

Fairness Modulo Theory: A New Approach to LTL Software Model Checking

Daniel Dietsch^(✉), Matthias Heizmann,
Vincent Langenfeld, and Andreas Podelski

University of Freiburg, Freiburg im Breisgau, Germany
dietsch@informatik.uni-freiburg.de



Abstract. The construction of a proof for unsatisfiability is less costly than the construction of a ranking function. We present a new approach to LTL software model checking (i.e., to statically analyze a program and verify a temporal property from the full class of LTL including general liveness properties) which aims at exploiting this fact. The idea is to select finite prefixes of a path and check these for infeasibility before considering the full infinite path. We have implemented a tool which demonstrates the practical potential of the approach. In particular, the tool can verify several benchmark programs for a liveness property just with finite prefixes (and thus without the construction of a single ranking function).

1 Introduction

The long line of research on *software model checking*, i.e., on tools that statically analyze a given program in order to automatically verify a given temporal property, was initially restricted to safety properties [2, 3, 11, 20, 37, 45, 51]. It was later extended to termination [9, 21, 26, 27, 36, 40, 49, 50, 52]. The relative maturity of this research is reflected by the fact that software model checking tools successfully participate in the software verification competition SV-Comp [10], for safety [29, 33, 41, 47] as well as for termination [33, 55, 56].

In a more recent trend, approaches to software model checking are emerging for the general class of LTL properties, and in particular general liveness properties [5, 22–24]. In this paper, we introduce an approach to LTL software model checking which is based on *fairness modulo theory*, an extension of *reachability modulo theory* as introduced by Lal and Qadeer [42].

In the setting of [42], the existence of a program execution that violates a given safety property is proven via the reachability of an error location of the program along a *feasible* path. A path is feasible if the sequence of statements along the path is executable. This condition is checked by checking whether the corresponding logical formula is *satisfiable modulo theory* (i.e., satisfiable in the logical theory of integers, arrays, etc.). Today, quite efficient *SMT solvers* exist which can not only prove unsatisfiability but also compute *interpolants* [13, 14, 17, 19, 46]. Interpolants can be used to generalize the proof of unsatisfiability in

order to show the infeasibility of more and more paths and eventually prove the unreachability of an error location (which is the underlying idea in the approach to program verification of [35, 36]).

We extend the setting of [42] to LTL by defining the construction of a new kind of program (a *Büchi program*) from the input program and the LTL property. The control flow graph of a Büchi program comes with a distinguished set of nodes which is used to define (infinite) *fair* paths (a path is fair if it visits the distinguished set of nodes infinitely often). Now, in our extension of the setting in [42], the existence of a program execution that violates a given LTL property is proven via the existence of a feasible *fair* path.

In general, to show that the infinite sequence of statements along a path is not executable, one needs to construct a *ranking function*. For example, for each of the two infinite sequences of statements below, one may construct the ranking function r defined by $r(x, y) = x - y$.

$$\begin{aligned} \tau_1 : & \quad \text{x--} \quad \text{x>y} \quad \text{x--} \quad \text{x>y} \quad \text{x--} \quad \text{x>y} \quad \dots \\ \tau_2 : & \quad \text{x:=y} \quad \text{x>y} \quad \text{x--} \quad \text{x>y} \quad \text{x--} \quad \text{x>y} \quad \dots \end{aligned}$$

Every finite prefix of τ_1 is executable. In contrast, τ_2 has the prefix $\text{x:=y} \quad \text{x>y}$ which already is not executable.

In the case where an infinite sequence of statements has a finite prefix such that already the prefix is not executable, it is not necessary to construct a ranking function. Instead, it is sufficient to consider the prefix and prove the unsatisfiability of the logical formula corresponding to the finite sequence of statements in the prefix.

Tools exist that, given an infinite sequence of statements like $(\text{x>y} \quad \text{x--})^\omega$ or $\text{x:=y} \quad \text{x>y} \quad (\text{x>y} \quad \text{x--})^\omega$, can construct a ranking function like r above automatically [7, 12, 48]. Recent efforts go into improving the scope and the scalability of such tools [8, 25, 34, 43]. In comparison with proving unsatisfiability, the task of constructing a ranking function will always be more costly. Hence, substituting the construction of a ranking function by the construction of a proof of unsatisfiability carries an interesting potential for optimization. The goal of the work in this paper is to investigate whether this potential can be exploited practically. We develop a practical method and tool for LTL software model checking that shows that this is indeed the case.

In the remainder of the paper, after discussing an example, we introduce *Büchi programs* (as described above, we reduce the validity of an LTL property for a given program to the absence of a feasible fair path in a Büchi program). We present an algorithm that constructs such a Büchi program and checks whether it has a feasible fair path. The algorithm selects certain finite prefixes of a path for the check of feasibility before the full infinite path is considered. We then present the evaluation of a tool which implements the algorithm. Our evaluation shows the practical potential of our approach. In particular, the tool can verify several benchmark programs—for a liveness property—just with finite prefixes (and thus without the construction of a single ranking function).

2 Example

In this section we demonstrate how we apply our approach to the program \mathcal{P} depicted in Fig. 1a and the LTL property $\varphi = \Box(x > 0 \rightarrow \Diamond(y = 0))$.

We represent the program \mathcal{P} by the graph depicted in Fig. 1b. The edges of this graph are labeled with program statements. We use the Büchi automaton $\mathcal{A}_{\neg\varphi}$ depicted in Fig. 1c as representation of the negation of the LTL property φ .




Fig. 1. Program \mathcal{P} is shown in (a) as pseudocode and in (b) as control flow graph. The Büchi automaton $\mathcal{A}_{\neg\varphi}$ that represents the negation of the LTL property $\varphi = \Box(x > 0 \rightarrow \Diamond(y = 0))$ is shown in (c).

As a first step we construct the *Büchi program* \mathcal{B} depicted in Fig. 2. Afterwards we will show that this Büchi program \mathcal{B} has no path that is fair and feasible, thus proving that \mathcal{P} satisfies the LTL property φ .

A Büchi program is a program together with a fairness constraint: an execution is *fair* if a *fair location* is visited infinitely often. The fair locations of \mathcal{B} are highlighted by double circles. The locations of the Büchi program \mathcal{B} are pairs whose first element is a location of the program \mathcal{P} and whose second element is a state of the Büchi automaton $\mathcal{A}_{\neg\varphi}$. The edges of the Büchi program \mathcal{B} are labeled with sequential compositions of two statements where the first element is a statement of the program. The second element of the sequential composition is an assume statement that represents a letter of the Büchi automaton $\mathcal{A}_{\neg\varphi}$.

A key concept in our analysis is the notion of a *trace*. A trace is an infinite sequence of statements. We call a trace *fair* if it is the labeling of a path that visits some fair location infinitely often. A trace is *feasible* if it corresponds to some program execution. An example for a fair trace is $\tau_1\tau_2^\omega$ where τ_1 and τ_2 are as follows.

τ_1 : x := *; y := 1 !(y == 0) && (x > 0) !(x > 0) !(y == 0)
 τ_2 : x := *; y := 1 !(y == 0) !(x > 0) !(y == 0)

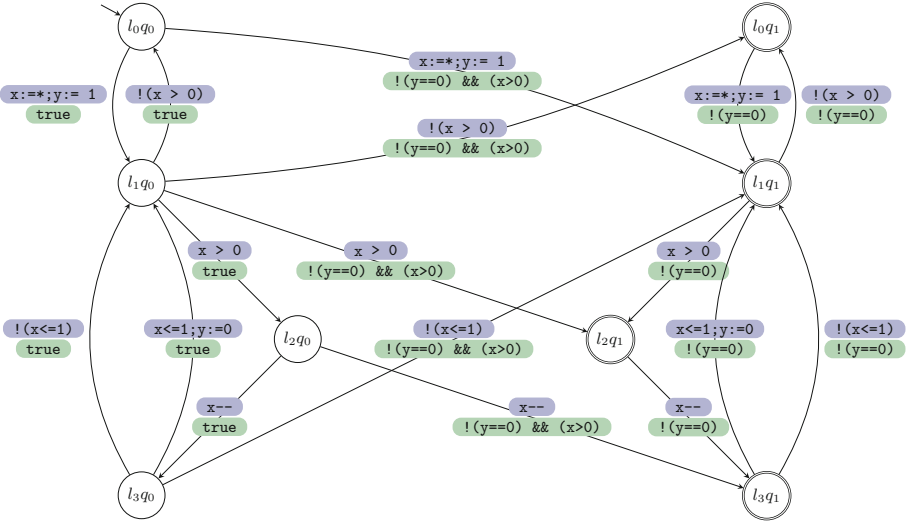


Fig. 2. The Büchi program \mathcal{B} constructed from the program \mathcal{P} (Fig. 1b) and the Büchi automaton representing $\neg\varphi$ (Fig. 1c). Each edge is labeled with the statements S_1 S_2 where S_1 comes from \mathcal{P} and S_2 comes from $\neg\varphi$. The fair locations are $l_{0q1}, l_{1q1}, l_{2q1}$ and l_{3q1} , i.e., all locations that contain the Büchi automaton’s accepting state q_1 .

This trace is not feasible because the second statement $!(y==0) \&\& (x>0)$ and the third statement $!(x>0)$ are contradicting each other.

Our algorithm constructs Büchi programs such that each fair and feasible trace of the Büchi program corresponds to a feasible trace of the original program that violates the LTL property.


In order to show that \mathcal{P} satisfies φ we show that no fair trace of the Büchi program \mathcal{B} is feasible. Thus, our algorithm tries to find arguments for infeasibility of fair traces in \mathcal{B} :

Local Infeasibility. In the Büchi program \mathcal{B} every trace that is the labeling of a path that contains the edge

$$l_{3q1} \quad x \leq 1; y := 0 \quad !(y == 0) \quad l_{1q1}$$


is infeasible, because the statements $y := 0$ and $!(y == 0)$ contradict each other. Another example for *local infeasibility* is the edge from l_{1q0} to l_{0q1} which is labeled with the two statements $!(x > 0)$ and $(x > 0)$ that contradict each other, too.

Infeasibility of a Finite Prefix. Every trace that is the labeling of a path that has the following finite prefix



is infeasible because $!(x <= 1)$ contradicts $!(x > 0)$. Another example for infeasibility of a finite prefix is the trace $\tau_1\tau_2^\omega$ that was discussed before.

ω -Infeasibility. Every trace that is the labeling of an infinite path that eventually loops along the following edges



is infeasible because $x--$ infinitely often decreases x . Thus, the value of x will eventually contradict $!(x <= 1)$. The formal termination argument is the ranking function $f(x) = x$.

Each fair trace of \mathcal{B} is infeasible for one of the reasons mentioned above. Hence, we can conclude that program \mathcal{P} indeed satisfies the LTL property φ .

All reasons for infeasibility that fall into the classes *Local infeasibility* or *infeasibility of a finite prefix* are comparatively cheap to detect. In this example we only needed to synthesize one ranking function, which is in general more expensive.

3 Preliminaries

Programs and Traces. In our formal exposition we consider a simple programming language whose statements are assignment, assume, and sequential composition. We use the syntax that is defined by the following grammar

$$s := \text{assume } \text{bexpr} \mid x := \text{expr} \mid s; s$$

where Var is a finite set of program variables, $\mathbf{x} \in Var$, expr is an expression over Var and bexpr is a Boolean expression over Var . For brevity we use bexpr to denote the assume statement $\text{assume } \text{bexpr}$.

We represent a *program* over a given set of statements $Stmt$ as a labeled graph $\mathcal{P} = (Loc, \delta, l_0)$ with a finite set of nodes Loc called locations, a set of edges labeled with statements, i.e., $\delta \subseteq Loc \times Stmt \times Loc$, and a distinguished node l_0 which we call the initial location.

In the following we consider only programs where each location has at least one outgoing edge, i.e. $\forall l \in Loc, \exists s \in Stmt, \exists l' \in Loc \bullet (l, s, l') \in \delta$. We note that each program can be transformed into this form by adding to each location without outgoing edges a selfloop that is labeled with $\text{assume } \text{true}$.

We call an infinite sequence of statements $\tau = s_0s_1s_2\dots$ a *trace of the program* \mathcal{P} if τ is the edge labeling of an infinite path that starts at the initial location l_0 . We define the set of all program traces formally as follows.

$$T(\mathcal{P}) = \{s_0s_1\dots \in Stmt^\omega \mid \exists l_1, l_2, \dots \bullet (l_i, s_i, l_{i+1}) \in \delta, \text{ for } i = 0, 1, \dots\}$$

Let \mathcal{D} be the set of values of the program's variables. We denote a program state σ as a function $\sigma : \text{Var} \rightarrow \mathcal{D}$ that maps program variables to values. We use S to denote the set of all program states. Each statement $s \in \text{Stmt}$ defines a binary relation ρ_s over program states which we call the *successor relation*. Let Expr be set of all expressions over the program variables Var . We assume a given interpretation function $\mathcal{I} : \text{Expr} \times (\text{Var} \rightarrow \mathcal{D}) \rightarrow \mathcal{D}$ and define the relation $\rho_s \subseteq S \times S$ inductively as follows.

$$\rho_s = \begin{cases} \{(\sigma, \sigma') \mid \mathcal{I}(\text{bexpr})(\sigma) = \text{true} \text{ and } \sigma = \sigma'\} & \text{if } s \equiv \text{assume bexpr} \\ \{(\sigma, \sigma') \mid \sigma' = \sigma[x \mapsto \mathcal{I}(\text{expr})(\sigma)]\} & \text{if } s \equiv x := \text{expr} \\ \{(\sigma, \sigma') \mid \exists \sigma'' \bullet (\sigma, \sigma'') \in \rho_{s_1} \text{ and } (\sigma'', \sigma') \in \rho_{s_2}\} & \text{if } s \equiv s_1 ; s_2 \end{cases}$$

Given a trace $\tau = s_0 s_1 s_2 \dots$, a sequence of program states $\pi = \sigma_0 \sigma_1 \sigma_2 \dots$ is called a *program execution of the trace* τ if each successive pair of program states is contained in the successor relation of the corresponding statement of the trace, i.e., $(\sigma_i, \sigma_{i+1}) \in \rho_{s_i}$ for $i \in \{0, 1, \dots\}$. We call a trace τ *infeasible* if it does not have any program execution, otherwise we call τ *feasible*. We use $\Pi(\tau)$ to denote the set of all program executions of τ . The set of all feasible trace of program \mathcal{P} is denoted by $T_{\text{feas}}(\mathcal{P})$, and the set of all program executions of \mathcal{P} is defined as follows.

$$\Pi(\mathcal{P}) = \bigcup_{\tau \in T_{\text{feas}}(\mathcal{P})} \Pi(\tau)$$

Büchi Automata and LTL Properties. We will not formally introduce linear temporal logic (LTL). Every LTL property can be expressed as a Büchi automaton [1]. In our formal presentation we use Büchi automata to represent LTL properties.

A *Büchi automaton* $\mathcal{A} = (\Sigma, Q, q_0, \longrightarrow, F)$ is a five tuple consisting of a finite alphabet Σ , a finite set of states Q , an initial state $q_0 \in Q$, a transition relation $\longrightarrow : Q \times \Sigma \times Q$, and a set of accepting states $F \subseteq Q$. A *word* over the alphabet Σ is an infinite sequence $w = a_0 a_1 a_2 \dots$ such that $a_i \in \Sigma$ for all $i \geq 0$. A *run* r of a Büchi automaton \mathcal{A} on w is an infinite sequence of states $q_0 q_1 \dots$, starting in the initial state such that for all $a_i \in w$ there is a transition $(q_i, a_i, q_{i+1}) \in \longrightarrow$. A run r is called *accepting* if r contains infinitely many accepting states. A word w is *accepted* by \mathcal{A} if there is an accepting run of \mathcal{A} on w . The *language* $\mathcal{L}(\mathcal{A})$ of a Büchi automaton \mathcal{A} is the set of all words that are accepted by \mathcal{A} .

An *atomic proposition* is a set of program states. An *LTL property* over a set of atomic propositions AP defines a set of words over the alphabet $\Sigma = 2^{AP}$. LTL properties are usually denoted by formulas, but several translations from formulas to equivalent Büchi automata are available [31, 32, 54]. We assume that we have given a Büchi automaton \mathcal{A}_φ for each LTL property φ .

A program state σ *satisfies* a symbol a of the alphabet 2^{AP} if σ is an element of all atomic propositions in a . A sequence of program states $\sigma_0 \sigma_1 \dots$ *satisfies* a word $a_0 a_1 a_2 \dots \in (2^{AP})^\omega$, if σ_{i+1} satisfies a_i for each $i \geq 0$. A sequence of

program states π *satisfies* the LTL property φ if π satisfies some word $w \in \mathcal{A}_\varphi$. A trace $\tau = s_0 s_1 \dots$ *satisfies* φ if it has at least one program execution and all program executions of the trace satisfy φ . A program \mathcal{P} *satisfies* φ if all program executions of \mathcal{P} satisfy φ . We will use the \models symbol to denote each of these “satisfies relations”, e.g., we will write $\mathcal{P} \models \varphi$ if the program \mathcal{P} satisfies the LTL property φ .

We note that these definitions do not put any restrictions on the initial state σ_0 of a sequence of program states. This accounts for the fact that our programs do not have to start in a given initial program state and allows programs that satisfy the LTL property $\Box(x = 0)$. For example, the program whose first statement sets the variable x to 0 and whose other statements do not modify x .

4 Büchi Program and Büchi Program Product

In this section we introduce the notion of a Büchi program, which is a program which is extended by a fairness constraint. We show that the problem whether a program satisfies an LTL property can be reduced to the problem whether a Büchi program has a fair program execution.

Definition 1 (Büchi Program). *A Büchi program $\mathcal{B} = (Stmt, Loc, \delta, l_0, Loc_{fair})$ is a program $\mathcal{P} = (Loc, \delta, l_0)$ whose set of statements is $Stmt$, with a distinguished subset of locations $Loc_{fair} \subseteq Loc$. We call the locations Loc_{fair} the fair locations of \mathcal{B} .*

An example for a Büchi program is the program depicted in Fig. 2 which was discussed in Sect. 2.

Definition 2 (Fair Trace). *A trace $s_0 s_1 s_2 \dots$ of a Büchi program \mathcal{B} is a fair trace if*

- *there exists a sequence of locations l_0, l_1, \dots such that $l_0 \xrightarrow{s_0} l_1 \xrightarrow{s_1} l_2 \xrightarrow{s_2} \dots$ is a path in \mathcal{B} , i.e., $(l_i, s_i, l_{i+1}) \in \delta$ for $i = 0, 1, \dots$, and*
- *the sequence l_0, l_1, \dots contains infinitely many fair locations.*

We use $T_{fair}(\mathcal{B})$ to denote the set of fair traces of \mathcal{B} .

If we consider the Büchi program $\mathcal{B} = (Stmt, Loc, \delta, l_0, Loc_{fair})$ as a Büchi automaton where the alphabet is the set of program statements $Stmt$, the set of states is the set of program locations Loc , the transition relation is the labeled edge relation δ the initial state is the initial location l_0 and the set of accepting states is the set of fair locations Loc_{fair} , then the language of this Büchi automaton is exactly the set of fair traces of the Büchi program.

Definition 3 (Fair Program Execution). *A program execution π of a Büchi program \mathcal{B} is a fair program execution of \mathcal{B} if π is the program execution of some fair trace of \mathcal{B} . We use $\Pi_{fair}(\mathcal{B})$ to denote the set of all fair program execution of \mathcal{B} .*

We note that traces that are fair and feasible have at least one fair program execution.

Boolean expressions over the set of program variables Var , and atomic propositions both define sets of program states. For a letter $a \in 2^{AP}$, we will use **assume** a to denote the assume statement whose expression evaluates to *true* for each state σ that satisfies a . Hence **assume** a has the following successor relation.

$$\{(\sigma, \sigma') \mid \sigma \models p \text{ for each } p \in a\}$$

Definition 4 (Büchi Program Product). Let $\mathcal{P} = (Loc, l_0, \delta_{\mathcal{P}})$ be a program over the set of statements $Stmt$, AP a set of atomic propositions over the program's variables Var , and let $\mathcal{A} = (\Sigma, Q, q_0, \rightarrow, F)$ be a Büchi automaton whose alphabet is $\Sigma = 2^{AP}$. The Büchi program product $\mathcal{P} \otimes \mathcal{A}$ is a Büchi program $\mathcal{B} = (Stmt_{\mathcal{B}}, Loc_{\mathcal{B}}, l_{0_{\mathcal{B}}}, \delta_{\mathcal{B}}, Loc_{F_{\mathcal{B}}})$ such that the set of statements consists of all sequential compositions of two statements where the first element is a statement of \mathcal{P} and the second element is a statement that assumes that a subset of atomic propositions is satisfied, i.e.,

$$Stmt_{\mathcal{B}} = \{s; \text{assume } a \mid s \in Stmt, a \in 2^{AP}\},$$

the set of locations is the Cartesian product of program locations and Büchi automaton states, i.e.,

$$Loc_{\mathcal{B}} = \{(l, q) \mid l \in Loc \text{ and } q \in Q\},$$

the initial location is the pair consisting of the program's initial location and the Büchi automaton's initial state, i.e.,

$$l_{0_{\mathcal{B}}} = (l_0, q_0),$$

the labeled edge relation is a product of the program's edge relation and the transition relation of the Büchi automaton such that an edge is labeled by the statement that is a sequential composition of the program's edge label and an assume statement obtained from the transition's letter, formally defined as follows

$$\delta_{\mathcal{B}} = \{((l, q), \{s; \text{assume } a, (l', q')\}) \mid (l, s, l') \in \delta_{\mathcal{P}} \text{ and } (q, a, q') \in \rightarrow\},$$

the set of fair locations contains all pairs where the second component is an accepting state of the Büchi automaton, i.e.,

$$Loc_{F_{\mathcal{B}}} = \{(l, q) \mid l \in Loc \text{ and } q \in F\}.$$

The following theorem shows how we can use the Büchi program product to check if a program satisfies an LTL property.

Theorem 1. *The program \mathcal{P} satisfies the LTL property φ if and only if the Büchi program product $\mathcal{B} = \mathcal{P} \otimes \mathcal{A}_{\neg\varphi}$ does not have a trace that is fair and feasible, i.e.,*

$$\mathcal{P} \models \varphi \quad \text{iff} \quad T_{\text{fair}}(\mathcal{B}) \cap T_{\text{feas}}(\mathcal{B}) = \emptyset$$

Proof. For brevity, we give only a sketch of the proof. A more detailed proof is available in an extended version of this paper [30]. First, we use the definition of the Büchi program product to show the following connection between traces of \mathcal{B} , traces of \mathcal{P} and words over 2^{AP} . s_0 ; **assume** a_0 s_1 ; **assume** $a_1 \dots \in T_{fair}(\mathcal{B})$ if and only if $s_0 s_1 s_2 \dots \in T(\mathcal{P})$ and $a_0 a_1 a_2 \dots \in \mathcal{L}(\mathcal{A}_{\neg\varphi})$. Next, we use this equivalence to show that for a sequence of program states the following holds. $\pi \in \Pi_{fair}(\mathcal{B})$ if and only if $\pi \in \Pi(\mathcal{P})$ and $\pi \models \mathcal{A}_{\neg\varphi}$. A Büchi program has a fair program execution if and only if it has a fair and feasible trace. We conclude that the intersection $T_{fair}(\mathcal{B}) \cap T_{feas}(\mathcal{B})$ is empty if and only if each program execution of \mathcal{P} satisfies the LTL property φ . \square

5 LTL Software Model Checking

In this section we describe our LTL software model checking algorithm. The algorithm is based on counter example guided abstraction refinement (CEGAR) in the fashion of [35] extended by a check for termination of fair traces and a corresponding abstraction refinement.

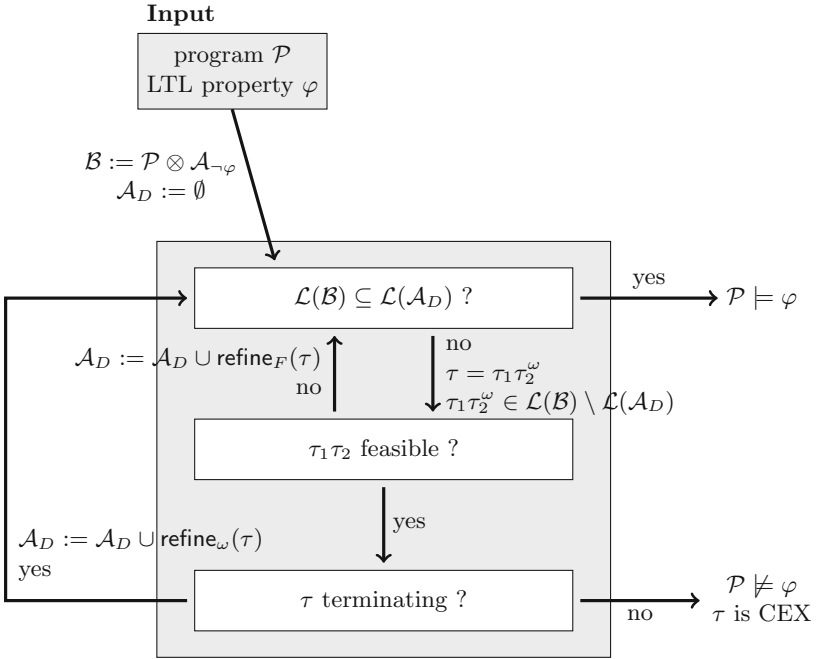


Fig. 3. The model checking algorithm. We use an automata-based approach that collects generalizations of infeasible traces in a Büchi automaton \mathcal{A}_D . The three inner boxes represent the three checks, which lead either to a refinement of \mathcal{A}_D , a result, or to a timeout (not shown).

Figure 3 shows an overview of the algorithm. The general idea is to create and continuously enlarge a Büchi automaton \mathcal{A}_D whose language contains all fair traces of \mathcal{B} that are already known to be infeasible. The algorithm starts by constructing a Büchi program \mathcal{B} with the product construction from Sect. 4. Initially, \mathcal{A}_D is a Büchi automaton that recognizes the empty language.

We use the similarities between Büchi programs and Büchi automata, i.e., that $\mathcal{L}(\mathcal{B}) = T_{feas}(\mathcal{B})$, throughout the whole algorithm. For example, in the first step of our CEGAR loop we check whether the set of fair traces represented by \mathcal{A}_D is a superset of the fair traces of \mathcal{B} (the first box in Fig. 3). This check for *trace inclusion* can be done with only Büchi automata operations.

If the set of fair traces of \mathcal{A}_D is indeed a superset of the set of fair traces of \mathcal{B} , we know that there is no fair and feasible trace in \mathcal{B} and our algorithm returns *safe*.

As the trace inclusion check is performed by computing $\mathcal{L}(\mathcal{B}) \setminus \mathcal{L}(\mathcal{A}_D)$, we will receive a fair trace τ of \mathcal{B} that witnesses that the set of fair traces of \mathcal{A}_D is no superset of the set of fair traces of \mathcal{B} . In this case, τ is always of the form $\tau_1\tau_2^\omega$.

Next, our algorithm tries to decide whether τ is feasible or not. This is done by first checking various finite prefixes for feasibility. More precisely, the stem τ_1 , the loop τ_2 and then the concatenation $\tau_1\tau_2$ are checked for feasibility in that order. If none of those finite prefixes is infeasible, our algorithm tries to prove that the full infinite trace terminates. The termination analysis (inner lower box) tries to find a ranking function to prove that the loop will terminate eventually. When non-termination can be proven, we conclude that τ is feasible. Therefore, τ is a fair and feasible trace in \mathcal{B} and thus a counterexample for the property φ . If instead termination can be shown, we know that τ is infeasible and the algorithm continues to the next step.

Note that the checks for feasibility of τ_1 and τ_2 as well as the termination analysis are based on – in general – undecidable methods. It is possible that they do not terminate. In such cases, our algorithm runs into a timeout and returns *unknown* as answer.

In the last step of the CEGAR loop we want to refine \mathcal{A}_D by adding more fair and infeasible traces. We do this by replacing \mathcal{A}_D with a Büchi automaton that is the union of the old Büchi automaton \mathcal{A}_D and a new Büchi automaton which we create from trace τ . This new Büchi automaton recognizes all fair traces of \mathcal{B} that are infeasible for the same reason for which trace τ is infeasible. Depending on the reason for infeasibility of trace τ , we use different methods for the construction of this new Büchi automaton: if τ was infeasible because we found an infeasible finite prefix, we use the method `refineF`, if τ was infeasible because we found a ranking function, we use `refineω`.

The methods `refineF` and `refineω` generalize a single trace to a set of traces. The input of these methods is the trace τ together with an infeasibility proof (resp. termination proof). The output is a Büchi automaton that accepts a set of traces whose infeasibility (resp. termination) can be shown by this infeasibility proof (resp. termination proof). `refineF` and `refineω` guarantee that at least the

single trace is contained in the language, but usually recognize a much larger set of traces. As the generalization performed by these methods is quite involved, it is not in the scope of this paper. We refer the interested reader to [35,36] for a detailed description.

6 Implementation and Evaluation

We implemented the algorithm from Sect. 5 as `ULTIMATE LTLAUTOMIZER` in the program analysis framework `ULTIMATE` [16]. This allowed us to use different, already available components for our implementation:

- a parser for ANSI C extended with specifications written in ACSL [6],
- various source-to-source transformations that optimize and simplify the input program,
- an implementation of the Trace Abstraction algorithm [35] to determine feasibility of finite trace prefixes,
- an implementation of a ranking function synthesis algorithm based on [34] to prove termination of fair traces in the Büchi program, and
- various automata operations like union, complementation and intersection of Büchi automata.

For the LTL property we use a custom annotation compatible to the ACSL format. After parsing, we transform the LTL property with `LTL2BA` [32] to a Büchi automaton, which is then together with an initial program the input for the product algorithm.

Our implementation of the product construction already contains some optimizations. For one, we already described that we remove locally infeasible traces by removing infeasible edges during the construction. We also convert the expression `e` of `assume e` statements to disjunctive normal form. If this results in edges labeled with more than one disjuncts, i.e. with `assume e1||e2||...||en`, we convert them to n edges labeled with `assume ei`. This improves the performance of the ranking function synthesis algorithm considerably.

Table 1 shows a comparison of our implementation against the benchmarks and the data provided by [23], in which the authors compare their novel LTL-checking approach based on decision predicates (DP) against a `TERMINATOR`-like procedure with an extension for fairness [21] (`Term.`). The set of benchmarks contains examples from “[...] the I/O subsystem of the windows kernel, the back-end infrastructure of the PostgreSQL database server, and the Apache web server”, as well as “some toy examples”. As the tools that were used in [23] are not publicly available, we could not re-run their implementations on our machine. Therefore, the results in the columns “`Term.`” and “`DP`” are verbatim from the original publication.

We could solve most of the benchmarks in under five seconds. Notable exceptions are “Windows OS 5”, where the other tools run into a timeout, and “Windows OS 8” where we performed much slower than DP. We are still unclear about the `OOM` result in “Apache `accept()`”, but we suspect a bug in our tool.

Table 1. The results of the comparison with the benchmarks from [23]. “Program”, “Lines”, and “ φ ” contain the name of the benchmark, the lines of code of the program, and the checked property (atomic propositions have been abbreviated). “Result” states whether the tool proved the property (\checkmark), produced a valid counterexample (\times), ran out of time (**T.O.**) or out of memory (**OOM**). **N.R.** shows the instance where we could not use the benchmark because the property was not specified explicitly and could not be guessed from the comments in the file. “Time” contains the runtime of the respective tool in seconds. For ULTIMATE LTLAUTOMIZER, there are additional statistics columns: “ $|r_F|$ ” states how many traces were refined using refine_F , and analogous “ $|r_\omega|$ ” for refine_ω . “Inc.” shows how much the product increased in size compared to the original CFG of the program. The timeout for “Term.” and “DP” was four hours, our timeout was 20 min. Our memory limit was 8GB.

Program	Lines φ	Term. [21]		DP [23]		ULTIMATE LTLAUTOMIZER				
		Time (s)	Re-sult	Time (s)	Re-sult	Time (s)	Re-sult	$ r_F $	$ r_\omega $	Inc. (%)
Ex. Sec. 2 of [23]	5 $\diamond\Box p$	2.32	\checkmark	1.98	\checkmark	0.51	\checkmark	1	0	122
Ex. Fig. 8 of [21]	34 $\Box(p \Rightarrow \diamond q)$	209.64	\checkmark	27.94	\checkmark	0.72	\checkmark	2	0	186
Toy acquire/release	14 $\Box(p \Rightarrow \diamond q)$	103.48	\checkmark	14.18	\checkmark	0.44	\checkmark	1	1	129
Toy linear arith. 1	13 $p \Rightarrow \diamond q$	126.86	(\checkmark)	34.51	(\checkmark)	1.10	\times	5	1	0.28
Toy linear arith. 2	13 $p \Rightarrow \diamond q$	T.O.	T.O.	6.74	(\checkmark)	0.82	\checkmark	4	2	0.24
PostgreSQL strmsrv	259 $\Box(p \Rightarrow \diamond\Box q)$	T.O.	T.O.	9.56	\checkmark	1.04	\checkmark	2	0	216
PostgreSQL strmsrv + bug	259 $\Box(p \Rightarrow \diamond\Box q)$	87.31	(\times)	47.16	(\times)	0.66	\checkmark	2	0	216
PostgreSQL pgarch	61 $\diamond\Box p$	31.50	(\checkmark)	15.20	(\checkmark)	0.33	\times	2	0	209
PostgreSQL dropbuf	152 $\Box p$	T.O.	T.O.	1.14	(\checkmark)	3.57	\times	1	1	148
PostgreSQL dropbuf	152 $\Box(p \Rightarrow \diamond q)$	53.99	\checkmark	27.54	\checkmark	1.37	\checkmark	2	1	168
Apache accept()	314 $\Box p \Rightarrow \Box \diamond q$	T.O.	T.O.	197.41	\checkmark	502.15	OOM	-	-	209
Apache progress	314 $\Box(p \Rightarrow (\diamond q_1 \vee \diamond q_2))$	685.34	\checkmark	684.24	\checkmark	2.01	\checkmark	4	0	209
Windows OS 1	180 $\Box(p \Rightarrow \diamond q)$	901.81	\checkmark	539.00	\checkmark	43.59	\checkmark	1	1	178
Windows OS 2	158 $\diamond\Box p$	16.47	\checkmark	52.10	\checkmark	0.11	\checkmark	1	0	176
Windows OS 2 + bug	158 $\diamond\Box p$	26.15	(\times)	30.37	(\times)	0.22	(\times)	1	0	174
Windows OS 3	14 $\diamond\Box p$	4.21	\checkmark	15.75	\checkmark	0.08	\checkmark	2	0	220
Windows OS 4	327 $\Box(p \Rightarrow \diamond q)$	T.O.	T.O.	1,114.18	\checkmark	1.86	\checkmark	1	3	207
Windows OS 4	327 $(\diamond p) \vee (\diamond q)$	1,223.96	\checkmark	100.68	\checkmark	-	N.R.	-	-	-
Windows OS 5	648 $\Box(p \Rightarrow \diamond q)$	T.O.	T.O.	T.O.	T.O.	20.76	\checkmark	1	16	190
Windows OS 6	13 $\diamond\Box p$	149.41	\checkmark	59.56	\checkmark	T.O.	T.O.	6	8	158
Windows OS 6 + bug	13 $\diamond\Box p$	6.06	(\times)	22.12	(\times)	0.05	(\times)	0	0	61
Windows OS 7	13 $\diamond\Box p$	T.O.	T.O.	55.77	\checkmark	0.91	\checkmark	2	11	161
Windows OS 8	181 $\diamond\Box p$	T.O.	T.O.	5.24	\checkmark	53.55	\checkmark	4	55	168

In many instances with liveness properties we did not need to provide a ranking function, because the generalization from traces that are infeasible because of infeasible finite prefixes already excluded all fair traces of the Büchi program. For the remainder, the termination arguments were no challenge, except for “Windows OS 8”: we had difficulties to generalize from many terminating traces, which also resulted in the slowdown compared to DP.

The expected increase in size of the Büchi program compared to the initial program’s CFG (Inc.) was also manageable. Interestingly, in both instances of

“Toy linear arith.” the product was even smaller than the original CFG, because we could remove many infeasible edges.

On four benchmarks ULTIMATE LTLAUTOMIZER results are different from the data in [23]: we contacted the authors and confirmed that our result for “Toy linear arith. 1” is indeed correct. We also could not run the benchmark “Windows OS 4”, because the LTL property contained variables that were not defined in the source file. We did not yet receive a response regarding this issue as well as regarding the correctness of our results in the other three instances.

Table 2. Results of ULTIMATE LTLAUTOMIZER on other benchmark sets. “RERS” are the online problems from “The RERS Grey-Box Challenge 2012” [39] and “coolant” consists of toy examples modelled after real-world embedded systems with specifications based on the LTL patterns described in [53]. Each program set contains pairs of a file and a property. “Avg. Lines” states the average lines of code in the sample set, and |Set| the number of file-property pairs. In the next five columns we use the same symbols as in Table 1 except for \star , which represents abnormal termination of ULTIMATE LTLAUTOMIZER. The last four columns show the average runtime, the average number of refinements with refine_F and refine_ω , and how much the size of the optimized product increased on average compared to the original CFG. We used the same timeout and memory limits as in Table 1.

Program set	Avg. Lines	Set	Statistics for \checkmark and \times									
			\checkmark	\times	T.O.	OOM	\star (N.R.)	Avg. Time (s)	Avg. $ \mathbf{r}_F $	Avg. $ \mathbf{r}_\omega $	Inc. (%)	
RERS P14	514	50	19	21	2	0	8	107.21	21	< 1	329	
RERS P15	1353	50	24	0	11	12	3	103.46	17	< 1	369	
RERS P16	1304	50	15	1	16	14	4	297.34	32	< 1	362	
RERS P17	2100	50	26	0	9	9	6	56.38	12	< 1	324	
RERS P18	3306	50	21	0	17	10	2	262.03	24	< 1	297	
RERS P19	8079	50	0	0	28	17	5	-	-	-	-	
coolant	65	18	6	10	2	0	0	1.75	2	1	258	
Benchmarks from Tab. 1	157	23	15	5	1	1	0 (1)	16.78	2	5	184	

We also considered two other benchmark sets (see Table 2). First, we ran the on-site problems from the RERS Grey-Box Challenge 2012 [39] (RERS). RERS is about comparing program verification techniques on a domain of problems comparable to the ones seen in embedded systems engineering. For this, they generate control-flow-intensive programs that contain a so-called ECA-engine (event-condition-action): one non-terminating while loop which first reads an input symbol, then calls a function that based on the current state and the input calculates an output symbol, and finally writes this output symbol. We took all 6 problem classes from the on-site part of the challenge and tried to solve them with our tool. The classification (P14 to P19) encodes the size and assumed difficulty of the problem class: P14 and P15 are small, P16 and P17 are

medium, and P18 and P19 are large problems. Inside a size bracket, the larger number means a higher difficulty.

We were able to verify roughly 43% of the RERS benchmarks without any modifications. The RERS set also helped us finding a bug that one of our optimizations on the Büchi program product introduced and which is responsible for all but four of the ★ results. For the remaining four examples, ★ occurred because ULTIMATE LTLAUTOMIZER was unable to synthesize a ranking function. Interestingly, the RERS benchmarks did seldomly require generalizations with refine_ω . In most cases, the refine_F already excluded all fair traces from the Büchi program. This trend can also be observed in the number of refine_ω applications on the benchmarks that timed out (not shown in Table 2).

Second, we used a small toy example modeled after an embedded system, a coolant facility controller that encompasses two potentially non-terminating loops in succession. The first polls the user for the input of a sane temperature limit (except one example all versions of the coolant controller can loop infinitely in this step if the input is not suitable). The second loop polls the temperature, does some calculations, increments a counter and sets the “spoiled goods” flag if the temperature limit is exceeded. The LTL properties specify that the spoiled variable cannot be reset by the program (safety), that setup stages occur in the correct order (safety and liveness), and that the temperature controlling loop always progresses (safety and liveness). We then introduced various bugs in the original version of the program and checked against the property and its negation. Although the coolant examples are quite small, they contain complex inter-dependencies between traces which lead to timeouts in two cases.

An unexpected result of the evaluation was, that the initial size of the program does not seem to define the performance of the verification, both in time and success rate, as the larger programs from P17 and P18 had more results and were faster than their counterparts from P15 and P16. Also, the effective blow-up due to the product construction is no more than four times, which is still quite manageable.

The benchmark sets together with ULTIMATE LTLAUTOMIZER are available from [30].

7 Related Work

An earlier approach to LTL software model checking was done in [21]. There, the authors reduced the problem to fair termination checking. Our work can be seen as improvement upon this approach, as we also use fair termination checking, but only when it is necessary. We avoid a large number of (more costly) termination checks due to our previous check for infeasible finite prefixes and the resulting generalizing refinement.

In [23], the authors reduce the LTL model checking problem to the problem of checking $\forall\text{CTL}$ by first approximating the LTL formula with a suitable CTL formula, and then refining counterexamples that represent multiple paths

by introducing non-deterministic prophecy variables in their program representation. This non-determinism is then removed through a determinization procedure. By using this technique, they try to reduce their dependence on termination proofs, which they identified as the main reason for poor performance of automata-theoretic approaches. Our approach can be seen as another strategy to reduce the reliance on many termination proofs. By iteratively refining the Büchi program with different proof techniques, we often remove complex control structures from loops and thus reduce the strain on the termination proof engine.

There exist various publicly available finite-state model checking tools that support both LTL properties and programs, but are in contrast to `ULTIMATE LTLAUTOMIZER` limited to finite-state systems: `SPIN` [38] and `Divine` [4] are both based on the Vardi-Wolper product [57] for LTL model checking. `Divine` supports C/C++ via LLVM bytecode, `SPIN` can be used with different frontends that translate programs to finite-state models, e.g. with `Bandera` [28] for Java. `NuSMV` [18] and `Cadence SMV` [44] reduce LTL model checking to CTL model checking. `NuSMV` can use different techniques like BDD symbolic model checking using symbolic fixed point, computation with BDDs, or bounded model checking using MiniSat. `Cadence SMV` uses Mu-calculus with additional fairness constraints [15].

8 Conclusion and Future Work

The encoding of the LTL program verification problem through the infeasibility of *fair* paths in a Büchi program has allowed us to define a sequence of semi-tests which can be scheduled before the full test of infeasibility of an infinite path. The occurrence of a successful semi-test (the proof of infeasibility for a finite prefix, by the construction of a proof of unsatisfiability) makes the full test redundant and avoids the relatively costly construction of a ranking function. Our experiments indicate that the corresponding approach leads to a practical tool for LTL software model checking.

We see several ways to improve performance. We may try to use alternatives to `LTL2BA` such as `SPOT` [31]; see [54]. The technique of large block encoding [11] adapted to Büchi programs, may help to reduce memory consumption.

References

1. Baier, C., Katoen, J.-P., et al.: Principles of Model Checking, vol. 26202649. MIT Press, Cambridge (2008)
2. Ball, T., Majumdar, R., Millstein, T.D., Rajamani, S.K.: Automatic predicate abstraction of C programs. In: PLDI, pp. 203–213 (2001)
3. Ball, T., Podelski, A., Rajamani, S.K.: Boolean and cartesian abstraction for model checking C programs. *STTT* 5(1), 49–58 (2003)
4. Barnat, J., Brim, L., Havel, V., Havlíček, J., Kriho, J., Lenčo, M., Ročkai, P., Štill, V., Weiser, J.: DiVinE 3.0 – an explicit-state model checker for multithreaded C & C++ programs. In: Sharygina, N., Veith, H. (eds.) CAV 2013. LNCS, vol. 8044, pp. 863–868. Springer, Heidelberg (2013)

5. Bauch, P., Havel, V., Barnat, J.: LTL model checking of LLVM bitcode with symbolic data. In: Hliněný, P., Dvořák, Z., Jaroš, J., Kofroň, J., Kořenek, J., Matula, P., Pala, K. (eds.) MEMICS 2014. LNCS, vol. 8934, pp. 47–59. Springer, Heidelberg (2014)
6. Baudin, P., Filliâtre, J.-C., Marché, C., Monate, B., Moy, Y., Prevosto, V., et al.: ACSL: ANSI/ISO C specification language, Feb 2015. <http://frama-c.com/download.html>
7. Ben-Amram, A.M.: Size-change termination, monotonicity constraints and ranking functions. In: Bouajjani, A., Maler, O. (eds.) CAV 2009. LNCS, vol. 5643, pp. 109–123. Springer, Heidelberg (2009)
8. Ben-Amram, A.M., Genaim, S.: On the linear ranking problem for integer linear-constraint loops. In: POPL, pp. 51–62 (2013)
9. Berdine, J., Cook, B., Distefano, D., O’Hearn, P.W.: Automatic termination proofs for programs with shape-shifting heaps. In: Ball, T., Jones, R.B. (eds.) CAV 2006. LNCS, vol. 4144, pp. 386–400. Springer, Heidelberg (2006)
10. Beyer, D.: Software verification and verifiable witnesses. In: Baier, C., Tinelli, C. (eds.) TACAS 2015. LNCS, vol. 9035, pp. 401–416. Springer, Heidelberg (2015)
11. Beyer, D., Cimatti, A., Griggio, A., Keremoglu, M.E., Sebastiani, R.: Software model checking via large-block encoding. In: FMCAD, pp. 25–32, IEEE (2009)
12. Bradley, A.R., Manna, Z., Sipma, H.B.: Linear ranking with reachability. In: Etessami, K., Rajamani, S.K. (eds.) CAV 2005. LNCS, vol. 3576, pp. 491–504. Springer, Heidelberg (2005)
13. Brillout, A., Kroening, D., Rümmer, P., Wahl, T.: An interpolating sequent calculus for quantifier-free presburger arithmetic. *J. Autom. Reason.* **47**(4), 341–367 (2011)
14. Bruttomesso, R., Pek, E., Sharygina, N., Tsitovich, A.: The openSMT solver. In: Esparza, J., Majumdar, R. (eds.) TACAS 2010. LNCS, vol. 6015, pp. 150–153. Springer, Heidelberg (2010)
15. Burch, J.R., Clarke, E.M., McMillan, K.L., Dill, D.L., Hwang, L.-J.: Symbolic model checking: 10^{20} states and beyond. In: LICS, pp. 428–439, IEEE (1990)
16. Christ, J., Dietsch, D., Ermis, E., Heizmann, M., Hoenicke, J., Langenfeld, V., Leike, J., Musa, B., Nutz, A., Schilling, C.: The program analysis framework ultimate, Feb 2015. <http://ultimate.informatik.uni-freiburg.de>
17. Christ, J., Hoenicke, J., Nutz, A.: SMTInterpol: an interpolating SMT solver. In: Donaldson, A., Parker, D. (eds.) SPIN 2012. LNCS, vol. 7385, pp. 248–254. Springer, Heidelberg (2012)
18. Cimatti, A., Clarke, E., Giunchiglia, E., Giunchiglia, F., Pistore, M., Roveri, M., Sebastiani, R., Tacchella, A.: NuSMV 2: an opensource tool for symbolic model checking. In: Brinksma, E., Larsen, K.G. (eds.) CAV 2002. LNCS, vol. 2404, pp. 359–364. Springer, Heidelberg (2002)
19. Cimatti, A., Griggio, A., Schaafsma, B.J., Sebastiani, R.: The MathSAT5 SMT solver. In: Piterman, N., Smolka, S.A. (eds.) TACAS 2013 (ETAPS 2013). LNCS, vol. 7795, pp. 93–107. Springer, Heidelberg (2013)
20. Clarke, E., Grumberg, O., Jha, S., Lu, Y., Veith, H.: Counterexample-guided abstraction refinement. In: Emerson, E.A., Sistla, A.P. (eds.) CAV 2000. LNCS, vol. 1855, pp. 154–169. Springer, Heidelberg (2000)
21. Cook, B., Gotsman, A., Podelski, A., Rybalchenko, A., Vardi, M.Y.: Proving that programs eventually do something good. *ACM SIGPLAN Not.* **42**, 265–276 (2007). ACM

22. Cook, B., Khlaaf, H., Piterman, N.: Fairness for infinite-state systems. In: Baier, C., Tinelli, C. (eds.) TACAS 2015. LNCS, vol. 9035, pp. 384–398. Springer, Heidelberg (2015)
23. Cook, B., Koskinen, E.: Making prophecies with decision predicates. ACM SIGPLAN Not. **46**, 399–410 (2011). ACM
24. Cook, B., Koskinen, E., Vardi, M.Y.: Temporal property verification as a program analysis task - extended version. FMSD **41**(1), 66–82 (2012)
25. Cook, B., Kroening, D., Rümmer, P., Wintersteiger, C.M.: Ranking function synthesis for bit-vector relations. FMSD **43**(1), 93–120 (2013)
26. Cook, B., Podelski, A., Rybalchenko, A.: Termination proofs for systems code. In: PLDI, pp. 415–426, ACM (2006)
27. Cook, B., Podelski, A., Rybalchenko, A.: Proving thread termination. In: PLDI, pp. 320–330 (2007)
28. Corbett, J.C., Dwyer, M.B., Hatcliff, J., Laubach, S., Pasareanu, C.S., Zheng, H., et al.: Bandera: extracting finite-state models from Java source code. In: ICSE, pp. 439–448, IEEE (2000)
29. Dangl, M., Löwe, S., Wendler, P.: Cpachecker with support for recursive programs and floating-point arithmetic - (competition contribution). In: Baier, C., Tinelli, C. (eds.) TACAS 2015. LNCS, vol. 9035, pp. 423–425. Springer, Heidelberg (2015)
30. Dietsch, D., Heizmann, M., Langenfeld, V.: Ultimate LTLAutomizer website, Feb 2015. <http://ultimate.informatik.uni-freiburg.de/ltlautomizer>
31. Duret-Lutz, A., Poitrenaud, D.: SPOT: an extensible model checking library using transition-based generalized Büchi automata. In: MASCOTS, pp. 76–83, IEEE (2004)
32. Gastin, P., Oddoux, D.: Fast LTL to Büchi automata translation. In: Berry, G., Comon, H., Finkel, A. (eds.) CAV 2001. LNCS, vol. 2102, pp. 53–65. Springer, Heidelberg (2001)
33. Heizmann, M., Dietsch, D., Leike, J., Musa, B., Podelski, A.: Ultimate automizer with array interpolation - (competition contribution). In: Baier, C., Tinelli, C. (eds.) TACAS 2015. LNCS, vol. 9035, pp. 455–457. Springer, Heidelberg (2015)
34. Heizmann, M., Hoenicke, J., Leike, J., Podelski, A.: Linear ranking for linear lasso programs. In: Van Hung, D., Ogawa, M. (eds.) ATVA 2013. LNCS, vol. 8172, pp. 365–380. Springer, Heidelberg (2013)
35. Heizmann, M., Hoenicke, J., Podelski, A.: Software model checking for people who love automata. In: Sharygina, N., Veith, H. (eds.) CAV 2013. LNCS, vol. 8044, pp. 36–52. Springer, Heidelberg (2013)
36. Heizmann, M., Hoenicke, J., Podelski, A.: Termination analysis by learning terminating programs. In: Biere, A., Bloem, R. (eds.) CAV 2014. LNCS, vol. 8559, pp. 797–813. Springer, Heidelberg (2014)
37. Henzinger, T.A., Jhala, R., Majumdar, R., Sutre, G.: Lazy abstraction. In: POPL, pp. 58–70, ACM (2002)
38. Holzmann, G.J.: The SPIN Model Checker Primer and Reference Manual, vol. 1003. Addison-Wesley, Reading (2004)
39. Howar, F., Isberner, M., Merten, M., Steffen, B., Beyer, D.: The RERS grey-box challenge 2012: analysis of event-condition-action systems. In: Margaria, T., Steffen, B. (eds.) ISoLA 2012, Part I. LNCS, vol. 7609, pp. 608–614. Springer, Heidelberg (2012)
40. Kroening, D., Sharygina, N., Tsitovich, A., Wintersteiger, C.M.: Termination analysis with compositional transition invariants. In: Touili, T., Cook, B., Jackson, P. (eds.) CAV 2010. LNCS, vol. 6174, pp. 89–103. Springer, Heidelberg (2010)

41. Kroening, D., Tautschnig, M.: CBMC – C bounded model checker. In: Ábrahám, E., Havelund, K. (eds.) TACAS 2014 (ETAPS). LNCS, vol. 8413, pp. 389–391. Springer, Heidelberg (2014)
42. Lal, A., Qadeer, S.: Reachability modulo theories. RP **2013**, 23–44 (2013)
43. Leike, J., Heizmann, M.: Ranking templates for linear loops. In: Ábrahám, E., Havelund, K. (eds.) TACAS 2014 (ETAPS). LNCS, vol. 8413, pp. 172–186. Springer, Heidelberg (2014)
44. McMillan, K.: Cadence SMV. Cadence Berkeley Labs, CA (2000). <http://www.kennmcil.com/smv.html>
45. McMillan, K.L.: Lazy abstraction with interpolants. In: Ball, T., Jones, R.B. (eds.) CAV 2006. LNCS, vol. 4144, pp. 123–136. Springer, Heidelberg (2006)
46. McMillan, K.L.: Interpolants from Z3 proofs. In: FMCAD, pp. 19–27 (2011)
47. Ermis, E., Nutz, A., Dietsch, D., Hoenicke, J., Podelski, A.: Ultimate kojak. In: Ábrahám, E., Havelund, K. (eds.) TACAS 2014 (ETAPS). LNCS, vol. 8413, pp. 421–423. Springer, Heidelberg (2014)
48. Podelski, A., Rybalchenko, A.: A complete method for the synthesis of linear ranking functions. In: Steffen, B., Levi, G. (eds.) VMCAI 2004. LNCS, vol. 2937, pp. 239–251. Springer, Heidelberg (2004)
49. Podelski, A., Rybalchenko, A.: Transition invariants. In: LICS, pp. 32–41, IEEE Computer Society (2004)
50. Podelski, A., Rybalchenko, A.: Transition predicate abstraction and fair termination. In: POPL, pp. 132–144, ACM (2005)
51. Podelski, A., Rybalchenko, A.: ARMC: the logical choice for software model checking with abstraction refinement. In: Hanus, M. (ed.) PADL 2007. LNCS, vol. 4354, pp. 245–259. Springer, Heidelberg (2007)
52. Podelski, A., Rybalchenko, A., Wies, T.: Heap assumptions on demand. In: Gupta, A., Malik, S. (eds.) CAV 2008. LNCS, vol. 5123, pp. 314–327. Springer, Heidelberg (2008)
53. Post, A.C.: Effective Correctness Criteria for Real-time Requirements. Shaker, Aachen (2012)
54. Rozier, K.Y., Vardi, M.Y.: LTL satisfiability checking. In: Božnački, D., Edelkamp, S. (eds.) SPIN 2007. LNCS, vol. 4595, pp. 149–167. Springer, Heidelberg (2007)
55. Ströder, T., Aschermann, C., Frohn, F., Hensel, J., Giesl, J.: Aprove: termination and memory safety of C programs - (competition contribution). In: Baier, C., Tinelli, C. (eds.) TACAS 2015. LNCS, vol. 9035, pp. 417–419. Springer, Heidelberg (2015)
56. Urban, C.: FuncTion: an abstract domain functor for termination. In: Baier, C., Tinelli, C. (eds.) TACAS 2015. LNCS, vol. 9035, pp. 464–466. Springer, Heidelberg (2015)
57. Vardi, M.Y., Wolper, P.: An automata-theoretic approach to automatic program verification. In: LICS, pp. 322–331, IEEE Computer Society (1986)