

Verifying CTL-live Properties of Infinite State Models using an SMT solver (Version 1)

Amirhossein Vakili and Nancy A. Day Cheriton School of Computer Science
 University of Waterloo
 Waterloo, Ontario, Canada, N2L 3G1
 {avakili, nday}@uwaterloo.ca

Technical Report: CS-2014-10
 May, 2014



Abstract—The ability to create and analyze abstract models is an important step in conquering software complexity. In this paper, we show that it is practical to verify dynamic properties of infinite state models expressed in a subset of CTL directly using an SMT solver without iteration, abstraction, or human intervention. We call this subset CTL-live and it consists of the operators of CTL expressible using the least fixed point operator of the mu-calculus, which are commonly considered liveness properties (e.g., AF, EU). We show that using this method the verification of an infinite state model can sometimes complete more quickly than verifying a finite version of the model. We also examine modelling techniques to represent abstract models in first-order logic that facilitate this form of model checking.

1 INTRODUCTION

Abstraction is a key element to conquering complexity in the development of software [1], [2]. We need tools that support reasoning about abstract models of systems in order to better understand our models and to detect errors earlier in the development process. Abstract models are often expressed using infinite or complex data structures. Temporal logic model checking [3] of the dynamic behaviour of models with infinite state spaces without the use of abstraction is usually considered beyond the realm of first-order logic (FOL) reasoners because of the iterative nature of the fixed point (or transitive closure) computation. However, with the recent advances in SMT (satisfiability modulo theories) solvers that have turned first-order reasoners into powerful, efficient verification tools, it is worth taking another look at the problem of how to express the temporal logic model checking problem in FOL. Some results use an SMT solver iteratively to analyze invariants of infinite state systems (e.g., [4], [5]). These methods are guaranteed to terminate without approximation only if the property is not satisfied.

In recent work [6], we showed that the validity of properties within a subset of the temporal logic CTL

(Computation Tree Logic) can be expressed in FOL directly without the use of iteration. We called this subset CTL-live, and it consists of operators that are commonly used to describe liveness properties, i.e., those expressible using the least-fixed point operator of the mu-calculus (e.g., AF, EU). We also showed that CTL-live is maximal with respect to FOL in the sense that CTL operators that are not within CTL-live (e.g., invariants) are not expressible in first-order logic.

Our FOL theory for CTL-live creates the possibility of the following practical use: model the system as a (potentially infinite) Kripke structure in FOL, add automatically generated constraints based on the CTL-live property, and give the problem to an SMT solver to solve by itself. If the property is valid, theoretically with enough resources, the SMT solver can complete the analysis because FOL is recursively enumerable. This method is elegant in its simplicity: no iteration is required, no abstraction is necessary, and no user intervention is needed to determine reachability constraints (inductive invariants).

In this paper, we evaluate the practical application of this theory through a set of case studies. We address three open questions:

- 1) Will this method work in practice? In other words, are state-of-the-art SMT solvers efficient enough to analyze properties of the dynamic behaviour of infinite state systems?
- 2) How efficient is the model checking of an infinite state model in comparison to the analysis of a finite version of the same model?
- 3) Are there modelling techniques that facilitate the use of SMT solvers for model checking?

We have chosen a varied collection of four case studies drawn from different sources. Each has an infinite state

$c : \text{Integer}$	
Initial condition: $c := 0;$	
P1: $c := c + 2;$	P2: $c := c + 3;$

Fig. 1. A simple counter




Fig. 2. The Kripke structure represented by the system in Figure 1

space through the use of integers or more complex datatypes. Our results show that our approach does work in practice and can verify liveness properties of infinite state models quickly using the SMT solvers Z3 [7] and CVC4 [8]. In fact, for some of the case studies, we show that the verification of the infinite state system completes more quickly than the verification of the same problem with limited ranges in finite solvers such as Alloy [9] and Cadence SMV [10]. Throughout the paper, we consider questions regarding modelling techniques to facilitate this form of dynamic analysis.

Finally, we address the verification of safety properties (which are not part of CTL-live). We show that the inductive invariant approach to verifying invariants is not complete for infinite state systems: an invariant cannot always be strengthened to become an inductive invariant.

We believe our results regarding the model checking of infinite state systems automatically are an exciting step forward in the quest to provide automatic reasoning tools for abstract models of dynamic systems.

The rest of this paper is organized as follows: Section 2 provides the background material needed to understand our results. Section 3 describes the process and the chain of tools that we use to verify our case studies. The case studies are presented in Section 4. Section 5 discusses modelling choices that have an effect on the performance of the tools we use. Section 6 presents our theoretical result on the method of finding inductive invariants for model checking safety properties. Section 7 describes related work, and Section 8 concludes the paper.

2 BACKGROUND

In this section, we briefly present the background concepts needed to understand our results. A Kripke struc-

ture is a basic way of modelling the dynamic behaviour of a system. A *Kripke structure* $K = \langle S, I, N, \mathbb{P} \rangle$ is a four tuple, where S is a set of states, I is a non-empty subset of S called the initial states, N is a binary relation over S , and \mathbb{P} is a set of labelling predicates, where each labelling predicate is a subset of S . A Kripke structure is infinite if and only if its set of states is infinite. Kripke structures are used to define the semantics of temporal logics.

Computation tree logic (CTL) is a temporal logic that allows us to describe properties over possible computation paths of a system [11]. CTL contains all the logical connectives of propositional logic and a set temporal connectives. Each temporal connective consists of two parts, a *path quantifier* and a *state quantifier*. Path quantifiers are A (for all) and E (exists). The state quantifiers are X (next state), F (eventually), G (globally), and U (strong until). The syntax of CTL is defined for a given set of labelling predicates \mathbb{P} :

$$\begin{aligned}
 \varphi &::= P \mid \neg\varphi \mid \varphi_1 \vee \varphi_2 \mid \varphi_1 \wedge \varphi_2 \quad \text{where } P \in \mathbb{P} \\
 &::= EX\varphi \mid AX\varphi \mid EF\varphi \mid AF\varphi \mid EG\varphi \mid AG\varphi \\
 &::= \varphi_1 AU\varphi_2 \mid \varphi_1 EU\varphi_2
 \end{aligned} \tag{1}$$

A Kripke structure defines a set of infinite computation paths, where each path represents a trace of execution. A computation path starting at state $s \in S$ is a sequence of states, $s_0 \mapsto s_1 \mapsto \dots$ such that $s = s_0$ and for every $i \geq 0$, $N(s_i, s_{i+1})$. The satisfiability relation for CTL, \models_c , is used to give meaning to CTL formulae. The notation $K, s \models_c \varphi$ denotes that the state s of the Kripke structure K satisfies the CTL formula φ and $K, s \not\models_c \varphi$ is used when $K, s \models_c \varphi$ does not hold. The satisfiability relation for CTL, \models_c , is defined by structural induction on φ :

$$\begin{aligned}
 K, s \models_c P &\iff P(s) \text{ holds, where } P \in \mathbb{P} \\
 K, s \models_c \neg\varphi &\iff K, s \not\models_c \varphi \\
 K, s \models_c \varphi_1 \vee \varphi_2 &\iff K, s \models_c \varphi_1 \vee K, s \models_c \varphi_2 \\
 K, s \models_c EX\varphi &\iff \exists s' \in S : N(s, s') \wedge K, s' \models_c \varphi \\
 K, s \models_c AF\varphi &\iff \text{for all paths } s_0 \mapsto s_1 \mapsto \dots \\
 &\quad \text{such that } s_0 = s \text{ there exists an} \\
 &\quad i \text{ such that } K, s_i \models_c \varphi. \\
 K, s \models_c \varphi_1 EU\varphi_2 &\iff \text{there exists a } j \text{ and a path,} \\
 &\quad s_0 \mapsto s_1 \mapsto \dots, \text{ such that} \\
 &\quad s = s_0, K, s_j \models_c \varphi_2, \text{ and} \\
 &\quad \text{for all } i < j \text{ } K, s_i \models_c \varphi_1.
 \end{aligned}$$

The connectives that are mentioned in here form a complete fragment for CTL; e.g., $EG\varphi$ is equivalent to $\neg AF\neg\varphi$. A Kripke structure K satisfies the CTL formula φ , denoted by $K \models_c \varphi$, iff for all $s \in I$ we have $K, s \models_c \varphi$.

Example 1: Figure 1 represents a simple asynchronous counter system. There are two processes in this system, P1 and P2. They both have access to a shared variable of type integer, c . At each moment in time, one of the processes changes the value of c : P1 increments c by 2 units and P2 by 3 units. In this asynchronous system, the order in which the processes are executed is not known; as a result, when $c=100$, in the next step $c=102$ or $c=103$ depending on which process has been executed. The initial value of c in this

system is 0. We want to use CTL to study the behaviour of this system. Figure 2 represents a part of the Kripke structure of the counter system in Figure 1. The set of states of this Kripke structure K is the set of all positive integers, $S = \text{Integer}$. This Kripke structure has only one initial state, $I = \{0\}$, and the transition relation has the following property:

$$(c, c') \in N \quad \text{iff} \quad c' = c + 2 \quad \vee \quad c' = c + 3$$

This Kripke structure satisfies $EF \ c=5$, i.e., $K \models_c EF \ c=5$, since there exists a path from the initial state to a state where c becomes 5; on the other hand $K \not\models_c AF \ c=5$, since not all paths eventually make c equal to 5. If process $P1$ is executed for the first 3 steps, c never becomes 5. Since each process increases the value of c , we can see that $K \models_c AF \ c>5$ holds.

2.1 Kripke Structures in FOL and SMT solvers

First-order logic (FOL) provides quantifiers along with propositional logic connectives to describe properties over *relational* and *functional symbols* [12]. An *interpretation* determines the content of each relational and functional symbol. A problem in FOL consists of a set of formulae. Every interpretation that satisfies all these formulae is called a satisfying interpretation. Using the concept of satisfiability, logical entailment is defined as follows:

Definition 1: (Semantic entailment) Suppose Γ is a set of FOL formulae and Φ is a FOL formula: Γ entails Φ , denoted by $\Gamma \models \Phi$, iff every interpretation that satisfies all the formulae in Γ also satisfies Φ .

From this definition, we can prove that $\Gamma \models \Phi$ iff $\Gamma \cup \{\neg\Phi\}$ is unsatisfiable.

A satisfiability modulo theories (SMT) solver, for short SMT solver, is an automatic tool to check the satisfiability of a set of FOL formulae [13]. An SMT solver differs from a general purpose FOL satisfiability checker in one major way: if a built-in type such as `Integer` is used in a formula, the SMT solver only considers the “standard” interpretation for that type and the defined operations over it. SMT-LIB is a standard notation that state-of-the-art SMT solvers accept as input [14]. Specification of a problem in SMT-LIB consists of four parts: 1) declaration of user-defined types, 2) declaration of functional symbols used in the model¹, 3) definitions that are used to simplify the model, and 4) a set of constraints, where each constraint is a formula. SMT-LIB does not distinguish between terms and formulae. A formula is a term of type `Bool`. To ease the parsing of SMT-LIB specifications by SMT solvers, each SMT-LIB specification is a sequence of S-expressions.

Example 2: Figure 3 presents the SMT-LIB specification of the Kripke structure of Figure 2. This specification does not contain user-defined types. Lines 1 and 2 declare that `Init` and `Next` are relational symbols over `Int` and `Int × Int`

Symbol	Meaning
$\Gamma \models \Phi$	FOL entailment
$\mathcal{K} \models_c \varphi$	CTL model checking
$\mathcal{D} \models_v \varphi$	Universal model checking

Fig. 4. Summary of satisfiability notations

respectively. To increase the readability of this specification, we have defined $P1$ and $P2$ in Lines 3 and 4. A definition is essentially a macro. In Lines 3-6, c represents the current state and cn the value of c in the next state. Line 5 is a constraint stating that the state c is an initial state iff it is equal to zero: ($= \ c \ 0$). The constraint in Line 6 states that cn is the next value of c iff either $P1$ holds between them or $P2$.

As this example suggests, to model a Kripke structure in FOL, we need at least two relational symbols: `Init` representing the set of initial states and `Next` representing the transition relation. The types of `Init` and `Next` are

`S → Bool` and `S × S → Bool` respectively. In these declarations, `S` depends on the types of the variables used in the specification.

Definition 2: (Declarative model) A declarative model of a Kripke structure, for short declarative model, is a set of declarations and logical formulae that contains the following two relation symbol for some S .

$$\text{Init} : S \rightarrow \text{Bool}, \quad \text{Next} : S \times S \rightarrow \text{Bool}$$

The SMT-LIB specification in Figure 3 is a declarative model. The formula in Line 6 of Figure 3 in infix form using the classical symbols for FOL connectives is the following:

$$\forall c, cn : \text{Int}. \text{Next}(c, cn) \Leftrightarrow P1(c, cn) \vee P2(c, cn)$$

In this case, the next state relation is uniquely defined. However, in general, a declarative model can represent a set of Kripke structures, rather than a single one. The Class of Kripke structures represented by a declarative model \mathcal{D} is denoted by $CK(\mathcal{D})$:

$$CK(\mathcal{D}) = \{\mathcal{K} \mid \mathcal{K} \text{ is satisfying interpretation of } \mathcal{D}\}$$

Under-specification of the next state relation and the use of user-defined types and operations that are not fully interpreted are the main reasons that a declarative model can represent multiple Kripke structures. In all our case studies, we uniquely define the next state relation.

For the set $CK(\mathcal{D})$, we defined two model checking questions [15]: 1) *do all the Kripke structures in $CK(\mathcal{D})$ satisfy the property?* 2) *is there a Kripke structure in $CK(\mathcal{D})$ that satisfies the property?* In this paper, we focus on the first of these questions, which we call universal model checking:

Definition 3: (Universal model checking) The universal model checking of a declarative model \mathcal{D} and a CTL for-

1. A relational symbol is a functional symbol of type `Bool`.

```

1) (declare-fun Init (Int) Bool)
2) (declare-fun Next (Int Int) Bool)
3) (define-fun P1 ((c Int) (cn Int)) Bool (= cn (+ c 2)))
4) (define-fun P2 ((c Int) (cn Int)) Bool (= cn (+ c 3)))
5) (assert (forall ((c Int)) (= (Init c) (= c 0))))
6) (assert (forall ((c Int) (cn Int)) (= (Next c cn) (or (P1 c cn) (P2 c cn)))))

```

Fig. 3. SMT-LIB specification of the model in Figure 1

Temporal part	
$\varphi ::=$	$\pi \mid EX\varphi \mid AX\varphi \mid EF\varphi \mid AF\varphi$
$::=$	$\varphi_1 \vee \varphi_2 \mid \varphi_1 \wedge \varphi_2$
$::=$	$\varphi_1 EU\varphi_2 \mid \varphi_1 AU\varphi_2$
Propositional part	
$\pi ::=$	$P \mid \neg\pi \mid \pi_1 \vee \pi_2$
	where P is a labelling predicate.

Fig. 6. CTL-live

mula φ , denoted by $\mathcal{D} \models_{\forall} \varphi$, is defined as checking whether all the Kripke structures in $CK(\mathcal{D})$ satisfy φ :

$$\mathcal{D} \models_{\forall} \varphi \iff \forall \mathcal{K} \in CK(\mathcal{D}) : \mathcal{K} \models_c \varphi$$

Figure 4 is a summary of the satisfiability notations used in this article.

2.2 CTL-live Verification as a FOL theory

In recent work [6], we presented a subset of CTL that we called CTL-live and described how to represent the verification of a CTL-live property as a semantic entailment problem in FOL. CTL-live is presented in Figure 6. The grammar of CTL-live does not allow a temporal connective to be in the scope of negation: the formula $AF\neg\varphi$ is part of CTL-live, but $\neg AF\varphi$ is not.

CTL-live includes the CTL connectives whose semantics in the mu-calculus are defined using the *least fixed-point* operator. The intuition behind reducing CTL-live model checking to FOL semantic entailment is that model checking is about verifying whether the set of initial states is included in the set of states that satisfies a property. If the CTL property under study is expressible as the *smallest* set that satisfies some FOL formulae, then checking whether the set of initial states is a subset of the smallest one is equivalent to checking whether the set of initial states is a subset of *all* of them:

$$I \subseteq \bigcap_{X \in \Theta} X \iff I \subseteq X \text{ for every } X \in \Theta$$

In this equation, Θ contains all the sets that satisfy some property and $\bigcap_{X \in \Theta} X$ is the smallest one. This property has a higher-order quantifier over sets, which is not available in FOL, but it is implicitly available in the quantification over interpretations in the definition of semantic entailment in FOL (Definition 1).

To model check a declarative model \mathcal{D} , and a CTL-live formula φ , we use functions called `theory` and `axiom`,

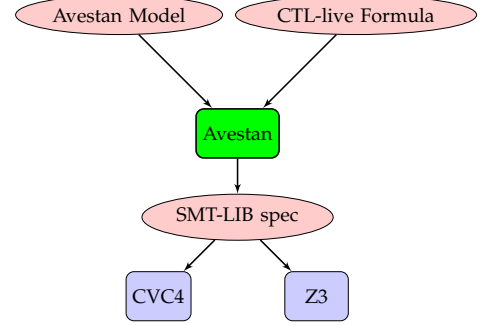


Fig. 7. Overview of our method

shown in Figure 5, to create a declarative model that is an enriched version of \mathcal{D} . The function `theory` recurses over the structure of φ . For each logical connective of CTL-live, the constraints that are added to \mathcal{D} by `theory` are defined by the (non-recursive) function `axiom`. For every sub-formula φ' of φ , `axiom` creates one or two FOL formulae and a new relational symbol $[\varphi']$ of type $S \rightarrow \text{Bool}$, which are added to \mathcal{D} . The complexity of `theory` is linear with respect to the size of φ .

The following theorem allows us to reduce CTL-live model checking to semantic entailment in FOL:

Theorem 1: (Model checking CTL-live) Let \mathcal{D} be a declarative model and φ a CTL-live formula; for universal model checking:

$$\mathcal{D} \models_{\forall} \varphi \text{ iff } \text{theory}(\mathcal{D}, \varphi) \models \forall s. I(s) \Rightarrow [\varphi](s)$$

where $[\varphi]$ is a relational symbol added to \mathcal{D} by `theory` and I is a predicate describing the initial set of states [6].

3 METHOD

Our theoretical result regarding the ability to express the verification of CTL-live properties in FOL makes it possible to create a CTL-live verification problem as a problem that can be directly solved by an SMT solver without iteration or human intervention. The approach that we use to implement our method is described in Figure 7.

A model is created in FOL using a tool that we call Avestan. Our current version of Avestan is a complete reengineering of our earlier tool [16] (also called Avestan), which was a language and tool to support the creation of declarative models. It was strongly based on Alloy [9], but the tool translated the model into an

<u>theory</u> (\mathcal{D}, φ) :	
	case φ of
	1) $P \rightarrow \mathcal{D}$
	2) $\Box\psi \rightarrow \text{theory}(\mathcal{D}, \psi) \cup \text{axiom}(\varphi)$
	3) $\psi_1 \diamond \psi_2 \rightarrow \text{theory}(\mathcal{D}, \psi_1) \cup \text{theory}(\mathcal{D}, \psi_2) \cup \text{axiom}(\varphi)$
<u>axiom</u> (φ) :	
	case φ of
1) P	$\rightarrow \{ \}$ where P is a labelling predicate
2) $\neg\psi$	$\rightarrow \{ \forall s. [\varphi](s) \Leftrightarrow \neg[\psi](s) \}$
3) $\psi_1 \vee \psi_2$	$\rightarrow \{ \forall s. [\varphi](s) \Leftrightarrow [\psi_1](s) \vee [\psi_2](s) \}$
4) $\psi_1 \wedge \psi_2$	$\rightarrow \{ \forall s. [\varphi](s) \Leftrightarrow [\psi_1](s) \wedge [\psi_2](s) \}$
5) $EX\psi$	$\rightarrow \{ \forall s. [\varphi](s) \Leftrightarrow (\exists s'. N(s, s') \wedge [\psi](s')) \}$
6) $AX\psi$	$\rightarrow \{ \forall s. [\varphi](s) \Leftrightarrow (\forall s'. N(s, s') \Rightarrow [\psi](s')) \}$
7) $EF\psi$	$\rightarrow \{ \forall s. [\psi](s) \Rightarrow [\varphi](s), \forall s, s'. N(s, s') \wedge [\psi](s') \Rightarrow [\varphi](s) \}$
8) $AF\psi$	$\rightarrow \{ \forall s. [\psi](s) \Rightarrow [\varphi](s), \forall s. (\forall s'. N(s, s') \Rightarrow [\varphi](s')) \Rightarrow [\varphi](s) \}$
9) $\psi_1 EU\psi_2$	$\rightarrow \{ \forall s. [\psi_2](s) \Rightarrow [\varphi](s), \forall s. [\psi_1](s) \wedge (\exists s'. N(s, s') \wedge [\varphi](s')) \Rightarrow [\varphi](s) \}$
10) $\psi_1 AU\psi_2$	$\rightarrow \{ \forall s. [\psi_2](s) \Rightarrow [\varphi](s), \forall s. [\psi_1](s) \wedge (\forall s'. N(s, s') \Rightarrow [\varphi](s')) \Rightarrow [\varphi](s) \}$

Fig. 5. Definitions of theory and axiom (from [6]). $\Box \in \{\neg, EX, AX, EF, AF\}$, $\diamond \in \{\vee, \wedge, EU, AU\}$ and φ is a CTL-live formula.

- 1) (declare-fun af (Int) Bool)
- 2) (assert (forall ((c Int)) (=> (> c 5) (af c))))
- 3) (assert (forall ((c Int)) (=> (forall ((cn Int)) (=> (Next c cn) (af cn))) (af c))))
- 4) (assert (not (forall ((c Int)) (=> (Init c) (af c)))))

Fig. 8. Declarations and formulae that are added to Figure 3 to model check $AF \ c > 5$

SMT-LIB specification. Our new tool is implemented in Python3 [17] and uses Python as both the object and meta-language for expressing models in FOL. It continues to produce specifications in SMT-LIB for analysis by an SMT solver.

We implemented the functions of Figure 5 in Avestan to create the constraints needed for the verification of a CTL-live property. Using Avestan, we transform a model plus these constraints into an SMT-LIB specification and check the verification problem as a satisfiability problem using both CVC4 (version 1.3) and Z3 (version 4.3.1). The following is an example that illustrates the method of applying the result of Theorem 1 to verify a declarative model using an SMT solver.

Example 3: Suppose we want to prove that in the Kripke structure of Figure 2, c eventually becomes larger than 5, by using an SMT solver. We need to prove that this Kripke structure satisfies $AF \ c > 5$. AF is part of CTL-live, therefore, we can use the result of Theorem 1. According to this theorem, we need to compute $\text{theory}(\mathcal{D}, AF \ c > 5)$, where \mathcal{D} is the declarative model in Figure 3. By the definition of theory in Figure 5, we have:

$$\text{theory}(\mathcal{D}, AF \ c > 5) = \mathcal{D} \cup \text{axiom}(AF \ c > 5)$$

Following the definition of axiom at Line 8, $\text{axiom}(AF \ c > 5)$ is a set with two constraints:

- 1) $\forall s. [\psi](s) \Rightarrow [\varphi](s)$
- 2) $\forall s. (\forall s'. N(s, s') \Rightarrow [\varphi](s')) \Rightarrow [\varphi](s)$

where $[\psi]$ is $c > 5$, N is $Next$ (of Figure 3), and $[\varphi]$ is a new relational symbol af of type $Int \rightarrow Bool$. Since the state of the system is represented by an integer, the quantification over states becomes quantification over integers. Written in terms of the model and property, these constraints are:

- 1) $\forall c: Int. c > 5 \Rightarrow af(c)$
- 2) $\forall c: Int. (\forall cn. Next(c, cn) \Rightarrow af(cn)) \Rightarrow af(c)$

Now, we need to check whether $\text{theory}(\mathcal{D}, AF \ c > 5)$ entails the following:

$$\forall c: Int. Init(c) \Rightarrow af(c) \quad (2)$$

We know that Γ entails Φ iff $\Gamma \cup \{\neg\Phi\}$ is unsatisfiable; therefore, we add the **negation** of the formula in Equation 2 to $\text{theory}(\mathcal{D}, AF \ c > 5)$ and run the SMT solver to check for the satisfiability: if it is unsatisfiable, then we can conclude that $AF \ c > 5$ holds. Figure 8 presents the declaration of af (Line 1), along with the three formulae (Lines 2-4) in SMT-LIB notation that need to be added to Figure 3 to model check $AF \ c > 5$. The output of Z3 on this model is *unsat*, which tells us that the original model satisfies $AF \ c > 5$.

4 CASE STUDIES

In this section, we present four case studies that test whether it is possible to use our theory and method to verify dynamic properties in CTL-live of abstract models using an off-the-shelf SMT solver. Our models were chosen from a variety of sources and domains. As we

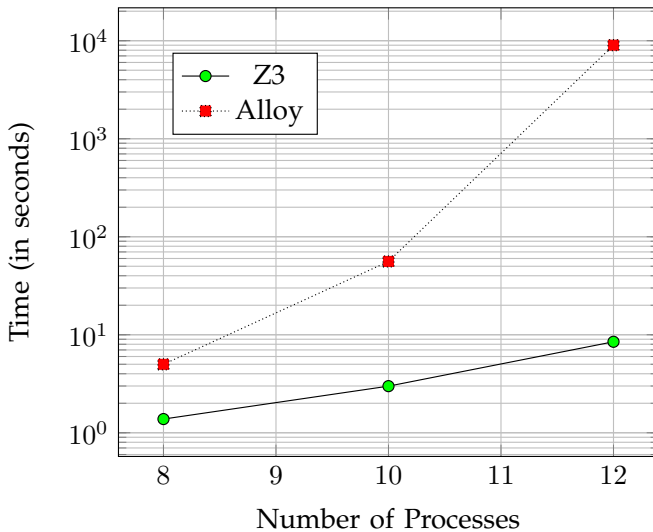


Fig. 9. Leader election model: Z3 vs Alloy

present each case study, we discuss how it is modelled in FOL and, when possible, we compare to how it was modelled and verified previously.

All our experiments were run on an Intel®Core™i7-3667U machine running Ubuntu 12.04 64-bit with up to 7.5GB of user memory. To analyze the case studies, we used the solvers in their default mode, without any flags or a customized configuration. The SMT-LIB specifications of the case studies and other models developed for this article are available on-line².

4.1 Case Study 1: Leader Election Protocol

The leader election model is a protocol to “elect” a process as the leader among a finite set of processes that form a ring [18]. A finite instance of it was previously verified by Jackson using the Alloy Analyzer [1]. In the leader election model, each process in the ring can only communicate with its successor and predecessor, and there is no centralized controller. Each process has a unique identifier (ID) and a value to represent who this process thinks is the leader of the ring (my_lead). The goal of the protocol is that every process (including the leader) will eventually recognize that the process with the greatest ID is the leader. We modelled a synchronous version of this protocol: at each moment in time, every process passes to its predecessor its value for my_lead and receives from its successor the successor’s value for my_lead . If the received value is greater than the process’ current value of my_lead , the process updates its value with the received one, otherwise, it is left unchanged. In the initial state, the value passed by a process is its own ID.

We used unbounded integers to model IDs and time. For each process, we declared a functional symbol my_lead of type $Int \rightarrow Int$. We have a fixed number of processes. The ring topology is enforced by an

ordering on the processes, where the successor of the last process is the 0th process and for any other processes such as i , the successor is $i + 1$.

The properties we verified are that every process will eventually recognize the leader:

$$AF \ (my_lead_i = lead_id)$$

where $lead_id$ is the largest ID among the current processes, my_lead_i the value of my_lead for the i th process. Thus, for i processes, we have i properties, which we checked one at a time. In this model, the set Int , which is used to represent time, is also the state space of this system. The following table shows the performance of Z3 for different numbers of processes:

12	14	16	18	20
8.48s	44.38s	3m24.64s	50m44.09s	2h37m11.69s

CVC4 with even 2 processes could not finish the verification.

When we modelled this problem with an unbounded number of processes, the verification in either SMT solver does not complete. Verification for an unbounded number of processes would likely require user intervention to deduce an invariant that would help the SMT solver verify the problem.

We also modelled this synchronous version of the algorithm in Alloy [1]. To verify the liveness properties using Alloy, we needed to finitize all sets, including time. We set the bounds on time and IDs to be the number of processes. Figure 9, which is in logarithmic scale to increase the readability of the plot, compares the performance of Z3 on models where there are no bounds on time and IDs to the performance of the Alloy Analyzer (version 4.2 using minisat) where time and IDs are bounded. As this figure shows, our approach to verification of this protocol with an infinite state space is much faster than Alloy where every set needs to be finitized.

While our verification does require a bound on the number of processes, it is significant that it does not require a bound on time. When we finitize time (as in the Alloy model), we are doing bounded model checking (BMC) [19]. When using BMC to verify a liveness property, spurious counterexamples can result because the bound is insufficient to conclude liveness. In general, computing a sufficient bound to get a reliable result is hard, and in some infinite cases it is impossible. In our SMT-LIB models, we use unbounded integers to represent time. Since SMT solvers check satisfiability with respect to standard interpretations and this interpretation for integers guarantees that Int has an infinite set of elements, our technique does not produce spurious counterexamples.

4.2 Case Study 2: Bakery Algorithm

The bakery algorithm ensures mutual exclusion between two processes that run concurrently and asynchronously [4]. Bultan, Gerber, and Pugh verified that

2. <https://cs.uwaterloo.ca/~avakili/artifact/fse14.tar>

in this algorithm the two processes cannot get into their critical sections at the same time [4]. Their method is an iterative approach that uses a Presburger arithmetic solver.

In the bakery algorithm model, the state of a process is determined by its control state value and a ticket. The value of a control state is either *Thinking*, *Waiting*, or *Critical*. A ticket is a non-negative unbounded integer. Since we have two processes, the state space of this system, S , is the following:

$$S = \{T, W, C\} \times \text{Int} \times \{T, W, C\} \times \text{Int}$$

We modelled the set $\{T, W, C\}$ as an uninterpreted type, named `ControlState`, where T, W , and C are three distinct constants of type `ControlState`. The following is a fragment of the SMT-LIB specification that models `ControlState` ensuring that each value is distinct:

```

1) (declare-sort ControlState 0)
2) (declare-fun T () ControlState)
3) (declare-fun W () ControlState)
4) (declare-fun C () ControlState)
5) (assert (not (= T W)))
6) (assert (not (= T C)))
7) (assert (not (= W C)))

```

Besides comparing the value of the tickets, this algorithm also manipulates the value of tickets using the addition operation on integers; as a result, an uninterpreted type with a total ordering would not be sufficient to express this model. Each transition in our model is defined as a functional symbol of type $S \times S \rightarrow \text{Bool}$. By combining these transitions, we modelled the next state relation.

For this case study, we verified that any process, e.g., process 1, that is waiting to get into its critical section, will eventually succeed:

$$AG (c1 = W \Rightarrow AF c1 = C) \quad (3)$$

This is an invariant property, therefore to verify this property, we needed to show that every reachable state satisfies $c1 = W \Rightarrow AF c1 = C$. AG is not part of CTL-live, therefore we cannot ask the SMT solver to prove this property directly. Instead, we created a more general property that implies the formula of Equation 3: we proved that the set of *all* states, which includes the reachable states, satisfies the following property:

$$c1 = W \Rightarrow AF (c1 = C \vee \text{dead_end})$$

where `dead_end` is true of a state iff that state does not have any next state. The reason we needed to add this part is that according to the semantics of CTL, a state that does not have any next state satisfies AX , AF , and AU properties. These connectives are satisfied if for *all* computation paths some properties hold; if there is no path starting from a state, that state vacuously satisfies the property. This simple technique eliminates the need to make a transition relation total by adding some extra behaviour to the model that is not part of the original

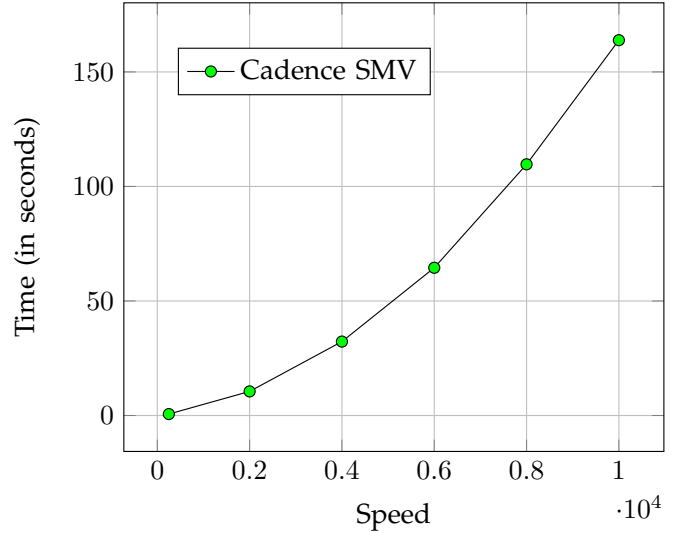


Fig. 10. Collision Avoidance model in Cadence SMV

specification. On the other hand, for CTL-live properties that use an existential path quantifier, we must not use a `dead_end` predicate.

We stated this property by making the set of initial states be the set of all states. This revised property is part of CTL-live. Z3 verified this property in 0.08 seconds and CVC4 in 8.64 seconds.

Another algorithm studied by Bultan, Gerber, and Pugh is the ticket mutual exclusion algorithm [4], [20]. We tried to verify an invariant property similar to Equation 3 for this model using a similar technique to the Bakery algorithm; however, neither SMT solver terminates within a threshold of 3 hours. It is likely that this property of this model is only satisfied within the reachable set of states and therefore it does not hold for the entire set of states.

4.3 Case Study 3: Collision Avoidance Stateflow Model

Our third case study is a Stateflow model of a collision avoidance feature used in a modern vehicle [21], [22]. It was previously used with other feature models to check for feature interactions using Cadence SMV.

This case study has control state complexity in a hierarchical, non-current state-transition model (9 basic states). However, there are two variables manipulated by the transitions of the model: speed and threshold, which determine when collision avoidance needs to be engaged. These variables are used in the triggers of transitions and thus, affect the control logic of the system and therefore are not removed by standard cone of influence reductions. In our model, the speed of a vehicle is modelled as an unbounded integer, and threshold is a constant positive integer. We verified that every basic state is reachable without a bound on speed and threshold. This property is a conjunction of 9 EF formulae. Z3 verifies all these properties together in 0.58 seconds.

CVC4 terminates in 0.89 seconds having UNKNOWN as output. The UNKNOWN means the solver cannot verify nor refute the property.

Because we had access to the original models, we can compare our results to using Cadence SMV to analyze the Stateflow models for different finite bounds on speed and threshold. Figure 10 presents these results. As this figure shows, the performance of Cadence SMV degrades as the size of speed and threshold is increased.

4.4 Case Study 4: File system

Our last case study is a file system that was originally modelled in Z [23]. Woodcock and Davies use natural deduction to prove manually that this model is consistent.

The state of the file system is represented as a partial function from `Keys` to `Data` named `content`. There are three operations that change the state of the file system: adding a new entry, deleting an existing entry and writing a new data to an existing key.

The major difference between the file system model and our other case studies is in its state space: each state is a function whereas in the other case studies, a state is a tuple that includes an infinite element. Since quantification over functions is not allowed in FOL, we cannot directly use our technique to model check a CTL-live property of this model.

Borrowing a technique used in Alloy models [1], in our model, we explicitly introduced the state space as a new uninterpreted set `State` and declared `content` as follows:

```
content: State × Key -> Data
```

where $\text{content}(s, k) = d$ is interpreted as the content of the file system at state s for the key k is d . To model the fact that `content` is a *partial* function from `Key` to `Data`, we declare a constant `NULL` of type `Data`: the value of $\text{content}(s, k)$ being equal to `NULL` means that the content of the file system at state s for the key k is empty. In Alloy, this technique manifests itself in the use of a “State” object to encapsulate the elements of the state.

The disadvantage of explicitly introducing the set `State` is that it is uninterpreted, and it may result in spurious counterexamples. For example, the following property is not entailed by this model:

$$\text{content}(s, k) \neq \text{NULL} \Rightarrow \exists s': \text{delete}(k, s, s') \quad (4)$$

This property states that if at state s the content of key k is not empty, then we can delete k from it and go to some state s' . The spurious counterexample for this property is a single state with a non-empty content. We need to ensure that interpretations that do not include enough states are eliminated from the analysis. To eliminate these spurious counterexamples, we need to “interpret” `State` by adding some axioms to the model. These

TABLE 1
Run time of Z3 and CVC4 for each case study in seconds (NA: Not applicable)

Case study	Z3	CVC4
Leader election, 12 process	8.48	NA
Leader election, 14 process	44.38	NA
Leader election, 16 process	204.64	NA
Leader election, 18 process	3044.09	NA
Leader election, 20 process	9431.69	NA
Bakery algorithm	0.08	8.64
Collision avoidance	0.58	NA
File system	0.15	0.69

axioms are called *generator axioms* [1]. For our file system model where only a performed operation can change the state, a set of standard generator axioms exist: for every operation we needed to add a formula stating that if an operation `OP` is applicable on a state s_1 , then there exists another state such as s_2 that is the result of performing `OP` on s_1 ; in other words, we needed to state that all the operations are total. For example the generator axiom for `delete` is same as the formula in Equation 4 except s and k are bounded by universal quantifiers.

We verified a bisimilarity property that the operation `write` can be simulated by some combination of `add` and `delete` for all possible states of the file system. For this purpose, we created two models with the same state space: one that includes all operations (model #1) and one that includes only `add` and `delete` (model #2). We assume that some state s_2 is the result of writing something to the file system at some state s_1 ; then, we check in model #2 that s_2 is reachable from s_1 :

$$(\text{write}(k, d, s_1, s_2) \wedge s = s_1) \Rightarrow \text{EF } s = s_2$$

Z3 verified this property in 0.15 seconds and CVC4 in 0.69 seconds.

4.5 Conclusions

Our case studies show that our method is practical for a variety of different examples. In all our models, we were able to leave some element of the model state unbounded and complete verification of a property in CTL-live. We used unbounded integers, user-declared sorts, and a partial function as part of the state. Table 1 summarizes the run times of Z3 and CVC4 for all the case studies. Z3 clearly performs better than CVC4 for the data types used in our case studies.

5 MODELLING OPTIMIZATIONS

In this section, we provide some insights about factors that can be used by a modeller to develop models that are more efficient to analyze in SMT solvers.

First, we consider the trade-off in the number of variables and the number of constraints. For the leader election case study, we have two choices for expressing the ID of the leader:

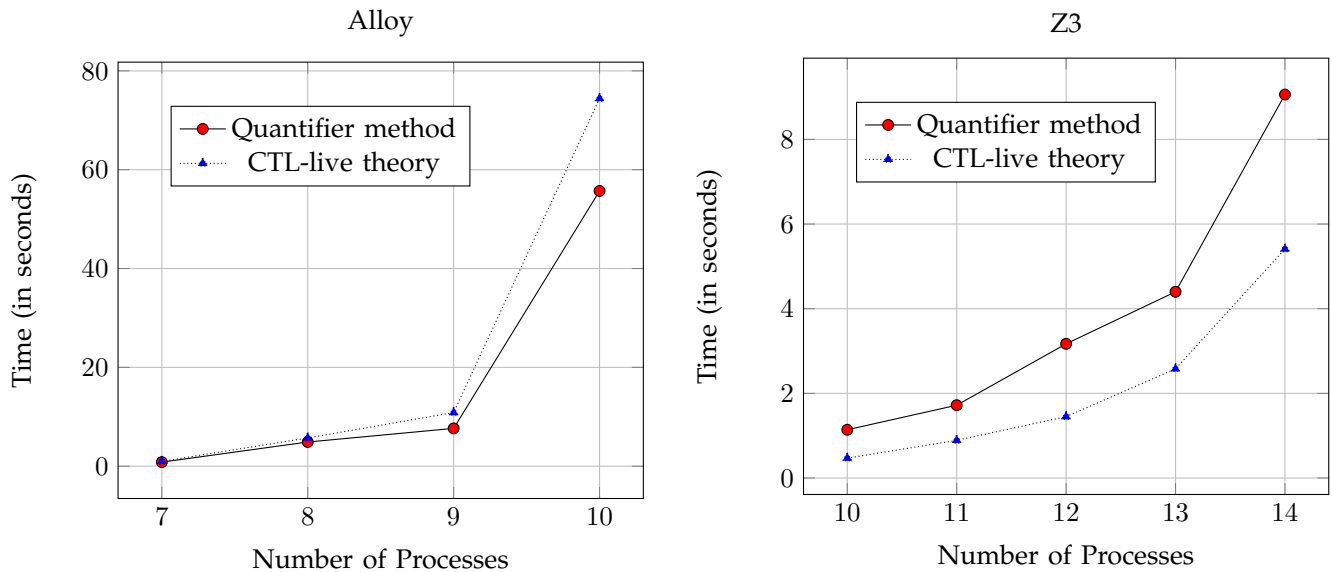


Fig. 11. Alloy and Z3 with different approaches for the leader election case study

- 1) declare a new constant and assert that this constant is equal to the ID of some process and that it is greater or equal to all the IDs; or
- 2) use the if-then-else construct in SMT-LIB and compare all the IDs with each other to determine the largest.

The first approach adds $2 \times n$ constraints and a new variable, where n is the number of processes. The second approach does not introduce any new constraints or variables, but the term that represents the greatest ID is complex. Plot ITE of Figure 12 shows the sum of the times for verifying n properties using the ITE modelling approach, where n is the number of processes. Plot 1-1 shows the same problem using the first modelling approach. Clearly, the approach of creating one more complicated constraint performed less efficiently than having a number of simple constraints with more variables in this case.

In addition, we can compare verifying n properties together (as a conjunction of constraints), or verifying each property individually. Plot ALL of Figure 12 is the result of verifying the conjunction of the properties. It clearly performs more poorly than plot 1-1, which is the sum of the times to verify each property individually. This result again supports the hypothesis that simple constraints are better for SMT solvers than complex ones.

Next, we consider the effect of the use of quantifiers in these problems. Since we use integers to model time in the leader election case study, rather than using our CTL-live theory, the eventuality property can be expressed using an existential quantifier as in:

```
∃t:Int. t > 0 ∧ my_lead_i(t) = leader_id
```

Since Alloy's input language is as expressive as FOL, we can use our `theory` function to model check a CTL-live property using the Alloy Analyzer. We set the size

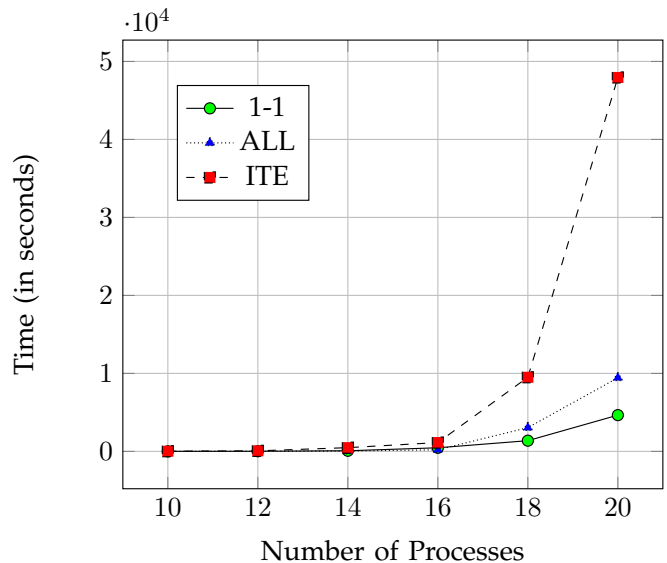


Fig. 12. Z3 on different models for the leader election problem

of all the sets in the Alloy model equal to the number of processes. Figure 11 presents the result of trying these two approaches both for Alloy and Z3. As this figure shows, the Alloy Analyzer is on average 1.27X faster when using the quantifier method to express the properties compared to our CTL-live theory in Alloy. On the other hand, Z3 on the SMT-LIB models that used our CTL-live theory was on average 1.98X faster than using the quantifier method on the model. Our conclusion from this observation is that the modelling methods also depend on the analysis tool that is used. However, Z3 using the CTL-live theory with unbounded integers was the most efficient method.

6 INDUCTIVE INVARIANTS

The verification of invariants is often of interest for safety properties of models. A property P is an *invariant* iff it holds in every *reachable* state of a Kripke structure. According to the semantics of CTL, P being an invariant of a Kripke structure \mathcal{K} is equivalent to \mathcal{K} satisfying $AG P$. $AG P$ is not part of CTL-live and previously we proved that its model checking is not reducible to FOL entailment checking:

Theorem 2: (Maximality of CTL-live) CTL-live is the largest fragment of CTL that its model checking is reducible to entailment checking in FOL; in other words, the temporal part of CTL-live cannot be extended with EG , AG , or \neg for universal model checking in FOL [6]. In our proof of this theorem, we showed that the complement of the halting problem on an empty tape for a deterministic Turing machine (DTM) is reducible to universal model checking of $EG (AG)$. The complement of the halting problem is not recursively enumerable, and as a result, it cannot be reduced to entailment checking in FOL, which is a recursively enumerable problem. We also know that $EG\varphi$ is equivalent to $\neg(\top AU\neg\varphi)$. Since AU is included in CTL-live, \neg cannot be added as well. Therefore, the verification of an invariant cannot be done by an SMT solver directly and alternative techniques are needed. A common approach to this problem is to find an inductive invariant. A property P is an *inductive invariant* for a Kripke structure \mathcal{K} iff it satisfying the following two constraints:

- 1) $\forall s : I(s) \Rightarrow P(s)$
- 2) $\forall s, s' : P(s) \wedge N(s, s') \Rightarrow P(s')$

The first constraint states that every initial state satisfies P , and the second one states that if the state s satisfies P and s' is reachable from s in one step, then s' satisfies P . It is easy to see that every inductive invariant is also an invariant of a Kripke structure, but every invariant is not necessarily an inductive invariant. Checking if a property is an inductive invariant is computationally easier than checking if it is an invariant.

According to Theorem 2, model checking AG is not recursively enumerable, whereas, inductive invariant checking is. Motivated by this fact, the inductive invariant method to check if a property P is an invariant has gained popularity for both finite and infinite Kripke structures. Many results have found inductive invariants by hand. The method of IC3 [24] is a way to find automatically inductive invariants for finite systems, and this approach has been generalized in nuXmv [5] in an incomplete approach to finding automatically inductive invariants for infinite state systems.

Generally speaking, the goal is to find an inductive invariant that is strong enough to prove the original invariant of interest. This method is essentially as follows: to prove that P is an invariant, first check if it is an inductive invariant; if it is not, then try to compute or guess an R so that $P \wedge R$ is an inductive invariant, and therefore, P is proved to be an invariant. The formula R

tries to eliminate unreachable states that do not allow P to be an inductive invariant.

A important question is “does an R always exist when P is an invariant?” For finite Kripke structures, the answer is “yes” since the number of states is finite, R can enumerate all reachable states. However, for infinite state systems, we can now show that R is not guaranteed to exist.

Theorem 3: (Incompleteness of inductive invariant method) There exists a Kripke structure \mathcal{K} and a property P such that P is an invariant of \mathcal{K} and there is no formula R such that $P \wedge R$ is an inductive invariant for \mathcal{K} .

Proof: We have shown that proving a DTM does *not* halt on an empty tape is reducible to proving that a formula named $\neg\text{halt}$ is an invariant [6]. If an R exists then we can enumerate all R 's and check if $\neg\text{halt} \wedge R$ is an inductive invariant in parallel; therefore, a semi-decision procedure for the complement of the halting problem exists and it is recursively enumerable. This is a contradiction, and as a result, such an R does not always exist. \square

7 RELATED WORK

SAT and SMT solvers have been used for bounded model checking [19], [25]. These methods use a reasoner directly for model checking by expanding the transition relation for a finite number of steps.

K-induction is a technique for unbounded model checking of safety properties [26]. This technique extends bounded model checking by proving that bounded model checking for the bound K is sufficient. The number K is dominated by the diameter of a Kripke structure. The diameter is computed iteratively using a SAT solver to check the equivalence of two formulae: the equivalence holds iff no new state can be reached by taking more than K steps. In [26], termination is guaranteed due to the finiteness of the Kripke structures under study.

Bultan, Gerber, and Pugh use Presburger formulae to represent infinite sets of states symbolically [4]. Their model checking approach for invariants requires a fix-point calculation, and termination is achieved by using conservative approximation. This approach allows false negatives.

Recently, IC3 has been generalized using SMT solvers to verify safety properties of infinite systems iteratively [27], [28]. These approaches also incorporate abstraction techniques to gain better performance.

In comparison to these approaches, our approach does not require using an SMT solver iteratively, but it is only applicable to a subset of CTL. Also, our approach is theoretically guaranteed to terminate when the property is valid, whereas the other approaches terminate when the property is not satisfied.

Compositional model checking [29] and abstraction [30] are techniques that can be applied to model check an infinite state system. Our approach could be

used along with these techniques to verify safety and liveness properties of an infinite system.

8 CONCLUSION

In this paper, we have shown that it is practical to use SMT solvers, in particular Z3, to verify CTL-live properties of infinite state models without the need for iteration, abstraction, or human intervention. The system is modelled as a (potentially infinite) Kripke structure in FOL, a theory of FOL constraints is automatically generated based on the property, and the problem is given to an SMT solver to solve by itself. The decidability of analysis is based on the subset of FOL used to express the model. Because FOL is recursively enumerable, with enough resources, the analysis will terminate if the property is valid. We have also shown that the analysis of infinite state systems using an SMT solver can be more efficient than the analysis of a finite version of the model. SMT solvers must use some deductive analysis (rather than just state space search) and therefore can take advantage of structures found in abstract models. We discussed modelling techniques that facilitate efficient model checking using SMT solvers. Finally, we proved that inductive invariants do not always exist for safety properties of infinite state systems.

In the future, we plan to investigate the scalability of the approach for modelling larger examples. However, even though the textual size of our case studies are all fairly small, the modelling efficiencies gained by representing systems abstractly offset somewhat concerns with respect to scalability. We are also exploring less primitive ways to write FOL models, the analysis of models with richer data types and quantifiers, and ways to more easily understand counterexamples produced by SMT solvers for temporal logic properties.

REFERENCES

- [1] D. Jackson, *Software Abstractions - Logic, Language, and Analysis*. MIT Press, 2006.
- [2] J. Kramer, "Is Abstraction the Key to Computing?" *CACM*, vol. 50, no. 4, pp. 36–42, Apr. 2007.
- [3] E. Clarke, O. Grumberg, and D. A. Peled, *Model Checking*. MIT Press, 1999.
- [4] T. Bultan, R. Gerber, and W. Pugh, "Symbolic Model Checking of Infinite State Systems Using Presburger Arithmetic," in *CAV*, ser. LNCS. Springer, 1997, vol. 1254, pp. 400–411.
- [5] R. Cavada, A. Cimatti, M. Dorigatti, A. Mariotti, A. Micheli, S. Mover, A. Griggio, M. Roveri, and S. Tonetta, "The nuXmv Symbolic Model Checker," Tech. Rep., Feb. 2014.
- [6] A. Vakili and N. A. Day, "Reducing CTL-live Model Checking to Semantic Entailment in First-Order Logic (Version 1)," Cheriton School of Computer Science, University of Waterloo, Tech. Rep. CS-2014-05, 2014.
- [7] L. de Moura and N. Björner, "Z3: An Efficient SMT Solver," in *TACAS*, ser. LNCS. Springer, 2008, vol. 4963, pp. 337–340.
- [8] K. L. Barrett and C. Tinelli, "CVC3," in *CAV*. Springer-Verlag, 2007, pp. 298–302.
- [9] D. Jackson, "Alloy: a lightweight object modelling notation," *ACM TOSEM*, vol. 11, no. 2, pp. 256–290, 2002.
- [10] K. L. McMillan, "The SMV system," Nov. 06 1992. [Online]. Available: <http://www.kenmcmil.com/language.ps>
- [11] E. M. Clarke, E. A. Emerson, and A. P. Sistla, "Automatic Verification of Finite-State Concurrent Systems Using Temporal Logic Specifications," *ACM TOPLAS*, pp. 244–263, 1986.
- [12] M. Huth and M. Ryan, *Logic in Computer Science, Modelling and Reasoning about Systems*, 2nd ed. Cambridge University Press, 2004.
- [13] C. Barrett, R. Sebastiani, S. Seshia, and C. Tinelli, *Satisfiability Modulo Theories*, ser. Frontiers in Artificial Intelligence and Applications. IOS Press, February 2009, vol. 185, ch. 26, pp. 825–885.
- [14] C. Barrett, A. Stump, and C. Tinelli, *The SMT-LIB Standard Version 2.0 Reference Manual*, Jan. 2010. [Online]. Available: <http://www.smt-lib.org>
- [15] A. Vakili and N. Day, "Temporal Logic Model Checking in Alloy," in *ABZ*, ser. LNCS. Springer, 2012, vol. 7316, pp. 150–163.
- [16] —, "Avestan: A declarative modeling language based on SMT-LIB," in *ICSE Workshop on Modeling in Software Engineering (MISE)*, June 2012, pp. 36–42.
- [17] "Python," <http://www.python.org>.
- [18] E. Chang and R. Roberts, "An Improved Algorithm for Decentralized Extrema-finding in Circular Configurations of Processes," *CACM*, vol. 22, no. 5, pp. 281–283, 1979.
- [19] A. Biere, A. Cimatti, E. Clarke, and Y. Zhu, "Symbolic Model Checking without BDDs," in *TACAS*, ser. LNCS. Springer, 1999, vol. 1579, pp. 193–207.
- [20] G. R. Andrews, *Concurrent Programming: Principles and Practice*. Benjamin-Cummings Publishing Co., Inc., 1991.
- [21] A. L. Juarez Dominguez, N. A. Day, and R. T. Fanson, "A Preliminary Report on Tool Support and Methodology for Feature Interaction Detection," Cheriton School of Computer Science, University of Waterloo, Tech. Rep. CS-2007-44, 2007.
- [22] A. L. J. Dominguez, "Detection of Feature Interactions in Automotive Active Safety Features," Ph.D. dissertation, Cheriton School of Computer Science, University of Waterloo, May 2012.
- [23] J. Woodcock and J. Davies, *Using Z: Specification, Refinement, and Proof*. Prentice-Hall, Inc., 1996.
- [24] A. Bradley, "SAT-Based Model Checking without Unrolling," in *Verification, Model Checking, and Abstract Interpretation*, ser. LNCS. Springer Berlin Heidelberg, 2011, vol. 6538, pp. 70–87.
- [25] T. Schüle and K. Schneider, "Bounded model checking of infinite state systems," *Formal Methods in System Design*, pp. 51–81, 2007.
- [26] M. Sheeran, S. Singh, and G. Stålmarck, "Checking Safety Properties Using Induction and a SAT-Solver," in *FMCAD*, ser. LNCS. Springer, 2000, vol. 1954, pp. 127–144.
- [27] A. Cimatti and A. Griggio, "Software Model Checking via IC3," in *CAV*, ser. LNCS. Springer Berlin Heidelberg, 2012, vol. 7358, pp. 277–293.
- [28] A. Cimatti, A. Griggio, S. Mover, and S. Tonetta, "IC3 modulo theories via implicit predicate abstraction," *CoRR*, vol. abs/1310.6847, 2013.
- [29] K. McMillan, "Verification of infinite state systems by compositional model checking," in *Correct Hardware Design and Verification Methods*, ser. LNCS. Springer Berlin Heidelberg, 1999, pp. 219–237.
- [30] E. M. Clarke, O. Grumberg, and D. E. Long, "Model checking and abstraction," *ACM Trans. Program. Lang. Syst.*, vol. 16, no. 5, pp. 1512–1542, 1994.