathbb{V}^{\text{Var}}$ the set of environments. A CFG is a graph (N, E, s, e) , where N is a finite set of nodes, representing program points, $s, e \in N$ are the start and end nodes, respectively, and $E \subseteq N \times F \times N$ is a set of edges, labelled over a set of transfer (additive) functions $F \subseteq \mathcal{P}(\text{Env})^{\mathcal{P}(\text{Env})}$. For example, the program c in Figure 2a is decorated with program points $n \in N = \{s, 1, 2, 3, e\}$, has variables $\text{Var} = \{x, y, z, w\}$, and values ranging in the finite domain $\mathbb{V} = \mathbb{Z}_k$ of integers modulo a given $k > 0$. The CFG of c is depicted in Figure 2b.

Predicate abstraction allows to approximate program invariants [32]. Given a set of predicates $\text{Pred} \subseteq \mathcal{P}(\text{Env})$ (where any $p \in \text{Pred}$ has a representation $p = \{\rho \in \text{Env} \mid \rho \models p\}$), the predicate abstraction domain \mathbb{P} is defined by adding to Pred the complement predicates $\overline{\text{Pred}} \triangleq \{\bar{p} \mid p \in \text{Pred}\}$, and then by closing $\text{Pred} \cup \overline{\text{Pred}}$ under logical conjunction. We define a function $\pi : \mathcal{P}(\text{Env}) \rightarrow \mathbb{P}$ that associates to each set of environments the strongest predicate it satisfies. For example, given the two predicates $p \triangleq (z = 0)$ and $q \triangleq (x = y)$, the predicate abstraction domain \mathbb{P} induced by the set $\text{Pred} = \{p, q\}$ is depicted in Figure 2c. Correspondingly, the product abstraction \mathbb{P}^N allows us to represent the abstract state of the program as a function that associates to each program point n the strongest predicate in \mathbb{P} that holds at n . Hence, for instance, the set of all possible initial states of c , which is $\{(s, xyzw) \mid x, y, z, w \in \mathbb{V}\}$, is represented by the function $(s \mapsto \top, 1 \mapsto \perp, 2 \mapsto \perp, 3 \mapsto \perp, e \mapsto \perp)$, often written $(s \mapsto \top)$, omitting the program points which are mapped to \perp . \dashv

An abstract interpreter computes in the underlying abstract domain through correct (and effective) abstract approximations of concrete functions. Given $A_{\alpha, \gamma} \in \text{Abs}(C)$ and a function $f : C \rightarrow C$, an abstract function $f^\# : A \rightarrow A$ is a correct approximation of f if $\alpha \circ f \leq f^\# \circ \alpha$, and it is a *complete* approximation of f when $\alpha \circ f = f^\# \circ \alpha$. The *best correct approximation* (bca) of f in A , is defined as $f^A \triangleq \alpha \circ f \circ \gamma : A \rightarrow A$, and turns out to be the most precise correct abstraction, i.e., $f^A \leq f^\#$ for any other correct approximation $f^\#$ of f . When $f^\#$ is complete, then $f^\# = f^A$, thus making completeness an abstract domain property defined by the equation $\alpha \circ f = \alpha \circ f \circ \gamma \circ \alpha$. In program analysis, abstract domains are commonly

endowed with correct but incomplete abstract transfer functions. When completeness holds, the abstract interpreter is as precise as possible for the given abstract domain and cannot raise false alarms when verifying properties that are expressible in the abstract domain [27].

2.2 Transition Systems and Logics

Temporal logics are typically interpreted over unlabeled, finite, directed graphs, whose nodes and edges model states and transitions between them (i.e., state changes), respectively.

A *transition system* T is a tuple $(\Sigma, \mathbf{I}, \mathbf{P}, \rightarrow, \vdash)$, where Σ is a finite set of states ranged over by σ , $\mathbf{I} \subseteq \Sigma$ is the set of initial states, $\rightarrow \subseteq \Sigma \times \Sigma$ is the transition relation, \mathbf{P} is a (finite) set of atomic propositions, ranged over by p , and $\vdash \subseteq \Sigma \times \mathbf{P}$ is the satisfaction relation.

We write $\sigma \rightarrow \sigma'$ instead of $(\sigma, \sigma') \in \rightarrow$ and let $\text{post}(\sigma) = \{\sigma' \mid \sigma \rightarrow \sigma'\}$ denote the set of direct successors of $\sigma \in \Sigma$. As common in model checking [13], we consider systems whose transition relation is total, i.e., for all $\sigma \in \Sigma$ there exists $\sigma' \in \Sigma$ such that $\sigma \rightarrow \sigma'$.

A path is an infinite denumerable state sequence $(\sigma_i)_{i \in \mathbb{N}}$ such that $\sigma_i \rightarrow \sigma_{i+1}$ for all $i \in \mathbb{N}$.

► **Example 2.3** (Control flow graphs as transition systems). A CFG (N, E, s, e) can be viewed as a transition system with states $\Sigma = N \times \text{Env}$ and transition relation defined by $(n, \rho) \rightarrow (n', \rho')$ if there is an edge $(n, f, n') \in E$ such that $\rho' \in f(\{\rho\})$. Additionally, in order to make the transition relation total we add self-loops to all the states (e, ρ) involving the end node e . For instance, by using the shorthand $(n, xyzw)$ for the states in Example 2.2, the transition $(s, 0111) \rightarrow (1, 1101)$ is induced by the edge $(s, x := y \ z := 0, 1)$ in Figure 2b, while $(3, 1111) \rightarrow (e, 1111)$ by $(3, x = y, e)$ in Figure 2b.

Let us point out that the predicate abstraction in Example 2.2 naturally lifts to the power-set of states with codomain $A = \mathbb{P}^N$, with the abstraction map $\alpha(X)(n) \triangleq \pi(\{\rho \mid (n, \rho) \in X\})$ for $X \in \mathcal{P}(\Sigma) = \mathcal{P}(N \times \text{Env})$. Given $\sigma^\# \in A$, we define $\text{supp}(\sigma^\#) \triangleq \{n \in N \mid \sigma^\#(n) \neq \perp\}$. ◀

ACTL. ACTL is the fragment of CTL whose temporal formulae are universally quantified over all paths leaving the current state. Thus, given a set of atomic propositions $p \in \mathbf{P}$:

$$\text{ACTL} \ni \varphi ::= p \mid \neg p \mid \varphi_1 \vee \varphi_2 \mid \varphi_1 \wedge \varphi_2 \mid \text{AX } \varphi_1 \mid \text{AF } \varphi_1 \mid \text{AG } \varphi_1 \mid \varphi_1 \text{ AU } \varphi_2$$

► **Definition 2.4** (ACTL semantics). Given a transition system $T = (\Sigma, \mathbf{I}, \mathbf{P}, \rightarrow, \vdash)$, the semantics $\llbracket \varphi \rrbracket \subseteq \Sigma$ of ACTL formulae over T is as follows:

$$\begin{aligned} \llbracket p \rrbracket &\triangleq \{\sigma \in \Sigma \mid \sigma \vdash p\} & \llbracket \neg p \rrbracket &\triangleq \{\sigma \in \Sigma \mid \sigma \not\vdash p\} \\ \llbracket \varphi_1 \vee \varphi_2 \rrbracket &\triangleq \llbracket \varphi_1 \rrbracket \cup \llbracket \varphi_2 \rrbracket & \llbracket \varphi_1 \wedge \varphi_2 \rrbracket &\triangleq \llbracket \varphi_1 \rrbracket \cap \llbracket \varphi_2 \rrbracket \\ \llbracket \text{AX } \varphi_1 \rrbracket &\triangleq \{\sigma \in \Sigma \mid \forall \sigma'. \sigma \rightarrow \sigma' \Rightarrow \sigma' \in \llbracket \varphi_1 \rrbracket\} \\ \llbracket \text{AF } \varphi_1 \rrbracket &\triangleq \{\sigma_0 \in \Sigma \mid \text{for all path } (\sigma_i)_{i \in \mathbb{N}} \exists k \in \mathbb{N}. \sigma_k \in \llbracket \varphi_1 \rrbracket\} \\ \llbracket \text{AG } \varphi_1 \rrbracket &\triangleq \{\sigma_0 \in \Sigma \mid \text{for all path } (\sigma_i)_{i \in \mathbb{N}} \forall j \in \mathbb{N}. \sigma_j \in \llbracket \varphi_1 \rrbracket\} \\ \llbracket \varphi_1 \text{ AU } \varphi_2 \rrbracket &\triangleq \{\sigma_0 \in \Sigma \mid \text{for all path } (\sigma_i)_{i \in \mathbb{N}} \exists k \in \mathbb{N}. (\sigma_k \in \llbracket \varphi_2 \rrbracket \wedge \forall j < k. \sigma_j \in \llbracket \varphi_1 \rrbracket)\} \end{aligned}$$

Universal fragment of single variable μ -calculus. The modal μ -calculus is a well known extension of propositional modal logic with least and greatest fixed point operators. We will focus on its universal fragment only allowing the \Box modal operator that quantifies over all transitions. Moreover, for the sake of simplicity, we restrict to the single variable fragment where, roughly speaking, nested fixpoints cannot have mutual dependencies.

Given a set of atomic propositions $p \in \mathbf{P}$, the μ_{\Box} -calculus is defined as follows:

$$\mu_{\Box} \ni \varphi ::= p \mid \neg p \mid \varphi_1 \vee \varphi_2 \mid \varphi_1 \wedge \varphi_2 \mid \Box \varphi_1 \mid x \mid \mu x. \varphi_x \mid \nu x. \varphi_x$$

► **Definition 2.5** (μ_\square -calculus semantics). *Given $T = (\Sigma, \mathbf{I}, \mathbf{P}, \rightarrow, \vdash)$, and a valuation $\mathcal{V} : \text{Var} \rightarrow \mathcal{P}(\Sigma)$, the semantics $\llbracket \varphi \rrbracket_{\mathcal{V}} \subseteq \Sigma$ of μ_\square -calculus formulae over T is as follows:*

$$\begin{aligned} \llbracket p \rrbracket_{\mathcal{V}} &\triangleq \{\sigma \in \Sigma \mid \sigma \vdash p\} & \llbracket \neg p \rrbracket_{\mathcal{V}} &\triangleq \{\sigma \in \Sigma \mid \sigma \not\vdash p\} \\ \llbracket \varphi_1 \vee \varphi_2 \rrbracket_{\mathcal{V}} &\triangleq \llbracket \varphi_1 \rrbracket_{\mathcal{V}} \cup \llbracket \varphi_2 \rrbracket_{\mathcal{V}} & \llbracket \varphi_1 \wedge \varphi_2 \rrbracket_{\mathcal{V}} &\triangleq \llbracket \varphi_1 \rrbracket_{\mathcal{V}} \cap \llbracket \varphi_2 \rrbracket_{\mathcal{V}} \\ \llbracket \Box \varphi_1 \rrbracket_{\mathcal{V}} &\triangleq \{\sigma \in \Sigma \mid \forall \sigma'. \sigma \rightarrow \sigma' \Rightarrow \sigma' \in \llbracket \varphi_1 \rrbracket_{\mathcal{V}}\} & \llbracket x \rrbracket_{\mathcal{V}} &\triangleq \mathcal{V}(x) \\ \llbracket \mu x. \varphi_x \rrbracket_{\mathcal{V}} &\triangleq \text{lfp}(\lambda S. \llbracket \varphi_x \rrbracket_{\mathcal{V}[x \mapsto S]}) & \llbracket \nu x. \varphi_x \rrbracket_{\mathcal{V}} &\triangleq \text{gfp}(\lambda S. \llbracket \varphi_x \rrbracket_{\mathcal{V}[x \mapsto S]}) \end{aligned}$$

where $\mathcal{V}[x \mapsto S]$ is the usual notation for function update.

We write $\llbracket \varphi \rrbracket$ instead of $\llbracket \varphi \rrbracket_{\mathcal{V}}$ when the valuation is inessential, and $\sigma \models \varphi$ when $\sigma \in \llbracket \varphi \rrbracket$.

3 The Language MOKA

We define a meta-language, called MOKA (for Model checking as abstract interpretation of Kleene Algebras), as a (generalised) Kleene Algebra with a set of basic expressions suited for identifying counterexamples to the validity of temporal formulae.

KAF. We rely on Kozen's Kleene Algebra with tests [34] with a Fixpoint operator [36], KAF for short. Given a set Exp of basic expressions e , KAF is defined below:

$$\text{KAF} \ni r ::= 1 \mid 0 \mid e \mid r_1; r_2 \mid r_1 \oplus r_2 \mid r_1^* \mid X \mid \mu X. r_1$$

The term 1 represents the identity, i.e. no action, 0 represents divergence, $r_1; r_2$ represents sequential composition, $r_1 \oplus r_2$ represents non-deterministic choice, r_1^* represents the Kleene iteration of r_1 , i.e. r_1 performed zero or any finite number of times, X is a variable ranging in a set Var , and $\mu X. r_1$ represents the least fixpoint operator with respect to variable X .

Commands are interpreted as functions over a complete lattice C . Given a semantics $\llbracket \cdot \rrbracket : \text{Exp} \rightarrow C \rightarrow C$ for basic expressions, the semantics of regular expressions $\llbracket \cdot \rrbracket_{\eta} : \text{KAF} \rightarrow C \rightarrow C$ is inductively defined as follows, where $\eta : \text{Var} \rightarrow C \rightarrow C$ is an environment:

$$\begin{aligned} \llbracket 1 \rrbracket_{\eta} &\triangleq \lambda x. x & \llbracket 0 \rrbracket_{\eta} &\triangleq \lambda x. \perp \\ \llbracket e \rrbracket_{\eta} &\triangleq \langle e \rangle & \llbracket r_1; r_2 \rrbracket_{\eta} &\triangleq \llbracket r_2 \rrbracket_{\eta} \circ \llbracket r_1 \rrbracket_{\eta} \\ \llbracket r_1 \oplus r_2 \rrbracket_{\eta} &\triangleq \llbracket r_1 \rrbracket_{\eta} \vee \llbracket r_2 \rrbracket_{\eta} & \llbracket r_1^* \rrbracket_{\eta} &\triangleq \bigvee \{ \llbracket r_1 \rrbracket_{\eta}^k \mid k \in \mathbb{N} \} \\ \llbracket X \rrbracket_{\eta} &\triangleq \eta(X) & \llbracket \mu X. r_1 \rrbracket_{\eta} &\triangleq \text{lfp}(\lambda f : C \rightarrow C. \llbracket r_1 \rrbracket_{\eta[X \mapsto f]}) \end{aligned}$$

It can be seen that the Kleene star r^* can be encoded as a fixpoint $\mu X. (1 \oplus r; X)$, and, when the semantics of basic expressions $\langle \cdot \rangle$ is additive, also as $r^* = \mu X. (1 \oplus X; r)$. Although redundant, the inclusion of the Kleene star in our syntax is very convenient, as it enables significant simplifications in some encodings, particularly when the full expressiveness of least fixpoint calculations is not necessary (see the remark at the end of Section 4). The same applies to the proof logic studied in Section 6: while basic LCL suffices for the KAF fragment without least fixpoint operator, its extension μLCL is otherwise needed. For closed KAF terms the environment is inessential and we write just $\llbracket r \rrbracket$ instead of $\llbracket r \rrbracket_{\eta}$.

The Language MOKA. Given a transition system $T = (\Sigma, \mathbf{I}, \mathbf{P}, \rightarrow, \vdash)$, MOKA is an instance of KAF interpreted over stacks of frames, each frame representing a computation path.

► **Definition 3.1** (Frame, stack). *A frame is a pair $\langle \sigma, \Delta \rangle \in \Sigma \times \mathcal{P}(\Sigma)$. We denote by F_{Σ} the set of frames. A stack is a finite non-empty sequence of frames, i.e. an element of F_{Σ}^+ , denoted by $\langle \sigma, \Delta \rangle :: S$, where S is a stack or the empty sequence ε .*

A frame $\langle \sigma, \Delta \rangle$ represents a computation in T where σ is the current state and Δ is the set of traversed states, used for loop-detection. The order of the traversed states and their possible multiple occurrences are abstracted away as they are irrelevant when checking the satisfaction of a formula. Frames are stacked to deal with formulae with nested operators.

The language MOKA is defined as instance of KAF with basic expressions for extending and filtering frames, for (de)constructing stacks, and to keep track of fixed-point equations.

► **Definition 3.2** (MOKA language). *The language MOKA is an instance of KAF with the following basic expressions, where p ranges over atomic propositions:*

$$e ::= p? \mid \neg p? \mid \text{loop?} \mid \text{next} \mid \text{add} \mid \text{reset} \mid \text{push} \mid \text{pop}$$

The (concrete) semantics of MOKA is given over the powerset of the set of stacks, $C = \mathcal{P}(F_\Sigma^+)$, ordered by subset inclusion. Although we could focus on stacks of uniform length, we use a larger domain to simplify the notation.

The semantics of MOKA commands $\llbracket r \rrbracket_\eta : \mathcal{P}(F_\Sigma^+) \rightarrow \mathcal{P}(F_\Sigma^+)$ follows from the general definition for KAF, once we specify the semantics of its basic expressions.

► **Definition 3.3** (Basic expression semantics). *Given $T = (\Sigma, I, \mathbf{P}, \rightarrow, \vdash)$, the semantics of MOKA basic expressions is the additive extension of the functions below, where $\langle \sigma, \Delta \rangle :: S \in F_\Sigma^+$:*

$$\begin{aligned} \langle p? \rangle \langle \sigma, \Delta \rangle :: S &\triangleq \{ \langle \sigma, \Delta \rangle :: S \mid \sigma \vdash p \} & \langle \neg p? \rangle \langle \sigma, \Delta \rangle :: S &\triangleq \{ \langle \sigma, \Delta \rangle :: S \mid \sigma \not\vdash p \} \\ \langle \text{loop?} \rangle \langle \sigma, \Delta \rangle :: S &\triangleq \{ \langle \sigma, \Delta \rangle :: S \mid \sigma \in \Delta \} & \langle \text{next} \rangle \langle \sigma, \Delta \rangle :: S &\triangleq \{ \langle \sigma', \Delta \rangle :: S \mid \sigma \rightarrow \sigma' \} \\ \langle \text{add} \rangle \langle \sigma, \Delta \rangle :: S &\triangleq \{ \langle \sigma, \Delta \cup \{ \sigma \} \rangle :: S \} & \langle \text{reset} \rangle \langle \sigma, \Delta \rangle :: S &\triangleq \{ \langle \sigma, \emptyset \rangle :: S \} \\ \langle \text{push} \rangle \langle \sigma, \Delta \rangle :: S &\triangleq \{ \langle \sigma, \Delta \rangle :: \langle \sigma, \Delta \rangle :: S \} & \langle \text{pop} \rangle \langle \sigma, \Delta \rangle :: S &\triangleq \{ S \mid S \neq \varepsilon \} \end{aligned}$$

The basic expressions $p?$, $\neg p?$, loop? , next , add , reset operate on the top frame. The filter $p?$ (resp. $\neg p?$) checks the validity of the proposition p (resp. $\neg p$), loop? checks if the current state loops back to some state in the trace, next extends the trace by one step in all possible ways, add adds the current state to the trace and reset empties the trace. Instead, push and pop change the stack length: push extends the stack to start a new trace, while pop restores the previous trace from the stack. Given a set of states $X \subseteq \Sigma$, we write $\llbracket r \rrbracket_\eta X$ as a shorthand for $\llbracket r \rrbracket_\eta(\{ \langle \sigma, \emptyset \rangle \mid \sigma \in X \})$.

4 Formulae as MOKA programs

We show that ACTL and μ_\square -calculus formulae φ can be transformed into MOKA programs $\llbracket \varphi \rrbracket$ so that for any transition system T , the semantics of $\llbracket \varphi \rrbracket$ over T consists exactly of the counterexamples to the validity of φ , that is, the states which do not satisfy φ .

► **Theorem 4.1** (Model checking as program verification). *Given $T = (\Sigma, I, \mathbf{P}, \rightarrow, \vdash)$, for any ACTL or μ_\square formula φ and set of stacks $P \subseteq F_\Sigma^+$, it holds $\llbracket \llbracket \varphi \rrbracket \rrbracket_\eta P = \{ \langle \sigma, \Delta \rangle :: S \in P \mid \sigma \not\models \varphi \}$.*

It follows that $\llbracket \llbracket \varphi \rrbracket \rrbracket_\eta P \subseteq P$ and that $\llbracket \llbracket \varphi \rrbracket \rrbracket_\eta \{ \sigma \} = \emptyset$ iff $\sigma \models \varphi$. It is also worth noting that the semantics of MOKA programs encoding formulae is a lower closure on sets of states, being monotone, reductive and idempotent (and it preserves arbitrary unions) [39, Section 3.2.3]. In this respect, it behaves analogously to the collecting semantics of Boolean tests which filter out memory states, keeping only those that satisfy the condition. As a consequence, similarly to what happens when abstracting Boolean tests, in any abstract domain A approximating sets of states the identity $\lambda a \in A. a$, is a correct (over-)approximation of $\llbracket \llbracket \varphi \rrbracket \rrbracket_\eta$.

ACTL as MOKA. To each ACTL formula φ we assign the MOKA program $[\overline{\varphi}]$ as follows:

$$\begin{aligned}
[\overline{p}] &\triangleq \neg p? & [\overline{\neg p}] &\triangleq p? \\
[\overline{\varphi_1 \wedge \varphi_2}] &\triangleq [\overline{\varphi_1}] \oplus [\overline{\varphi_2}] & [\overline{\varphi_1 \vee \varphi_2}] &\triangleq [\overline{\varphi_1}]; [\overline{\varphi_2}] \\
[\overline{AX \varphi}] &\triangleq \text{push; next; } [\overline{\varphi}]; \text{pop} \\
[\overline{AF \varphi}] &\triangleq [\overline{\varphi}]; \text{push; reset; (add; next; } [\overline{\varphi}])^*; \text{loop?; pop} \\
[\overline{AG \varphi}] &\triangleq \text{push; next}^*; [\overline{\varphi}]; \text{pop} \\
[\overline{\varphi_1 AU \varphi_2}] &\triangleq [\overline{\varphi_2}]; \text{push; reset; (add; next; } [\overline{\varphi_2}])^*; (\text{loop?} \oplus [\overline{\varphi_1}]); \text{pop}
\end{aligned}$$

The intuition is that MOKA programs $[\overline{\varphi}]$ act as (negative) filters: applied to a set of frames $\langle \sigma, \Delta \rangle :: S$, each representing a computation that reached state σ , they filter out states where φ is satisfied, keeping only candidate counterexamples to the validity of φ . In general, when the check requires exploring the future of the current state, a new frame is started with **push**; if the search is successful, a closing **pop** recovers the starting state violating the formula. We explain the main clauses defining $[\overline{\varphi}]$ (see [1, Theorem B.1] for the proof of the correctness of this transform). It turns out that $\langle \sigma, \Delta \rangle :: S$ is a counterexample for

- **AX** φ if φ is violated in at least one successor state of σ (as computed by **next**; $[\overline{\varphi}]$).
- **AF** φ if there is an infinite trace where φ never holds. This is encoded by saying that φ does not hold in the current state, i.e. $[\overline{\varphi}]$, and after that, there is an infinite trace traversing only states (collected through the command $(\text{add; next; } [\overline{\varphi}])^*$) that do not satisfy φ . Since we deal with finite state systems, infinite traces can be identified with looping traces (whence the check **loop?**); for this to work, after the **push** we **reset** the past.
- **AG** φ if we can reach, by repeatedly applying **next**, a counterexample to φ .
- $\varphi_1 \text{ AU } \varphi_2$ if φ_2 does not hold in the current state, i.e. $[\overline{\varphi_2}]$, and progressing through states that do not satisfy φ_2 with $(\text{add; next; } [\overline{\varphi_2}])^*$, either we detect a maximal (looping) trace or a state where φ_1 does not hold.

► **Example 4.2** (Control flow graphs and ACTL encoding). In our running example (see Examples 2.2 and 2.3), we check the property stating that if the program **c** terminates then the variable z is zero, which in the CFG means “when the exit node e is reached, $z = 0$ ” holds. This is expressed by the ACTL formula $\varphi = \text{AG } (n = e \rightarrow z = 0)$, where we use $p \rightarrow \varphi'$ as syntactic sugar for $\neg p \vee \varphi'$. The corresponding MOKA program is: $m_\varphi \triangleq [\overline{\varphi}] = [\overline{\text{AG } (n = e \rightarrow z = 0)}] = \text{push; next}^*; n = e?; z \neq 0?; \text{pop}$. ▮

μ_\square -calculus as MOKA. To each μ_\square -formula φ we assign the MOKA program $[\overline{\varphi}]$ as follows:

$$\begin{aligned}
[\overline{p}] &\triangleq \neg p? & [\overline{\neg p}] &\triangleq p? \\
[\overline{\varphi_1 \wedge \varphi_2}] &\triangleq [\overline{\varphi_1}] \oplus [\overline{\varphi_2}] & [\overline{\varphi_1 \vee \varphi_2}] &\triangleq [\overline{\varphi_1}]; [\overline{\varphi_2}] \\
[\overline{\Box \varphi}] &\triangleq \text{push; next; } [\overline{\varphi}]; \text{pop} & [\overline{X}] &\triangleq X \\
[\overline{\mu x. \varphi_x}] &\triangleq \text{push; reset; } \mu X. (\text{loop?} \oplus (\text{add; } [\overline{\varphi_x}])) ; \text{pop} & [\overline{\nu x. \varphi_x}] &\triangleq \mu X. [\overline{\varphi_x}]
\end{aligned}$$

The second component of a frame $\langle \sigma, \Delta \rangle$ is still used to identify looping computations. This intervenes in the encoding of least fixpoints $\mu x. \varphi_x$: when checking the formula, the current state first logged to the current frame (**add**; $[\overline{\varphi_x}]$ branch); a counterexample is found when we try to verify the fixpoint property in a state where the check has already been tried (filtered by **loop?**). The encoding of greatest fixpoints $\nu x. \varphi_x$ instead is simpler: searching for counterexamples naturally translates to a least fixpoint, which is offered natively by MOKA.

The correctness of this transform (see [1, Theorem B.2 and Corollary B.3]) leverages a small variation of the tableau construction in [53], with judgements roughly of the form $\sigma, \Delta \vdash \varphi$, meaning that the formula φ holds in a state σ assuming that all the states in Δ

have been visited while checking the current fixpoint subformula. Then, it can be shown that given a formula φ and a stack $\langle \sigma, \Delta \rangle :: S$, if there is no successful tableau for $\sigma, \Delta \vdash \varphi$ then $\llbracket \llbracket \varphi \rrbracket \rrbracket_\eta \{ \langle \sigma, \Delta \rangle :: S \} = \{ \langle \sigma, \Delta \rangle :: S \}$ holds, while $\llbracket \llbracket \varphi \rrbracket \rrbracket_\eta \{ \langle \sigma, \Delta \rangle :: S \} = \emptyset$ otherwise.

► **Example 4.3** (Control flow graphs and μ_\square -calculus encoding). Consider again Examples 2.2 and 2.3. We now check the property “if after three steps we end up in program point 3 then we will loop forever, every 4 steps, on this program point 3.” It is known that regular properties of this kind cannot be expressed in ACTL (see, e.g., [54]). This property can be expressed in the μ_\square -calculus as $\psi \triangleq \square^3(n = 3 \rightarrow \nu x.(n = 3 \wedge \square^4 x))$, where \square^k is a shorthand for $\square \dots \square$ repeated k times. Letting r^k as a shorthand for a k times composition $r; \dots; r$ of a MOKA program r , the MOKA program encoding ψ is $\llbracket \psi \rrbracket = (\text{push}; \text{next})^3; (n = 3; \llbracket \nu x.(n = 3 \wedge \square^4 x) \rrbracket); \text{pop}^3$, with $\llbracket \nu x.(n = 3 \wedge \square^4 x) \rrbracket = \mu X. ([n \neq 3? \oplus ((\text{push}; \text{next})^4; X; \text{pop}^4)])$. \dashv

It is well known that all ACTL formulae can be expressed as μ_\square -calculus formulae. For example, $\text{AX } \varphi = \square \varphi$, $\text{AF } \varphi = \mu x. (\varphi \vee \square x)$, and $\text{AG } \varphi = \nu x. (\varphi \wedge \square x)$. Hence one could obtain programs generating counterexamples for ACTL by encoding ACTL in the μ_\square -calculus and then generating the corresponding program. However, this in general produces programs which are (unnecessarily) more complex than those produced for ACTL formulae. For instance, the program for $\text{AF } \varphi$ obtained through the μ_\square -calculus encoding above would be $\text{push}; \text{reset}; \mu X. (\text{loop}? \oplus (\text{add}; \llbracket \varphi \rrbracket; \text{push}; \text{next}; X; \text{pop})); \text{pop}$.

5 Abstract Interpretation of MOKA

We show how to lift an abstraction over the states of a transition system to an abstraction over the domain of stacks. This is achieved stepwise by first considering an abstraction applied to each single stack, and then by merging classes of stacks through a suitable equivalence. The resulting stack abstraction will induce the abstract interpretation of MOKA programs.

Lifting the abstraction. Let (Σ, I, \rightarrow) be a fixed transition system and let $\langle \alpha, \gamma \rangle : \mathcal{P}(\Sigma) \rightleftarrows A$ be an abstraction of state properties. We consider the lattice $F_A \triangleq A \times A$, with componentwise order, whose elements $\langle \sigma^\#, \delta^\# \rangle$ are called *abstract frames*, and, in turn, F_A^+ whose elements $\langle \sigma^\#, \delta^\# \rangle :: S^\#$ are called *abstract stacks*. As in the concrete case, we abbreviate $\langle \sigma^\#, \perp_A \rangle$ as $\sigma^\#$.

The abstraction map α on state properties can be extended to a frame abstraction, that by abusing the notation, we still denote by α , and is defined by $\alpha(\langle \sigma, \Delta \rangle) \triangleq \langle \alpha(\{\sigma\}), \alpha(\Delta) \rangle$. In turn, the abstraction is inductively defined on stacks as $\alpha(\langle \sigma, \Delta \rangle :: S^n) \triangleq \alpha(\langle \sigma, \Delta \rangle) :: \alpha(S^n)$.

Now, a set of stacks is abstracted to a set of abstract stacks by first applying α pointwise using the image adjunction (see Example 2.1) and then joining classes of abstract stacks in a way which is parameterised by a suitable equivalence. Given an equivalence $\sim \subseteq L \times L$, we let $[x]_\sim$ denote the equivalence class of $x \in L$ w.r.t. \sim . Given $x \in L$ we let $\downarrow x \triangleq \{y \in L \mid y \leq x\}$.

► **Definition 5.1** (Equivalence adjunction). *Given a complete lattice L , a compatible equivalence is an equivalence $\sim \subseteq L \times L$ such that for all $x \in L$, it holds that $[x]_\sim$ is closed by joins of non-empty subsets. Let $\mathcal{P}(L)_\sim \triangleq \{X \in \mathcal{P}(L) \mid \forall x \in X. [x]_\sim \cap X = \{x\}\}$, ordered as follows: $X \leq_\sim Y$ if for all $x \in X$ there is $y \in Y$ such that $x \sim y$ and $x \leq y$. Then, the pair $\langle \alpha_\sim, \gamma_\sim \rangle : \mathcal{P}(L) \rightleftarrows \mathcal{P}(L)_\sim$ defined, for $X \in \mathcal{P}(L)$, $Y \in \mathcal{P}(L)_\sim$, by*

$$\alpha_\sim(X) \triangleq \left\{ \bigvee (X \cap [x]_\sim) \mid x \in X \right\} \quad \gamma_\sim(Y) \triangleq \bigcup \{[y]_\sim \cap \downarrow y \mid y \in Y\}$$

is a Galois connection, called equivalence adjunction.

In words, $\mathcal{P}(L)_\sim \subseteq \mathcal{P}(L)$ consists of the subsets of L containing at most one representative for each equivalence class, and $\alpha_\sim(X)$ abstracts X in the subset (in $\mathcal{P}(L)_\sim$) consisting of the least upper bounds of \sim -equivalent elements in X .

We can now define the stack abstraction parameterised by an equivalence on the lattice of abstract frames. Given an equivalence $\sim \subseteq X \times X$ on any set X , \sim_+ denotes the equivalence on X^+ induced by \sim as follows: for all $x_1 \cdots x_n, y_1 \cdots y_m \in A^+$, $x_1 \cdots x_n \sim_+ y_1 \cdots y_m \in X^+$ if $n = m$ and $x_i \sim y_i$ for all $i \in \{1, \dots, n\}$.

► **Definition 5.2** (Stack abstraction). *Let \sim be a compatible equivalence on the abstract domain A and let us extend it to the lattice of abstract frames F_A by $\langle \sigma_1^\#, \delta_1^\# \rangle \sim \langle \sigma_2^\#, \delta_2^\# \rangle$ if $\sigma_1^\# \sim \sigma_2^\#$. The \sim -stack abstraction is the domain $A_\sim^s \triangleq \mathcal{P}(F_A^+)_\sim$ along with the Galois connection $\langle \alpha_\sim^s, \gamma_\sim^s \rangle : \mathcal{P}(F_\Sigma^+) \rightleftarrows A_\sim^s$ defined by $\alpha_\sim^s \triangleq \alpha_{\sim_+} \circ \alpha^>$ and $\gamma_\sim^s \triangleq \alpha^< \circ \gamma_{\sim_+}$.*

In words, given a set X of stacks, its abstraction $\alpha_\sim^s(X)$ first applies pointwise the underlying α to each stack in X , and then joins equivalent abstract stacks. As corner cases we can have the identity relation $\sim = \text{id}$, which joins abstract frames with identical first component, and $\sim = A \times A$, the trivial relation, which joins all abstract stacks into one.

An abstract interpreter for counterexamples. Given an abstraction for state properties $\langle \alpha, \gamma \rangle : \mathcal{P}(\Sigma) \rightleftarrows A$ and a compatible equivalence on A , following the standard approach in abstract interpretation [16], we consider an abstract semantics $\llbracket r \rrbracket_\eta^\# : A_\sim^s \rightarrow A_\sim^s$, defined inductively as explained in § 3 and using as abstract semantics for basic expressions e their BCAs on A_\sim^s , denoted by $\langle e \rangle^{A_\sim^s}$. We write $\llbracket r \rrbracket^\#$ instead of $\llbracket r \rrbracket_\eta^\#$ when the abstract environment is inessential. Notably, the BCAs of basic expressions independent from the underlying system can be effectively defined, and some of them result to be complete.

► **Theorem 5.3** (Basic abstract operations). *Let \sim be a compatible equivalence on the abstract domain A . The BCAs of the basic expressions **add**, **reset**, **push**, **pop** for the \sim -stack abstraction are as follows: for all $\langle \sigma^\#, \delta^\# \rangle :: S^\# \in F_A^+$*

$$\begin{aligned} \langle \text{add} \rangle^{A_\sim^s} \{ \langle \sigma^\#, \delta^\# \rangle :: S^\# \} &= \{ \langle \sigma^\#, \delta^\# \vee \sigma^\# \rangle :: S^\# \} & \langle \text{reset} \rangle^{A_\sim^s} \{ \langle \sigma^\#, \delta^\# \rangle :: S^\# \} &= \{ \langle \sigma^\#, \perp \rangle :: S^\# \} \\ \langle \text{push} \rangle^{A_\sim^s} \{ \langle \sigma^\#, \delta^\# \rangle :: S^\# \} &= \{ \langle \sigma^\#, \delta^\# \rangle :: \langle \sigma^\#, \delta^\# \rangle :: S^\# \} & \langle \text{pop} \rangle^{A_\sim^s} \{ \langle \sigma^\#, \delta^\# \rangle :: S^\# \} &= \{ S^\# \} \end{aligned}$$

and the operations **reset**, **push**, **pop** are globally complete. Moreover

$$\begin{aligned} \langle \text{p?} \rangle^{A_\sim^s} \{ \langle \sigma^\#, \delta^\# \rangle :: S^\# \} &\leq_\sim \{ \langle \sigma^\# \wedge \text{p?}^{A_\sim^s}, \delta^\# \rangle :: S^\# \} \\ \langle \neg \text{p?} \rangle^{A_\sim^s} \{ \langle \sigma^\#, \delta^\# \rangle :: S^\# \} &\leq_\sim \{ \langle \sigma^\# \wedge \neg \text{p?}^{A_\sim^s}, \delta^\# \rangle :: S^\# \} \\ \langle \text{loop?} \rangle^{A_\sim^s} \{ \langle \sigma^\#, \delta^\# \rangle :: S^\# \} &\leq_\sim \{ \langle \sigma^\#, \delta^\# \rangle :: S^\# \mid \sigma^\# \wedge \delta^\# \neq \perp \} \end{aligned}$$

where $\text{p?}^{A_\sim^s} = \alpha(\llbracket \text{p?} \rrbracket \Sigma)$ is the abstraction of the set of concrete states satisfying p (and similarly for $\neg \text{p?}^{A_\sim^s}$).

For **loop?** one can get $\langle \text{loop?} \rangle^{A_\sim^s} \{ \langle \sigma^\#, \delta^\# \rangle :: S^\# \} = \{ \langle \sigma^\# \wedge \delta^\#, \delta^\# \rangle :: S^\# \mid \sigma^\# \wedge \delta^\# \neq \perp \}$ under additional conditions on \sim (cf. [1, Definition C.6]).

The soundness-by-design of this abstract semantics entails a sound program verification.

► **Proposition 5.4** (Program verification for satisfaction). *Given a transition system $T = (\Sigma, I, \mathbf{P}, \rightarrow, \vdash)$, for all ACTL or μ_\square -calculus formulae φ , abstract domain A , and a compatible equivalence \sim on A , if $\llbracket \llbracket \varphi \rrbracket \rrbracket_\eta^\#(\alpha_\sim^s(I)) = \perp$ then, for all $\sigma \in I$, we have $\sigma \models \varphi$.*

As in the concrete case, the abstract semantics of programs encoding formulae is a lower closure. This, together with the observation that the semantics only depends on the top frame of a stack, is relevant for the concrete implementation.

► **Remark 5.5.** We recalled in § 4 that for any Boolean test b and abstract domain A , the identity function $\lambda a \in A. a$ is a correct approximation of the filtering semantics $\llbracket b \rrbracket$ of b . This function can be enhanced to $\lambda a \in A. \alpha(\llbracket b \rrbracket \Sigma) \wedge_A a$, which is also a correct approximation of $\llbracket b \rrbracket$ [39, Section 4.1]. If A is a partitioning abstraction then one can show that $\lambda a \in A. \alpha(\llbracket b \rrbracket \Sigma) \wedge_A a$ is indeed the *best* correct approximation of $\llbracket b \rrbracket$. However, this property does not hold in general, even assuming that A is a disjunctive abstraction, i.e., the additivity of γ (this is substantiated in [1, Example C.10]). ◻

► **Example 5.6** (Control flow graphs and abstract frames). In our running example, abstract frames are of the shape $A \times A = \mathbb{P}^N \times \mathbb{P}^N$. The abstraction of singletons is $\alpha(\{(n, \rho)\}) = (n \mapsto a)$ with $a(n') = \perp$ for all $n' \neq n$. With the aim of joining past states when they correspond to the same program point, we consider the equivalence $\sim \subseteq \mathbb{P}^N \times \mathbb{P}^N$ on abstract frames defined by $\sigma_1^\# \sim \sigma_2^\#$ when $\text{supp}(\sigma_1^\#) = \text{supp}(\sigma_2^\#)$. The property $\psi = \Box^3(n = 3 \rightarrow \nu x. (n = 3 \wedge \Box^4 x))$ from Example 4.3 holds in the system, and we can prove it with this abstraction. Let $\sigma^\# = (s \mapsto \top) = \alpha(\{(s, xyzw)\})$ we can compute $\llbracket \llbracket \bar{\psi} \rrbracket \rrbracket_\eta^\#(\sigma^\#) = \perp$, implying that the formula holds from any initial states (convergence is after a single full iteration). Instead, the abstract computation for $\varphi = \text{AG}(n = e \rightarrow z = 0)$ from Example 4.2 yields a false positive. ◻

6 Locally Complete Analyses

If the abstract interpretation of a MOKA program returns an alarm, i.e., $\llbracket \llbracket \bar{\psi} \rrbracket \rrbracket^\# \alpha_s^s(I) \neq \perp$, then any initial state in the concretisation of $\llbracket \llbracket \bar{\psi} \rrbracket \rrbracket^\# \alpha_s^s(I)$ is a candidate counterexample to the validity of ψ . However, due to over-approximation, we cannot distinguish spurious counterexamples from true ones. Here, we discuss how to combine under- and over-approximation for the analysis of MOKA programs $\llbracket \bar{\psi} \rrbracket$ to overcome this problem. In particular we leverage Local Completeness Logic (LCL) [7, 9] possibly paired with Abstract Interpretation Repair (AIR) strategies [8], to improve the analysis precision.

LCL. Most abstract domains are not globally complete for program analysis, so that the corresponding analyses may well yield false alarms. Accordingly, [7, 9] studies how completeness can be locally weakened to an analysis of interest: given a function $f : C \rightarrow C$, an abstract domain $A_{\alpha, \gamma} \in \text{Abs}(C)$ is *locally complete on a value* $c \in C$, denoted by $\mathbb{C}_c^A(f)$, when $\alpha \circ f(c) = \alpha \circ f \circ \gamma \circ \alpha(c)$ (hence global completeness amounts to $\mathbb{C}_c^A(f)$ for all $c \in C$). Intuitively, the absence of false alarms in an abstract computation comes as a consequence of the local completeness of the abstract transfer functions on the traversed concrete states.

Moreover, local completeness is a convex property, and this allows to check local completeness on suitable under-approximations $u \leq c$ such that $\alpha(u) = \alpha(c)$, as $\mathbb{C}_u^A(f)$ implies $\mathbb{C}_c^A(f)$. The work [7, 9] implements this idea through LCL, an under-approximating program logic (in the style of incorrectness logic [43]), parameterised by the abstract domain A which provides an over-approximation. The LCL proof rules for KAT programs are recalled in Table 1. A provable LCL triple $\vdash_A [P] \text{ r } [Q]$ ensures that each state satisfying Q is reachable from some state satisfying P , and also guarantees that Q and $\llbracket r \rrbracket P$ have the same abstraction in A . This means that the LCL program logic is sound w.r.t. the following notion of validity.

► **Definition 6.1** (Valid LCL triples). *Let $P, Q \in C$ and r be a KAT program. A program triple $\vdash_A [P] \text{ r } [Q]$ is valid if $\mathbb{C}_P^A(r) \wedge Q \leq \llbracket r \rrbracket P \wedge \alpha(\llbracket r \rrbracket P) = \alpha(Q)$.*

The condition $\mathbb{C}_P^A(r)$ states that A is locally complete for $\llbracket r \rrbracket$ on P , while $Q \leq \llbracket r \rrbracket P$ that Q under-approximates $\llbracket r \rrbracket P$ and $\alpha(\llbracket r \rrbracket P) = \alpha(Q)$ that $\gamma(\alpha(Q))$ over-approximates $\llbracket r \rrbracket P$.

■ **Table 1** Rules of the Local Completeness Logic [9].

$$\begin{array}{c}
\frac{\mathbb{C}_P^A(e)}{\vdash_A [P] e \llbracket e \rrbracket [P]} \text{ (transfer)} \qquad \frac{P' \leq P \leq \gamma(\alpha(P')) \quad Q \leq Q' \leq \gamma(\alpha(Q))}{\vdash_A [P'] r [Q']} \text{ (relax)} \\
\\
\frac{\vdash_A [P] r_1 [W] \quad \vdash_A [W] r_2 [Q]}{\vdash_A [P] r_1; r_2 [Q]} \text{ (seq)} \qquad \frac{\vdash_A [P] r_1 [Q_1] \quad \vdash_A [P] r_2 [Q_2]}{\vdash_A [P] r_1 \oplus r_2 [Q_1 \vee Q_2]} \text{ (join)} \\
\\
\frac{\vdash_A [P] r [W] \quad \vdash_A [P \vee W] r^* [Q]}{\vdash_A [P] r^* [Q]} \text{ (rec)} \qquad \frac{\vdash_A [P] r [Q] \quad Q \leq \gamma(\alpha(P))}{\vdash_A [P] r^* [P \vee Q]} \text{ (iterate)}
\end{array}$$

■ **Table 2** Novel rules for μLCL .

$$\frac{}{\vdash_A [P] \mu^0 X.r [\perp]} (\mu^0) \qquad \frac{\vdash_A [P] r[\mu^n X.r/X] [Q]}{\vdash_A [P] \mu^{n+1} X.r [Q]} (\mu^+) \qquad \frac{\llbracket \mu X.r \rrbracket^\sharp \alpha(P) \leq_A \alpha(Q)}{\vdash_A [P] \mu X.r [Q]} \text{ (fix)}$$

LCL for MOKA. The LCL proof system has been defined for KAT programs only, hence it is enough for MOKA programs induced by ACTL formulae. For analysing general KAF and thus MOKA programs, we extend it to μLCL including the inference rules in Table 2, where $r[s/X]$ denotes the capture-avoiding substitution of the free occurrences of X for s in r and the term $\mu^n X.r$ represents the n -th fixpoint approximant with the expected semantics ($\llbracket \mu^0 X.r \rrbracket_\eta \triangleq \lambda x. \perp$ and $\llbracket \mu^{n+1} X.r \rrbracket_\eta \triangleq \llbracket r \rrbracket_{\eta[X \mapsto \llbracket \mu^n X.r \rrbracket_\eta]}$). The syntax $\mu^n X.r$ is used only in μLCL derivations. Intuitively, since $\mu^0 X.r$ behaves as 0, the unique valid judgement is $\vdash_A [P] \mu^0 X.r [\perp]$. The rule (μ^+) serves just to unfold $\mu^{n+1} X.r$ as many times as needed. The key rule is (fix) which can be used whenever the n -th approximant provides enough information: in case the premise holds, by local completeness we will have $\llbracket \mu^n X.r \rrbracket^\sharp \alpha(P) = \alpha(Q)$. Since it is always the case that $\llbracket \mu^n X.r \rrbracket^\sharp \alpha(P) \leq_A \llbracket \mu X.r \rrbracket^\sharp \alpha(P)$, if the side condition $\llbracket \mu X.r \rrbracket^\sharp \alpha(P) \leq_A \alpha(Q)$ holds, then local completeness for the fixpoint term $\mu X.r$ can be inferred. These novel rules are sound, in the sense of preserving validity (Definition 6.1).

► **Remark 6.2.** When a least fixpoint appears in $\llbracket \varphi \rrbracket$, we can exploit the following heuristic for choosing the value of n : if the abstract domain satisfies the Ascending Chain Condition (ACC) then, for every $P \in C$ there is $n_P \in \mathbb{N}$ such that $\llbracket \mu^{n_P} X.r \rrbracket^\sharp \alpha(P) = \llbracket \mu X.r \rrbracket^\sharp \alpha(P)$. Then, by taking any $n \geq n_P$ such that $\vdash_A [P] \mu^n X.r [Q]$ is provable, the condition $\llbracket \mu X.r \rrbracket^\sharp \alpha(P) \leq_A \alpha(Q)$ is readily satisfied, so that $\vdash_A [P] \mu X.r [Q]$ is valid. More precisely, the rule below is sound:

$$\frac{\vdash_A [P] \mu^n X.r [Q] \quad \llbracket \mu^n X.r \rrbracket^\sharp \alpha(P) = \llbracket \mu X.r \rrbracket^\sharp \alpha(P)}{\vdash_A [P] \mu X.r [Q]} \text{ (afix)} \quad \lrcorner$$

The following result relating μLCL proofs with validity of formulae follows as an easy consequence of [9, Corollary 5.6] and Theorem 4.1.

► **Corollary 6.3 (Precision).** *Let $T = (\Sigma, I, \mathbf{P}, \rightarrow, \vdash)$ be a transition system. For all ACTL or μ_\square -calculus formulae φ , abstract domain A and compatible equivalence \sim on A , if the triple $\vdash_{A^\sim} [I] \llbracket \varphi \rrbracket [Q]$ is derivable, then $Q \subseteq I$ and*

- $Q = \emptyset$ if and only if for all $\sigma \in I$ we have $\sigma \models \varphi$;
- if $Q \neq \emptyset$ then for all $\sigma \in Q$ we have $\sigma \not\models \varphi$.

► **Example 6.4** (Control flow graphs and LCL derivations). As pointed out in Example 5.6 the computation of the program m_φ , which encodes in MOKA the property $\varphi = \text{AG } (n = e \rightarrow z = 0)$ (see Example 4.2), yields a false alarm when we use the abstract domain A depicted in Figure 2c. Consider as set of initial states $I_0 = \{\langle (s, 0100), \emptyset \rangle, \langle (s, 0011), \emptyset \rangle\}$, whose abstraction covers the set of possible initial states $\{\langle (s, xyzw), \emptyset \rangle \mid x, y, z, w \in \mathbb{Z}_k\}$, because $\alpha_{\sim}^s(I_0) = (s \mapsto \top)$. Since the property holds, trying to derive a triple $\vdash_{A_{\sim}} [I_0] m_\varphi [\emptyset]$ leads to the failure of some local completeness assumption, which can drive the refinement of A . The imprecision is found in the fourth iteration of the `next` operator, in which the set $\{(3, 1100), (3, 0111)\}$ is abstracted to $(3, xyzw)$, losing the relation between p and q being either both valid or both invalid. Several domain repairs are possible. Following [8] we can repair the domain by adding the abstract point $q \rightarrow p$ (a more abstract repair than adding the element $q \leftrightarrow p$, as proposed in [4]). Here we can also exploit a different route, by refining the equivalence so that $\sigma_1^\# \sim \sigma_2^\#$ when $\text{supp}(\sigma_1^\#) = \text{supp}(\sigma_2^\#) \neq 3$. This intuitively corresponds to abstract separately the states at program point $n = 3$. In both cases, it holds $\llbracket \overline{\varphi} \rrbracket^\# \alpha_{\sim}^s(I_0) = \perp$ in the refined abstract domain A' . \lrcorner

► **Example 6.5** (Traffic light example, LCL derivation and repair, μ -calculus version). Consider the toy traffic light example from Figure 1a and the property $\psi = \text{AG } (g \rightarrow \text{AX } d)$, briefly discussed in the introduction. We give some additional details in the light of the results presented. We first translate the formula in $\mu\Box$ -calculus as $\psi_\mu = \nu x. ((g \rightarrow \Box d) \wedge \Box x)$ whose encoding is $\llbracket \psi_\mu \rrbracket = \mu X. (r_1 \oplus \text{push}; \text{next}; X; \text{pop}) \triangleq \mu X. r_x$, where $r_1 \triangleq g?; \text{push}; \text{next}; \neg d?; \text{pop}$. Then we try to derive the triple $\vdash_{A_{\sim}} [\{rs\}] \mu X. r_x [\emptyset]$ in the abstract domain A (see Figure 1c) by applying the rule (afix) with $n = 2$ (the number of iterations needed for convergence of the abstract fixpoint in A). By definition, $\mu^0 X. r_x = 0$; then the first approximant is $\mu^1 X. r_x = r_1 \oplus \text{push}; \text{next}; 0; \text{pop}$, but we can omit the branch that contains 0, so $\mu^1 X. r_x = r_1$; and the second approximant is $r_2 \triangleq \mu^2 X. r_x = r_1 \oplus \text{push}; \text{next}; r_1; \text{pop}$. The attempt to derive the triple $\vdash_{A_{\sim}} [\{rs\}] r_2 [\emptyset]$ is sketched below, where the mid row gives an indication of the triples labelling the leaves of the derivation, reporting for each basic command involved, the pre-condition and post-condition of the triple, while the top row reports their abstractions.

$$\begin{array}{cccccccccccc} \textcolor{red}{a} & \perp & \perp & \perp & \perp & \perp & \textcolor{red}{a} & \textcolor{red}{a} & \textcolor{red}{a} \vee \textcolor{blue}{c} & \textcolor{blue}{c} & \textcolor{blue}{c} & \\ \textcolor{red}{[rs]} & [\emptyset] & [\emptyset] & [\emptyset] & [\emptyset] & [\emptyset] & \textcolor{red}{[rs]} & \textcolor{red}{[rs]} & \textcolor{red}{[rs]} & \textcolor{red}{[rs]} & \textcolor{red}{[rs]} & \textcolor{red}{[rs]} \\ \vdash_{A_{\sim}} [\textcolor{red}{rs}] & g?; \text{push}; \text{next}; \neg d?; \text{pop} & \oplus & \text{push}; \text{next}; & g?; \text{push}; \text{next}; \neg d?; \text{pop}; \text{pop} & [\emptyset] & & & & & & \end{array}$$

While we can derive $\vdash_{A_{\sim}} [\{rs\}] r_1 [\emptyset]$, the source of local incompleteness is found in the second branch, at the $!$ position, where the proof obligation for $\mathbb{C}_{\{gs\}}^{A_{\sim}}(\text{next})$ fails, since $\alpha_{\sim}^s \circ \llbracket \text{next} \rrbracket (\{gs\}) = \alpha_{\sim}^s(\{gd, yd\}) = b \wedge c$, while $\alpha_{\sim}^s \circ \llbracket \text{next} \rrbracket \circ \gamma_{\sim}^s \circ \alpha_{\sim}^s(\{gs\}) = \alpha_{\sim}^s \circ \llbracket \text{next} \rrbracket \{gs, gd, gs, ys\} = \alpha_{\sim}^s(\{gs, gd, gs, ys, rs\}) = a \vee c$. We can thus repair the abstract domain A . Following the procedure in [8], we add the new abstract point c_1 :

$$c_1 \triangleq \bigvee \{T \subseteq \Sigma \mid T \subseteq \gamma \circ \alpha \{gs\}, \llbracket \text{next} \rrbracket T \subseteq \llbracket \text{next} \rrbracket (\{gs\})\} = \{gs, gd\}$$

The repaired domain is $A_1 \triangleq A \cup \{c_1, b \wedge c_1\}$, where $b \wedge c_1$ is also added because the domain must be closed under meets. Now, in A_1 , the proof obligation for $\mathbb{C}_{\{gs\}}^{A_1}(\text{next})$ holds true: $\llbracket \text{next} \rrbracket (\{gs\}) = \alpha_1^s \circ \llbracket \text{next} \rrbracket \circ \gamma \circ \alpha_1^s(\{gs\}) = b \wedge c$. In fact, in A_1 we just have $\llbracket \overline{\psi} \rrbracket^\# \alpha_1^s(\{rs\}) = \perp$ (we omit the details for the first branch r_1 that are as above):

$$\begin{array}{cccccccccccc} \textcolor{red}{a} & \perp & \textcolor{red}{a} & \textcolor{red}{a} & \textcolor{red}{a} \vee \textcolor{blue}{c} & c_1 & c_1 & b \wedge c & \perp & \perp & \perp & \\ \textcolor{red}{[rs]} & [\emptyset] & \textcolor{red}{[rs]} & \textcolor{red}{[rs]} & \textcolor{red}{[rs]} & \textcolor{red}{[rs]} & \textcolor{red}{[rs]} & \textcolor{red}{[rs]} & \textcolor{red}{[rs]} & \textcolor{red}{[rs]} & \textcolor{red}{[rs]} & \textcolor{red}{[rs]} \\ \vdash_{A_1} [\textcolor{red}{rs}] & r_1 & \oplus & \text{push}; \text{next}; & g?; \text{push}; \text{next}; \neg d?; \text{pop}; \text{pop} & [\emptyset] & & & & & & \end{array}$$

7 Related Work

Abstract interpretation and model checking have been applied to and combined with each other in several ways [2, 4, 5, 10–12, 14, 19–22, 25, 26, 28–30, 33, 35, 37, 38, 44–47, 49, 50] (see the survey [23] for further references). Abstract model checking broadly refers to checking a (temporal) specification in an abstract rather than a concrete model. On the one hand, in state partition-based approaches, the abstract model is itself a transition system, possibly induced by a behavioural state equivalence such as (bi)simulation, so that the verification method for the abstract model remains unaltered, e.g., [10–12, 14, 37]. On the other hand, if the abstract model derives from any, possibly nonpartitioning, state abstraction then the verification on this abstract model relies on abstract interpretation, e.g., [5, 21, 25, 38, 45, 46]. Hybrid abstraction techniques, in between these two approaches, have also been studied, e.g., predicate abstraction [28, 33], and the mixed transition systems with several universal and/or existential abstract transition relations in [3, 22]. Moreover, several works investigated the role of complete/exact/strongly preserving abstractions in model checking, e.g., [20, 26, 37, 44], and how to refine abstract models, e.g., [11, 12, 19, 21, 25, 29, 30, 45, 47, 50]. Let us also mention that [2] develops a theory of approximation for systems of fixpoint equations in the style of abstract interpretation, showing that up-to techniques can be interpreted as abstractions.

Our refinement technique somehow resembles CounterExample-Guided Abstraction Refinement (CEGAR) [12], a popular method for automatising the abstraction generation in model checking. CEGAR deals with state partition abstractions, thus merging sets of equivalent states: one starts from a rough abstraction which is iteratively refined on the basis of spurious counterexample traces arising due to over-approximation. The approach is sound for safety properties, i.e., no false positives can be found, and complete for a significant fragment of ACTL*. It turns out that state partition abstractions are a specific instance in our approach. The attempt of proving the absence of counterexamples in LCL using an initial coarse abstraction will yield a computation by some means similar to CEGAR: failing LCL proof obligations lead to abstraction refinements which can be more general than partitions.

The general idea that model checking can be expressed as static analysis has been first investigated in [40]. This work shows that model checking of ACTL can be reduced to a logic-based static analysis formulated within the flow logic approach [42], which is then computed through a solver for the alternation-free least fixed point logic designed in [41]. This reduction technique has been later extended to the μ -calculus in [55]. One major goal of [40] was to show the close relationship and interplay between model checking and static analysis, coupled with early work in [51] and, later, in [48, 49, 52], proving that static program analysis can be reduced to model checking of modal formulae. Let us remark that the reduction of [40, 55] to a flow logic-based static analysis is given for concrete model checking only and does not encompass the chance of dealing with abstract model checking and related abstraction refinement techniques such as CEGAR, that can be instead achieved by-design in our approach.

8 Conclusion and future work

We have introduced a framework where model checking of temporal formulae in ACTL or in the universal fragment of the modal μ -calculus can be reduced to program verification, paving the way to reuse the full range of abstract interpretation techniques. Formulae are mapped to programs of the MOKA language, that are then analysed through a sound-by-construction abstract interpretation. This exposes all the possible counterexamples, although false alarms can arise. We show how false alarms can be removed by inspecting the derivability of suitable

judgements in LCL, a program logic exploiting under- and over-approximations leveraging the notion of locally complete abstract interpretation. We expect that our approach, relying on KAF, naturally applies also to logics including operators based on regular expressions (see, e.g., [6, 31]). A first candidate is Propositional Dynamic Logic (PDL) [31], a modal logic closely connected to KAT. Its distinctive feature is a modal operator $[r]\varphi$ where r is a KAT expression, which is satisfied by a state σ when all the computations of r starting from σ end up in a state satisfying φ . For instance, the property expressed in the μ_{\square} -calculus in Example 4.3 can be equivalently written in PDL as $[\text{next}]^3(n = 3 \rightarrow ([\text{next}]^4)^*(n = 3))$, where, since we work in an unlabelled setting, “next” stands for the only action in the system. Indeed, PDL smoothly fits in our setting (we refer the interested reader to [1, Appendix E]).

Future Work. A number of interesting avenues of future research remain open. Our results are limited to universal fragments of temporal logics and finite state domains. A dual theory can be easily developed for existential fragments, focusing on the generation of witnesses rather than counterexamples. It would be interesting to combine the two approaches for dealing with universal and existential operators at the same time. Some ideas could come from [22] that, for solving the problem, works with two different abstract transition relations. The restriction to finite state domains is due to the fact that our encoding of logical formulae into MOKA programs relies on the `loop?` operator for detecting infinite traces which are identified with looping traces. Further work could overcome this restriction by considering an encoding that captures non-looping infinite traces in the concrete domain and by exploiting ACC domains for the abstraction.

The use of LCL allows us to track the presence of false alarms back to the failure of local completeness proof obligations, which can be resolved by refining the abstract domain. This can be done at different levels: either refining the abstraction over states or refining the equivalence on abstract frames. A proper theory of refinements, possibly identifying optimal ways of patching the domain, is a matter of future investigations.

We point out that providing a general bound for the complexity of the abstract model checking procedure is not straightforward, as it crucially depends on the choice of the abstract domain. Different domains may induce significantly different behaviors, especially for non-partitioning abstractions or when local completeness and domain refinement techniques are applied. A precise complexity analysis tailored to specific domains is an interesting subject for future work.

Concerning the automatisation, the abstract interpreter could be easily implementable leveraging the standard toolset of abstract interpretation. The abstract interpreter should be instrumented to report, when the result is not \perp , an abstract counterexample trace to check whether the counterexample is a false or true alarm. Making refinements effective requires working in a class of domains where local refinements are representable, e.g., by predicate abstractions or state partition abstractions. Developing a theory of refinements within specific subclasses of domains is a direction of future work.

References

- 1 Paolo Baldan, Roberto Bruni, Francesco Ranzato, and Diletta Rigo. Model checking as program verification by abstract interpretation (extended version). *CoRR*, abs/2506.05525, 2025. URL: <https://arxiv.org/abs/2506.05525>.
- 2 Paolo Baldan, Barbara König, and Tommaso Padoan. Abstraction, up-to techniques and games for systems of fixpoint equations. In *Proceedings of CONCUR 2020*, volume 171 of *LIPICs*, pages 25:1–25:20, 2020. doi:10.4230/LIPICs.CONCUR.2020.25.

- 3 Thomas Ball, Orna Kupferman, and Greta Yorsh. Abstraction for falsification. In *Proceedings of CAV 2005*, volume 3576 of *LNCS*, pages 67–81. Springer, 2005. doi:10.1007/11513988_8.
- 4 Thomas Ball, Andreas Podelski, and Sriram K. Rajamani. Boolean and cartesian abstraction for model checking C programs. In *Proceedings of TACAS 2001*, volume 2031 of *LNCS*, pages 268–283. Springer, 2001. doi:10.1007/3-540-45319-9_19.
- 5 Gourinath Banda and John P. Gallagher. Constraint-based abstract semantics for temporal logic: A direct approach to design and implementation. In *Proceedings of LPAR 2016*, volume 6355 of *LNCS*, pages 27–45. Springer, 2010. doi:10.1007/978-3-642-17511-4_3.
- 6 Ilan Beer, Shoham Ben-David, and Avner Landver. On-the-fly model checking of RCTL formulas. In *Proceedings of CAV 1998*, volume 1427 of *LNCS*, pages 184–194. Springer, 1998. doi:10.1007/BFB0028744.
- 7 Roberto Bruni, Roberto Giacobazzi, Roberta Gori, and Francesco Ranzato. A logic for locally complete abstract interpretations. In *Proceedings of LICS 2021*. IEEE, 2021. doi:10.1109/LICS52264.2021.9470608.
- 8 Roberto Bruni, Roberto Giacobazzi, Roberta Gori, and Francesco Ranzato. Abstract interpretation repair. In *Proceedings of PLDI 2022*, pages 426–441. ACM, 2022. doi:10.1145/3519939.3523453.
- 9 Roberto Bruni, Roberto Giacobazzi, Roberta Gori, and Francesco Ranzato. A correctness and incorrectness program logic. *Journal of the ACM*, 70(2):15:1–15:45, 2023. doi:10.1145/3582267.
- 10 Doron Bustan and Orna Grumberg. Simulation-based minimization. *ACM Transactions on Computational Logic*, 4(2):181–206, 2003. doi:10.1145/635499.635502.
- 11 Edmund M. Clarke, Orna Grumberg, Somesh Jha, Yuan Lu, and Helmut Veith. Counterexample-guided abstraction refinement. In *Proceedings of CAV 2000*, volume 1855 of *LNCS*, pages 154–169. Springer, 2000. doi:10.1007/10722167_15.
- 12 Edmund M. Clarke, Orna Grumberg, Somesh Jha, Yuan Lu, and Helmut Veith. Counterexample-guided abstraction refinement for symbolic model checking. *Journal of the ACM*, 50(5):752–794, 2003. doi:10.1145/876638.876643.
- 13 Edmund M. Clarke, Orna Grumberg, Daniel Kroening, Doron A. Peled, and Helmut Veith. *Model Checking, 2nd Edition*. MIT Press, 2018. URL: <https://mitpress.mit.edu/books/model-checking-second-edition>.
- 14 Edmund M. Clarke, Orna Grumberg, and David E. Long. Model checking and abstraction. *ACM Transactions on Programming, Languages and Systems*, 16(5):1512–1542, 1994. doi:10.1145/186025.186051.
- 15 Edmund M. Clarke, William Klieber, Milos Nováček, and Paolo Zuliani. Model checking and the state explosion problem. In *Tools for Practical Software Verification, LASER, International Summer School 2011*, volume 7682 of *LNCS*, pages 1–30. Springer, 2011. doi:10.1007/978-3-642-35746-6_1.
- 16 Patrick Cousot. *Principles of Abstract Interpretation*. MIT Press, 2021. URL: <https://mitpress.mit.edu/9780262044905/>.
- 17 Patrick Cousot and Radhia Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Proceedings of POPL 1977*, pages 238–252. ACM, 1977. doi:10.1145/512950.512973.
- 18 Patrick Cousot and Radhia Cousot. Systematic design of program analysis frameworks. In *Proceedings of POPL 1979*, pages 269–282. ACM Press, 1979. doi:10.1145/567752.567778.
- 19 Patrick Cousot and Radhia Cousot. Refining model checking by abstract interpretation. *Automated Software Engineering*, 6(1):69–95, 1999. doi:10.1023/A:1008649901864.
- 20 Patrick Cousot and Radhia Cousot. Temporal abstract interpretation. In *Proceedings of POPL 2000*, pages 12–25. ACM, 2000. doi:10.1145/325694.325699.
- 21 Patrick Cousot, Pierre Ganty, and Jean-François Raskin. Fixpoint-guided abstraction refinements. In *Proceedings of SAS 2007*, volume 4634 of *LNCS*, pages 333–348. Springer, 2007. doi:10.1007/978-3-540-74061-2_21.

- 22 Dennis Dams, Rob Gerth, and Orna Grumberg. Abstract interpretation of reactive systems. *ACM Transactions on Programming, Languages and Systems*, 19(2):253–291, 1997. doi:10.1145/244795.244800.
- 23 Dennis Dams and Orna Grumberg. Abstraction and abstraction refinement. In Edmund M. Clarke, Thomas A. Henzinger, Helmut Veith, and Roderick Bloem, editors, *Handbook of Model Checking*, pages 385–419. Springer, 2018. doi:10.1007/978-3-319-10575-8_13.
- 24 Brian A. Davey and Hilary A. Priestley. *Introduction to Lattices and Order, Second Edition*. Cambridge University Press, 2002. doi:10.1017/CB09780511809088.
- 25 Roberto Giacobazzi and Elisa Quintarelli. Incompleteness, counterexamples, and refinements in abstract model-checking. In *Proceedings of SAS 2001*, volume 2126 of *LNCS*, pages 356–373. Springer, 2001. doi:10.1007/3-540-47764-0_20.
- 26 Roberto Giacobazzi and Francesco Ranzato. Incompleteness of states w.r.t. traces in model checking. *Information and Computation*, 204(3):376–407, 2006. doi:10.1016/J.IC.2006.01.001.
- 27 Roberto Giacobazzi, Francesco Ranzato, and Francesca Scozzari. Making abstract interpretations complete. *Journal of the ACM*, 47(2):361–416, 2000. doi:10.1145/333979.333989.
- 28 Susanne Graf and Hassen Saïdi. Construction of abstract state graphs with PVS. In *Proceedings of CAV 1997*, volume 1254 of *LNCS*, pages 72–83. Springer, 1997. doi:10.1007/3-540-63166-6_10.
- 29 Orna Grumberg, Martin Lange, Martin Leucker, and Sharon Shoham. When not losing is better than winning: Abstraction and refinement for the full mu-calculus. *Information and Computation*, 205(8):1130–1148, 2007. doi:10.1016/J.IC.2006.10.009.
- 30 Anubhav Gupta and Ofer Strichman. Abstraction refinement for bounded model checking. In *Proceedings of CAV 2005*, volume 3576 of *LNCS*, pages 112–124. Springer, 2005. doi:10.1007/11513988_11.
- 31 David Harel, Jerzy Tiuryn, and Dexter Kozen. *Dynamic Logic*. MIT Press, Cambridge, MA, USA, 2000. URL: <https://mitpress.mit.edu/9780262527668/>.
- 32 Ranjit Jhala and Rupak Majumdar. Software model checking. *ACM Computing Surveys*, 41(4):21:1–21:54, 2009. doi:10.1145/1592434.1592438.
- 33 Ranjit Jhala, Andreas Podelski, and Andrey Rybalchenko. Predicate abstraction for program verification. In Edmund M. Clarke, Thomas A. Henzinger, Helmut Veith, and Roderick Bloem, editors, *Handbook of Model Checking*, pages 447–491. Springer, 2018. doi:10.1007/978-3-319-10575-8_15.
- 34 Dexter Kozen. Kleene algebra with tests. *ACM Transactions on Programming, Languages and Systems*, 19(3):427–443, 1997. doi:10.1145/256167.256195.
- 35 Kim Guldstrand Larsen and Bent Thomsen. A modal process logic. In *Proceedings of LICS 1988*, pages 203–210. IEEE Computer Society, 1988. doi:10.1109/LICS.1988.5119.
- 36 Hans Leiß. Towards Kleene algebra with recursion. In *Proceedings of CSL 1991*, volume 626 of *LNCS*, pages 242–256. Springer, 1991. doi:10.1007/BFB0023771.
- 37 Claire Loiseaux, Susanne Graf, Joseph Sifakis, Ahmed Bouajjani, and Saddek Bensalem. Property preserving abstractions for the verification of concurrent systems. *Formal Methods in System Design*, 6(1):11–44, 1995. doi:10.1007/BF01384313.
- 38 Damien Massé. Semantics for abstract interpretation-based static analyzes of temporal properties. In *Proceedings of SAS 2002*, volume 2477 of *LNCS*, pages 428–443. Springer, 2002. doi:10.1007/3-540-45789-5_30.
- 39 Antoine Miné. Tutorial on static inference of numeric invariants by abstract interpretation. *Foundations and Trends in Programming Languages*, 4(3-4):120–372, 2017. doi:10.1561/25000000034.
- 40 Flemming Nielson and Hanne Riis Nielson. Model checking *is* static analysis of modal logic. In *Proceedings of FoSSaCS 2010*, volume 6014 of *LNCS*, pages 191–205. Springer, 2010. doi:10.1007/978-3-642-12032-9_14.

- 41 Flemming Nielson, Helmut Seidl, and Hanne Riis Nielson. A succinct solver for ALFP. *Nordic Journal of Computing*, 9(4):335–372, 2002.
- 42 Hanne Riis Nielson and Flemming Nielson. Flow logic: A multi-paradigmatic approach to static analysis. In *The Essence of Computation, Complexity, Analysis, Transformation. Essays Dedicated to Neil D. Jones [on occasion of his 60th birthday]*, volume 2566 of *LNCS*, pages 223–244. Springer, 2002. doi:10.1007/3-540-36377-7_11.
- 43 Peter W. O’Hearn. Incorrectness logic. *Proceedings of POPL 2020*, 4:10:1–10:32, 2020. doi:10.1145/3371078.
- 44 Francesco Ranzato. On the completeness of model checking. In *Proceedings of ESOP 2001*, volume 2028 of *LNCS*, pages 137–154. Springer, 2001. doi:10.1007/3-540-45309-1_10.
- 45 Francesco Ranzato and Francesco Tapparo. Making abstract model checking strongly preserving. In *Proceedings of SAS 2002*, volume 2477 of *LNCS*, pages 411–427. Springer, 2002. doi:10.1007/3-540-45789-5_29.
- 46 Francesco Ranzato and Francesco Tapparo. Generalized strong preservation by abstract interpretation. *Journal of Logic and Computation*, 17(1):157–197, 2007. doi:10.1093/LOGCOM/EXL035.
- 47 David Schmidt. Binary relations for abstraction and refinement. Technical report, Kansas State University, 2001.
- 48 David A. Schmidt. Data flow analysis is model checking of abstract interpretations. In *Proceedings of POPL 1998*, pages 38–48. ACM, 1998. doi:10.1145/268946.268950.
- 49 David A. Schmidt and Bernhard Steffen. Program analysis as model checking of abstract interpretations. In *Proceedings of SAS 1998*, volume 1503 of *LNCS*, pages 351–380. Springer, 1998. doi:10.1007/3-540-49727-7_22.
- 50 Sharon Shoham and Orna Grumberg. Monotonic abstraction-refinement for CTL. In *Proceedings of TACAS 2004*, volume 2988 of *LNCS*, pages 546–560. Springer, 2004. doi:10.1007/978-3-540-24730-2_40.
- 51 Bernhard Steffen. Data flow analysis as model checking. In *Proceedings of TACS 1991*, volume 526 of *LNCS*, pages 346–365. Springer, 1991. doi:10.1007/3-540-54415-1_54.
- 52 Bernhard Steffen. Generating data flow analysis algorithms from modal specifications. *Science of Computer Programming*, 21(2):115–139, 1993. doi:10.1016/0167-6423(93)90003-8.
- 53 Colin Stirling and David Walker. Local model checking in the modal mu-calculus. In *Proceedings of TAPSOFT 1989*, volume 351 of *LNCS*, pages 369–383. Springer, 1989. doi:10.1007/3-540-50939-9_144.
- 54 Pierre Wolper. Temporal logic can be more expressive. *Information and Control*, 56(1/2):72–99, 1983. doi:10.1016/S0019-9958(83)80051-5.
- 55 Fuyuan Zhang, Flemming Nielson, and Hanne Riis Nielson. Model checking as static analysis: Revisited. In *Proceedings of IFM 2012*, volume 7321 of *LNCS*, pages 99–112. Springer, 2012. doi:10.1007/978-3-642-30729-4_8.