Abstract Interpretation of CTL Properties

Caterina Urban, Samuel Ueltschi, and Peter Müller

Department of Computer Science
ETH Zurich, Switzerland

Abstract. CTL is a temporal logic commonly used to express program properties. Most of the existing approaches for proving CTL properties only support certain classes of programs, limit their scope to a subset of CTL, or do not directly support certain existential CTL formulas. This paper presents an abstract interpretation framework for proving CTL properties that does not suffer from these limitations. Our approach automatically infers sufficient preconditions, and thus provides useful information even when a program satisfies a property only for some inputs.
We systematically derive a program semantics that precisely captures CTL properties by abstraction of the operational trace semantics of a program. We then leverage existing abstract domains based on piecewise-defined functions to derive decidable abstractions that are suitable for static program analysis. To handle existential CTL properties, we augment these abstract domains with under-approximating operators.
We implemented our approach in a prototype static analyzer. Our experimental evaluation demonstrates that the analysis is effective, even for CTL formulas with non-trivial nesting of universal and existential path quantifiers, and performs well on a wide variety of benchmarks.

1 Introduction

Computation tree logic (CTL) [6] is a temporal logic introduced by Clarke and Emerson to overcome certain limitations of linear temporal logic (LTL) [33] for program specification purposes. Most of the existing approaches for proving program properties expressed in CTL have limitations that restrict their applicability: they are limited to finite-state programs [7] or to certain classes of infinite-state programs (e.g., pushdown systems [36]), they limit their scope to a subset of CTL (e.g., the universal fragment of CTL [11]), or support existential path quantifiers only indirectly by considering their universal dual [8].

In this paper, we propose a new static analysis method for proving CTL properties that does not suffer from any of these limitations. We set our work in the framework of *abstract interpretation* [16], a general theory of semantic approximation that provides a basis for various successful industrial-scale tools (e.g., Astrée [3]). We generalize an existing abstract interpretation framework for proving termination [18] and other liveness properties [41].

Following the theory of abstract interpretation [14], we abstract away from irrelevant details about the execution of a program and systematically derive a program semantics that is *sound and complete* for proving a CTL property.

```
while [1]( rand() ) {
  [2]x := 1
  [3]n := rand()
  while [4]( n > 0 ) { [5]n := n − 1 }
  [6]x := 0
}
while [7]( true ) {}[8]
```

Fig. 1: Standard lock acquire/release-style program [12], where rand() is a random number generation function. Assignments $x := 1$ and $x := 0$ are acting as acquire and release, respectively. We want to prove the CTL property $\mathsf{AG}(x = 1 \Rightarrow \mathsf{A}(\mathit{true}\ \mathsf{U}\ x = 0))$ expressing that whenever a lock is acquired ($x = 1$) it is eventually released ($x = 0$). We assume that initially $x = 0$.

The semantics is a function defined over the programs states that satisfy the CTL formula. The value of the semantics for a CTL formula that expresses a liveness property (e.g., $\mathsf{A}(\mathit{true}\ \mathsf{U}\ \phi)$) gives an upper bound on the number of program execution steps needed to reach a desirable state (i.e., a state satisfying ϕ for $\mathsf{A}(\mathit{true}\ \mathsf{U}\ \phi)$). The semantics for any other CTL formula is the constant function equal to zero over its domain. We define the semantics inductively on the structure of a CTL formula, and we express it in a constructive fixpoint form starting from the functions defined for its sub-formulas.

Further sound abstractions suitable for static program analysis are derived by *fixpoint approximation* [14]. We leverage existing numerical abstract domains based on piecewise-defined functions [39], which we augment with novel under-approximating operators to directly handle existential CTL formulas. The piecewise-defined function for a CTL formula is automatically inferred through *backward analysis* by building upon the piecewise-defined functions for its sub-formulas. It over-approximates the value of the corresponding concrete semantics and, by under-approximating its domain of definition, yields a *sufficient precondition* for the CTL property. We prove the soundness of the analysis, meaning that all program executions respecting the inferred precondition indeed satisfy the CTL property. A program execution that does not respect the precondition might or might not satisfy the property.

To briefly illustrate our approach, let us consider the acquire/release-style program shown in Figure 1, and the CTL formula $\mathsf{AG}(x = 1 \Rightarrow \mathsf{A}(\mathit{true}\ \mathsf{U}\ x = 0))$. The analysis begins from the atomic propositions $x = 1$ and $x = 0$ and, for each program control point, it infers a piecewise-defined function that is only defined when x is one or zero, respectively. It then continues to the sub-formula $\mathsf{A}(\mathit{true}\ \mathsf{U}\ x = 0)$ for which, building upon the function obtained for $x = 0$, it infers the following interesting function at program point 4:

$$\lambda x.\lambda n. \begin{cases} 0 & x = 0 \\ 2 & x \neq 0 \wedge n \leq 0 \\ 2n + 2 & \text{otherwise} \end{cases} \tag{1.1}$$

The function indicates that the sub-formula $x = 0$ is either satisfied trivially (when x is already zero), or in at most 2 program execution steps when $n \leq 0$ (and thus the loop at program point 4 is not entered) and $2n + 2$ steps when $n > 0$ (and thus the loop is entered). The analysis then proceeds to $x = 1 \Rightarrow \mathsf{A}(\mathit{true}\ \mathsf{U}\ x = 0)$, i.e., $x \neq 1 \vee \mathsf{A}(\mathit{true}\ \mathsf{U}\ x = 0)$. The inferred function for the sub-formula $x \neq 1$ is only defined over the complement of the domain of the one obtained for $x = 1$. The disjunction combines this function with the one obtained for $\mathsf{A}(\mathit{true}\ \mathsf{U}\ x = 0)$ by taking the union over the function domains and the maximum over the function values. The result at program point 4 is the same function obtained for $\mathsf{A}(\mathit{true}\ \mathsf{U}\ x = 0)$. Finally, the analysis can proceed to the initial formula $\mathsf{AG}(x = 1 \Rightarrow \mathsf{A}(\mathit{true}\ \mathsf{U}\ x = 0))$. The function at program point 4 remains the same but its value now indicates the maximum number of steps needed until the *next state* that satisfies $x = 0$. The function inferred at the beginning of the program proves that the program satisfies the CTL formula $\mathsf{AG}(x = 1 \Rightarrow \mathsf{A}(\mathit{true}\ \mathsf{U}\ x = 0))$ unless x has initial value one. Indeed, in such a case, the program does not satisfy the formula since the loop at program point 1 might never execute. Thus, the inferred precondition is the weakest precondition for the CTL property $\mathsf{AG}(x = 1 \Rightarrow \mathsf{A}(\mathit{true}\ \mathsf{U}\ x = 0))$.

We implemented our approach in the prototype static analyzer FuncTion [13]. Our experimental evaluation demonstrates that the analysis is effective, even for CTL formulas with non-trivial nesting of universal and existential path quantifiers, and performs well on a wide variety of benchmarks.

2 Trace Semantics

We model the operational semantics of a program as a *transition system* $\langle \Sigma, \tau \rangle$ where Σ is a (potentially infinite) set of program states, and the transition relation $\tau \subseteq \Sigma \times \Sigma$ describes the possible transitions between states. The set of *final states* of the program is $\Omega \stackrel{\text{def}}{=} \{s \in \Sigma \mid \forall s' \in \Sigma \colon \langle s, s' \rangle \notin \tau\}$.

Given a transition system $\langle \Sigma, \tau \rangle$, the function $\mathrm{pre} \colon \mathcal{P}(\Sigma) \to \mathcal{P}(\Sigma)$ maps a given set of states X to the set of their predecessors with respect to τ: $\mathrm{pre}(X) \stackrel{\text{def}}{=} \{s \in \Sigma \mid \exists s' \in X : \langle s, s' \rangle \in \tau\}$, and the function $\widetilde{\mathrm{pre}} \colon \mathcal{P}(\Sigma) \to \mathcal{P}(\Sigma)$ maps a given set of states X to the set of states whose successors with respect to τ are all in X: $\widetilde{\mathrm{pre}}(X) \stackrel{\text{def}}{=} \{s \in \Sigma \mid \forall s' \in \Sigma : \langle s, s' \rangle \in \tau \Rightarrow s' \in X\}$.

In the following, given a set S, let $S^n \stackrel{\text{def}}{=} \{s_0 \cdots s_{n-1} \mid \forall i < n : s_i \in S\}$ be the set of all sequences of exactly n elements from S. We write ε to denote the empty sequence, i.e., $S^0 \stackrel{\text{def}}{=} \{\varepsilon\}$. Let $S^* \stackrel{\text{def}}{=} \bigcup_{n \in \mathbb{N}} S^n$ be the set of all finite sequences, $S^+ \stackrel{\text{def}}{=} S^* \setminus S^0$ be the set of all non-empty finite sequences, S^ω be the set of all infinite sequences, $S^{+\infty} \stackrel{\text{def}}{=} S^+ \cup S^\omega$ be the set of all non-empty finite or infinite sequences and $S^{*\infty} \stackrel{\text{def}}{=} S^* \cup S^\omega$ be the set of all finite or infinite sequences of elements from S. We write $\sigma\sigma'$ for the concatenation of two sequences $\sigma, \sigma' \in S^{*\infty}$ (with $\sigma\varepsilon = \varepsilon\sigma = \sigma$, and $\sigma\sigma' = \sigma$ if $\sigma \in S^\omega$), $T^+ \stackrel{\text{def}}{=} T \cap S^+$ for the selection of the non-empty finite sequences of $T \subseteq S^{*\infty}$, $T^\omega \stackrel{\text{def}}{=} T \cap S^\omega$ for the selection of the infinite sequences of $T \subseteq S^{*\infty}$, and $T\,;T' \stackrel{\text{def}}{=} \{\sigma s \sigma' \mid s \in S, \sigma s \in T, s\sigma' \in T'\}$

for the merging of sets of sequences $T \subseteq S^+$ and $T' \subseteq S^{+\infty}$, when a finite sequence in T terminates with the initial state of a sequence in T'.

Given a transition system $\langle \Sigma, \tau \rangle$, a *trace* is a non-empty sequence of program states described by the transition relation τ, that is, $\langle s, s' \rangle \in \tau$ for each pair of consecutive states $s, s' \in \Sigma$ in the sequence. The set of final states Ω and the transition relation τ can be understood as sets of traces of length one and two, respectively. The *maximal trace semantics* $\Lambda \in \mathcal{P}(\Sigma^{+\infty})$ generated by a transition system is the union of all non-empty finite traces that are terminating with a final state in Ω, and all infinite traces. It can be expressed as a least fixpoint in the complete lattice $\langle \mathcal{P}(\Sigma^{+\infty}), \sqsubseteq, \sqcup, \sqcap, \Sigma^\omega, \Sigma^+ \rangle$ [14]:

$$\Lambda = \mathrm{lfp}^{\sqsubseteq}\ \lambda T. \Omega \cup (\tau\ ;\ T) \tag{2.1}$$

where the computational order is $T_1 \sqsubseteq T_2 \stackrel{\text{def}}{=} T_1^+ \subseteq T_2^+ \wedge T_1^\omega \supseteq T_2^\omega$.

The maximal trace semantics carries all information about a program and fully describes its behavior. However, reasoning about a particular property of a program is facilitated by the design of a semantics that abstracts away from irrelevant details about program executions. In the paper, we use *abstract interpretation* [16] to systematically derive, by abstraction of the maximal trace semantics, a sound and complete semantics that precisely captures exactly and only the needed information to reason about CTL properties.

3 Computation Tree Logic

CTL is also known as *branching* temporal logic; its semantics is based on a branching notion of time: at each moment there may be several possible successor program states and thus each moment of time might have several different possible futures. Accordingly, the interpretation of CTL formulas is defined in terms of program states, as opposed to the interpretation of LTL formulas in terms of traces. This section gives a brief introduction into the syntax and semantics of CTL. We refer to [1] for further details.

We assume a set of atomic propositions describing properties of program states. Formulas in CTL are formed according to the following grammar:

$$\phi ::= a \mid \neg\phi \mid \phi \wedge \phi \mid \phi \vee \phi \mid \mathsf{AX}\phi \mid \mathsf{AG}\phi \mid \mathsf{A}(\phi\ \mathsf{U}\ \phi) \mid \mathsf{EX}\phi \mid \mathsf{EG}\phi \mid \mathsf{E}(\phi\ \mathsf{U}\ \phi)$$

where a is an atomic proposition. The universal quantifier (denoted A) and the existential quantifier (denoted E) allow expressing properties of *all* or *some* traces that start in a state. In the following, we often use Q to mean either A or E. The *next* temporal operator (denoted X) allows expressing properties about the next program state in a trace. The *globally* operator (denoted G) allow expressing properties that should hold always (i.e., for all states) on a trace. The *until* temporal operator (denoted U) allows expressing properties that should hold eventually on a trace, and always until then. We omit the *finally* temporal operator (denoted F) since a formula of the form $\mathsf{QF}\phi$ can be equivalently expressed as $\mathsf{Q}(\mathit{true}\ \mathsf{U}\ \phi)$.

The semantics of formulas in CTL is formally given by a satisfaction relation $\models$ between program states and CTL formulas. In the following, we write $s \models \phi$ if and only if the formula ϕ holds in the program state $s \in \Sigma$. We assume that the satisfaction relation for atomic propositions is given. The satisfaction relation for other CTL formulas is formally defined as follows:

$$\begin{array}{lcl} s \models \neg\phi & \Leftrightarrow & \neg(s \models \phi) \\ s \models \phi_1 \wedge \phi_2 & \Leftrightarrow & s \models \phi_1 \wedge s \models \phi_2 \\ s \models \phi_1 \vee \phi_2 & \Leftrightarrow & s \models \phi_1 \vee s \models \phi_2 \\ s \models \mathsf{A}\varphi & \Leftrightarrow & \forall \sigma \in T(s) \colon \sigma \models \varphi \\ s \models \mathsf{E}\varphi & \Leftrightarrow & \exists \sigma \in T(s) \colon \sigma \models \varphi \end{array} \tag{3.1}$$

where $T(s) \in \mathcal{P}(\Sigma^{+\infty})$ denotes the set of all program traces starting in the state $s \in \Sigma$. The semantics of trace formulas φ is defined below:

$$\begin{array}{lcl} \sigma \models \mathsf{X}\phi & \Leftrightarrow & \sigma[1] \models \phi \\ \sigma \models \mathsf{G}\phi & \Leftrightarrow & \forall 0 \leq i \colon \sigma[i] \models \phi \\ \sigma \models \phi_1 \ \mathsf{U} \ \phi_2 & \Leftrightarrow & \exists 0 \leq i \colon \sigma[i] \models \phi_2 \wedge \forall 0 \leq j < i \colon \sigma[j] \models \phi_1 \end{array} \tag{3.2}$$

where $\sigma[i]$ denotes the program state at position i on the trace $\sigma \in \Sigma^{+\infty}$. We refer to [1] for further details.

4 Program Semantics for CTL Properties

In the following, we derive a program semantics that is *sound and complete* for proving a CTL property. We define the semantics inductively on the structure of a CTL formula. More specifically, for each formula ϕ, we define the *CTL abstraction* $\alpha_\phi \colon \mathcal{P}(\Sigma^{+\infty}) \to (\Sigma \rightharpoonup \mathbb{O})$ which extracts a partial function $f \colon \Sigma \rightharpoonup \mathbb{O}$ from program states to ordinals from a given set of sequences $T \in \mathcal{P}(\Sigma^{+\infty})$ by building upon the CTL abstractions of the sub-formulas of ϕ. The domain of f coincides with the set of program states that satisfy ϕ. Ordinal values are needed to support programs with possibly unbounded non-determinism [18]. The definition of α_ϕ for each CTL formula is summarized in Figure 2 and explained in more detail below. We use the CTL abstraction to define the program semantics $\Lambda_\phi \colon \Sigma \rightharpoonup \mathbb{O}$ for a formula ϕ by abstraction of the maximal trace semantics Λ.

Definition 1. *Given a CTL formula ϕ and the corresponding CTL abstraction $\alpha_\phi \colon \mathcal{P}(\Sigma^{+\infty}) \to (\Sigma \rightharpoonup \mathbb{O})$, the program semantics $\Lambda_\phi \colon \Sigma \rightharpoonup \mathbb{O}$ for ϕ is defined as $\Lambda_\phi \stackrel{def}{=} \alpha_\phi(\Lambda)$, where Λ is the maximal trace semantics (cf. Equation 2.1).*

Remarks. It may seem unintuitive to define Λ_ϕ starting from program traces rather than program states (as in Section 3). The reason behind this deliberate choice is that it allows placing Λ_ϕ in the hierarchy of semantics defined by Cousot [14], which is a uniform framework that makes program semantics easily comparable and facilitates explaining the similarities and correspondences between semantic models. Specifically, this enables the comparison with existing semantics for termination [18] and other liveness properties [41] (cf. Section 7).

ϕ	$\alpha_\phi \colon \mathcal{P}\left(\Sigma^{+\infty}\right) \to (\Sigma \rightharpoonup \mathbb{O})$
a	$\alpha_a(T) \stackrel{\text{def}}{=} \lambda s \in \mathrm{st}(T).\ \begin{cases} 0 & s \models a \\ \text{undefined} & \text{otherwise} \end{cases}$
$\mathsf{Q}\mathsf{X}\phi$	$\alpha_{\mathsf{Q}\mathsf{X}\phi}(T) \stackrel{\text{def}}{=} \lambda s \in \mathrm{st}(T).\ \begin{cases} 0 & s \in \mathrm{trans}_{\mathsf{Q}}(\mathrm{dom}(\alpha_\phi(T))) \\ \text{undefined} & \text{otherwise} \end{cases}$
$\mathsf{Q}(\phi_1\ \mathsf{U}\ \phi_2)$	$\alpha_{\mathsf{Q}(\phi_1 \mathsf{U} \phi_2)}(T) \stackrel{\text{def}}{=} \alpha^{\mathrm{rk}}_{\mathsf{Q}}(\alpha^{\mathrm{sq}}_{\mathsf{Q}(\phi_1 \mathsf{U} \phi_2)}(T))$
$\mathsf{Q}\mathsf{G}\phi$	$\alpha_{\mathsf{Q}\mathsf{G}\phi}(T) \stackrel{\text{def}}{=} \mathrm{gfp}^{\sqsubseteq}_{\alpha_\phi(T)}\ \Theta_{\mathsf{Q}\mathsf{G}\phi}$ $\Theta_{\mathsf{Q}\mathsf{G}\phi}(f) \stackrel{\text{def}}{=} \lambda s.\ \begin{cases} f(s) & s \in \mathrm{dom}(f) \cap \mathrm{trans}_{\mathsf{Q}}(\mathrm{dom}(f)) \\ \text{undefined} & \text{otherwise} \end{cases}$
$\neg\phi$	$\alpha_{\neg\phi}(T) \stackrel{\text{def}}{=} \lambda s \in \mathrm{st}(T).\ \begin{cases} 0 & s \notin \mathrm{dom}(\alpha_\phi(T)) \\ \text{undefined} & \text{otherwise} \end{cases}$
$\phi_1 \wedge \phi_2$	$\alpha_{\phi_1 \wedge \phi_2}(T) \stackrel{\text{def}}{=} \lambda s \in \mathrm{st}(T).\ \begin{cases} \sup\{f_1(s), f_2(s)\} & s \in \mathrm{dom}(f_1) \cap \mathrm{dom}(f_2) \\ \text{undefined} & \text{otherwise} \end{cases}$
$\phi_1 \vee \phi_2$	$\alpha_{\phi_1 \vee \phi_2}(T) \stackrel{\text{def}}{=} \lambda s \in \mathrm{st}(T).\ \begin{cases} \sup\{f_1(s), f_2(s)\} & s \in \mathrm{dom}(f_1) \cap \mathrm{dom}(f_2) \\ f_1(s) & s \in \mathrm{dom}(f_1) \setminus \mathrm{dom}(f_2) \\ f_2(s) & s \in \mathrm{dom}(f_2) \setminus \mathrm{dom}(f_1) \\ \text{undefined} & \text{otherwise} \end{cases}$

Fig. 2: CTL abstraction $\alpha_\phi \colon \mathcal{P}\left(\Sigma^{+\infty}\right) \to (\Sigma \rightharpoonup \mathbb{O})$ for each CTL formula ϕ. The function $\mathrm{trans}_{\mathsf{Q}}$ stands for pre, if Q is E, or $\widetilde{\mathrm{pre}}$, if Q is A (cf. Section 2). The *state function* $\mathrm{st} \colon \mathcal{P}\left(\Sigma^{+\infty}\right) \to \mathcal{P}\left(\Sigma\right)$ collects all states of a given set of sequences T: $\mathrm{st}(T) \stackrel{\text{def}}{=} \{s \in \Sigma \mid \exists \sigma' \in \Sigma^*, \sigma'' \in \Sigma^{*\infty} : \sigma' s \sigma'' \in T\}$. The *ranking abstraction* $\alpha^{\mathrm{rk}}_{\mathsf{Q}} \colon \mathcal{P}\left(\Sigma^{+}\right) \to (\Sigma \rightharpoonup \mathbb{O})$ is defined in Equation 4.1, while the *subsequence abstraction* $\alpha^{\mathrm{sq}}_{\mathsf{Q}\mathsf{F}\phi} \colon \mathcal{P}\left(\Sigma^{+\infty}\right) \to \mathcal{P}\left(\Sigma^{+}\right)$ is defined in Equations 4.2 and 4.3. In the last two rows, $f_1 \stackrel{\text{def}}{=} \alpha_{\phi_1}(T)$ and $f_2 \stackrel{\text{def}}{=} \alpha_{\phi_2}(T)$.

It may also seem unnecessary to define Λ_ϕ to be a function. However, this choice yields a uniform treatment of CTL formulas independently of whether they express safety or liveness properties (or a combination of these). Additionally, it allows leveraging existing abstract domains [38,39] (cf. Section 5) to obtain a sound static analysis for CTL properties. In particular, the proof of the soundness of the static analysis (cf. Theorem 2 and [38] for more details) requires reasoning both about the domain of Λ_ϕ as well as its value.

Atomic Propositions. For an atomic proposition a, the CTL abstraction $\alpha_a \colon \mathcal{P}\left(\Sigma^{+\infty}\right) \to (\Sigma \rightharpoonup \mathbb{O})$ simply extracts from a given set T of sequences a partial function that maps the states of the sequences in T (i.e., $s \in \mathrm{st}(T)$) that satisfy a (i.e., $s \models a$) to the constant value zero, meaning that no program execution steps are needed until a is satisfied for all executions starting in those states. Thus, the domain of the corresponding program semantics $\Lambda_a \colon \Sigma \rightharpoonup \mathbb{O}$ is (cf. Definition 1) is the set of program states that satisfy a (since $\mathrm{st}(\Lambda) = \Sigma$).

Next-Formulas. Next-formulas $\mathsf{Q}\mathsf{X}\phi$ express that the next state of all traces (if Q is A) or at least one trace (if Q is E) satisfies ϕ.

The CTL abstraction $\alpha_{\mathsf{Q}\mathsf{X}\phi}\colon \mathcal{P}(\Sigma^{+\infty}) \to (\Sigma \rightharpoonup \mathbb{O})$ for a next-formula $\mathsf{Q}\mathsf{X}\phi$ (cf. Figure 2) maps a set T of sequences to a partial function defined over the states of the sequences in T (i.e., $s \in \mathrm{st}(T)$) that are the *predecessors* of the states that satisfy ϕ, that is, the predecessors of the states in the domain of the CTL abstraction for ϕ (i.e., $s \in \mathrm{trans}_{\mathsf{Q}}(\mathrm{dom}(\alpha_\phi(T)))$). The function has constant value zero over its domain, again meaning that no program execution steps are needed until $\mathsf{Q}\mathsf{X}\phi$ is satisfied for all executions starting in those states.

Thus, the domain of the program semantics $\Lambda_{\mathsf{Q}\mathsf{X}\phi}\colon \Sigma \rightharpoonup \mathbb{O}$ is the set of states inevitably (for $\Lambda_{\mathsf{A}\mathsf{X}\phi}$) or possibly (for $\Lambda_{\mathsf{E}\mathsf{X}\phi}$) leading to a state in the domain $\mathrm{dom}(\Lambda_\phi)$ of the program semantics of the sub-formula ϕ (cf. Definition 1).

Until-Formulas. Until-formulas $\mathsf{Q}(\phi_1\ \mathsf{U}\ \phi_2)$ express that some desired state (i.e., a state satisfying the sub-formula ϕ_2) is eventually reached during program execution, either in all traces (if Q is A) or in at least one trace (if Q is E), and the sub-formula ϕ_1 is satisfied in all program states encountered until then. Thus, we can observe that an until-formula is satisfied by *finite* subsequences of possibly *infinite* program traces. To reason about subsequences, we define the *subsequence function* $\mathrm{sq}\colon \mathcal{P}(\Sigma^{+\infty}) \to \mathcal{P}(\Sigma^{+})$ which extracts all finite subsequences of a given set of sequences T: $\mathrm{sq}(T) \stackrel{\text{def}}{=} \{\sigma \in \Sigma^{+} \mid \exists \sigma' \in \Sigma^{*}, \sigma'' \in \Sigma^{*\infty} : \sigma'\sigma\sigma'' \in T\}$. In the following, given a formula $\mathsf{Q}(\phi_1\ \mathsf{U}\ \phi_2)$, we define the corresponding *subsequence abstraction* $\alpha^{\mathrm{sq}}_{\mathsf{Q}(\phi_1\mathsf{U}\phi_2)}\colon \mathcal{P}(\Sigma^{+\infty}) \to \mathcal{P}(\Sigma^{+})$ which extracts the finite subsequences that satisfy $\mathsf{Q}(\phi_1\ \mathsf{U}\ \phi_2)$ from of a set of sequences T. We can then use $\alpha^{\mathrm{sq}}_{\mathsf{Q}(\phi_1\mathsf{U}\phi_2)}$ to define the CTL abstraction $\alpha_{\mathsf{Q}(\phi_1\mathsf{U}\phi_2)}\colon \mathcal{P}(\Sigma^{+\infty}) \to (\Sigma \rightharpoonup \mathbb{O})$ as shown in Figure 2. The *ranking abstraction* $\alpha^{\mathrm{rk}}_{\mathsf{Q}}\colon \mathcal{P}(\Sigma^{+}) \to (\Sigma \rightharpoonup \mathbb{O})$ is:

$$\alpha^{\mathrm{rk}}_{\mathsf{Q}}(T) \stackrel{\text{def}}{=} \alpha^{\mathrm{v}}_{\mathsf{Q}}(\vec{\alpha}\,(T)) \tag{4.1}$$

where $\vec{\alpha}\colon \mathcal{P}(\Sigma^{+}) \to \mathcal{P}(\Sigma) \times \mathcal{P}(\Sigma \times \Sigma)$ extracts from a given set of non-empty finite sequences T the smallest transition system $\langle S, r\rangle$ that generates T: $\vec{\alpha}\,(T) \stackrel{\text{def}}{=} \langle \mathrm{st}(T), \{\langle s, s'\rangle \in \Sigma \times \Sigma \mid \exists \sigma \in \Sigma^{*}, \sigma' \in \Sigma^{*\infty} : \sigma s s' \sigma' \in T\}\rangle$ and the function $\alpha^{\mathrm{v}}_{\mathsf{Q}}\colon \mathcal{P}(\Sigma) \times \mathcal{P}(\Sigma \times \Sigma) \to (\Sigma \rightharpoonup \mathbb{O})$ provides the rank of the elements in the domain of the transition relation of the transition system:

$$\alpha^{\mathrm{v}}_{\mathsf{Q}}\langle S, r\rangle s \stackrel{\text{def}}{=} \begin{cases} 0 & \forall s' \in S : \langle s, s'\rangle \notin r \\ \mathrm{bnd}_{\mathsf{Q}} \left\{ \alpha^{\mathrm{v}}_{\mathsf{Q}}\langle S, r\rangle s' + 1 \,\middle|\, \begin{array}{l} s \neq s', \langle s, s'\rangle \in r, \\ s' \in \mathrm{dom}(\alpha^{\mathrm{v}}_{\mathsf{Q}}\langle S, r\rangle) \end{array} \right\} & \text{otherwise} \end{cases}$$

where $\mathrm{bnd}_{\mathsf{Q}}$ stands for sup, if Q is A, or inf, if Q is E. The CTL abstraction $\alpha_{\mathsf{A}(\phi_1\mathsf{U}\phi_2)}$ (resp. $\alpha_{\mathsf{E}(\phi_1\mathsf{U}\phi_2)}$) maps the states $\mathrm{st}(T)$ of a given set of sequences T that satisfy $\mathsf{Q}(\phi_1\ \mathsf{U}\ \phi_2)$ to an upper bound (resp. lower bound) on the number of program execution steps until the sub-formula ϕ_2 is satisfied, for all (resp. at least one of the) executions starting in those states.

Existential Until-Formulas. The subsequence abstraction $\alpha^{\text{sq}}_{\mathsf{E}(\phi_1 \mathsf{U} \phi_2)}$ for a formula $\mathsf{E}(\phi_1 \mathsf{U} \phi_2)$ extracts from a given a set of sequences T the finite subsequence of states that terminate in a state satisfying ϕ_2 and all predecessor states satisfy ϕ_1 (and not ϕ_2). It is defined as follows:

$$\begin{aligned} \alpha^{\text{sq}}_{\mathsf{E}(\phi_1 \mathsf{U} \phi_2)}(T) &\stackrel{\text{def}}{=} \overline{\alpha}_{\mathsf{E}(\phi_1 \mathsf{U} \phi_2)}[\text{dom}(\alpha_{\phi_1}(T))][\text{dom}(\alpha_{\phi_2}(T))]T \\ \overline{\alpha}_{\mathsf{E}(\phi_1 \mathsf{U} \phi_2)}[S_1][S_2]T &\stackrel{\text{def}}{=} \{\sigma s \in \text{sq}(T) \mid \sigma \in (S_1 \setminus S_2)^*, s \in S_2\} \end{aligned} \tag{4.2}$$

where S_1 is the set of states that satisfy the sub-formula ϕ_1 (i.e., $\text{dom}(\alpha_{\phi_1}(T))$), and S_2 is the set of desired states (i.e, $\text{dom}(\alpha_{\phi_2}(T))$).

Universal Until-Formulas. A finite subsequence of states satisfies a universal until-formula $\mathsf{A}(\phi_1 \mathsf{U} \phi_2)$ if and only if it terminates in a state satisfying ϕ_2, all predecessor states satisfy ϕ_1, *and all other sequences with a common prefix also terminate in a state satisfying* ϕ_2 *(and all its predecessors satisfy* ϕ_1*)*, i.e., the program reaches a desired state (via states that satisfy ϕ_1) independently of the non-deterministic choices made during execution. We define the *neighborhood* of a sequence of states σ in a given set T as the set of sequences $\sigma' \in T$ with a common prefix with σ: $\text{nbhd}(\sigma, T) \stackrel{\text{def}}{=} \{\sigma' \in T \mid \text{pf}(\sigma) \cap \text{pf}(\sigma') \neq \emptyset\}$, where the *prefixes function* $\text{pf} \colon \Sigma^{+\infty} \to \mathcal{P}(\Sigma^{+\infty})$ yields the set of non-empty prefixes of a sequence $\sigma \in \Sigma^{+\infty}$: $\text{pf}(\sigma) \stackrel{\text{def}}{=} \{\sigma' \in \Sigma^{+\infty} \mid \exists \sigma'' \in \Sigma^{*\infty} \colon \sigma = \sigma'\sigma''\}$.

We can now defined the subsequence abstraction $\alpha^{\text{sq}}_{\mathsf{A}(\phi_1 \mathsf{U} \phi_2)}$:

$$\begin{aligned} \alpha^{\text{sq}}_{\mathsf{A}(\phi_1 \mathsf{U} \phi_2)}(T) &\stackrel{\text{def}}{=} \overline{\alpha}_{\mathsf{A}(\phi_1 \mathsf{U} \phi_2)}[\text{dom}(\alpha_{\phi_1}(T))][\text{dom}(\alpha_{\phi_2}(T))]T \\ \overline{\alpha}_{\mathsf{A}(\phi_1 \mathsf{U} \phi_2)}[S_1][S_2]T &\stackrel{\text{def}}{=} \left\{ \sigma s \in \text{sq}(T) \;\middle|\; \begin{array}{c} \sigma \in (S_1 \setminus S_2)^*, s \in S_2, \\ \text{nbhd}(\sigma, \text{sf}(T) \cap \overline{S_2}^{+\infty}) = \emptyset, \\ \text{nbhd}(\sigma, \text{sf}(T) \cap Z) = \emptyset \end{array} \right\} \end{aligned} \tag{4.3}$$

where the *suffixes function* $\text{sf} \colon \mathcal{P}(\Sigma^{+\infty}) \to \mathcal{P}(\Sigma^{+\infty})$ yields the set of non-empty suffixes of a set of sequences T: $\text{sf}(T) \stackrel{\text{def}}{=} \bigcup \{\sigma \in \Sigma^{+\infty} \mid \exists \sigma' \in \Sigma^* : \sigma'\sigma \in T\}$, and $Z \stackrel{\text{def}}{=} \left\{\sigma s \sigma' \in \Sigma^{+\infty} \mid \sigma \in \Sigma^* \wedge s \in \overline{S_1 \cup S_2} \wedge \sigma' \in \Sigma^{+\infty}\right\}$ is the set of sequences of states in which at least one state satisfies neither ϕ_1 nor ϕ_2. The last two conjuncts in the definition of the helper function $\overline{\alpha}_{\mathsf{A}(\phi_1 \mathsf{U} \phi_2)}[S_1][S_2]$ ensure that a finite subsequence satisfies $\mathsf{A}(\phi_1 \mathsf{U} \phi_2)$ only if it does not have a common prefix with any subsequence of T that never reaches a desired state in S_2 (i.e., $\text{nbhd}(\sigma, \text{sf}(T) \cap \overline{S_2}^{+\infty}) = \emptyset$) and with any subsequence that contains a state that does not belong to S_1 and S_2 (i.e, $\text{nbhd}(\sigma, \text{sf}(T) \cap Z) = \emptyset$).

Example 1. Let us consider again the acquire/release program of Figure 1 and let T be the set of its traces. The suffixes starting at program point **2** of the traces in T can be visualized as follows:

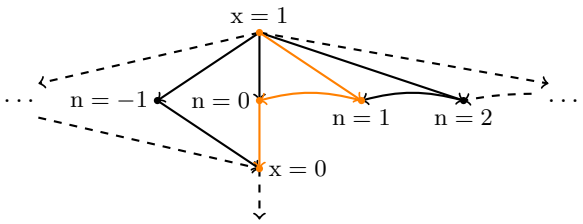

Observe that these sequences form a neighborhood in the set $\mathrm{sf}(T)$ of suffixes of T (i.e., the set of all these sequences is the neighborhood $\mathrm{nbhd}(\sigma, \mathrm{sf}(T))$ of any sequence σ in the set). In the following, we write x_i and n_i for the states denoted above by $\mathrm{x} = i$ and $\mathrm{n} = i$, respectively.

Let us consider the universal until-formula $\mathsf{A}(\mathrm{x} = 1\ \mathsf{U}\ \mathrm{x} = 0)$. The set of desired states is $S_2 = \{\mathrm{x}_0\}$ and $S_1 = \{\mathrm{x}_1\} \cup \{\mathrm{n}_i \mid i \in \mathbb{Z}\}$ is the set of states that satisfy $\mathrm{x} = 1$. All sequences in the neighborhood have prefixes of the form σs where $\sigma = \mathrm{x}_1 \cdots \in (S_1 \cap \overline{S_2})^*$ and $s = \mathrm{x}_0 \in S_2$. Thus, the neighborhood of any subsequence σs does not contain sequences in $\overline{S_2}^{+\infty}$ that never reach the desired state x_0 (i.e., $\mathrm{nbhd}(\sigma s, \mathrm{sf}(T) \cap \overline{S_2}^{+\infty}) = \emptyset$). Furthermore, the neighborhood does not contain sequences in Z in which at least one state neither satisfies $\mathrm{x} = 1$ nor $\mathrm{x} = 0$ (i.e., $\mathrm{nbhd}(\sigma, \mathrm{sf}(T) \cap Z) = \emptyset$). Therefore, the until-formula $\mathsf{A}(\mathrm{x} = 1\ \mathsf{U}\ \mathrm{x} = 0)$ is satisfied at program point 2.

Let us consider now the formula $\mathsf{A}(\mathrm{x} = 1 \wedge 0 \leq \mathrm{n}\ \mathsf{U}\ \mathrm{x} = 0)$. Again, all sequences in the neighborhood eventually reach the desired state x_0. However, in this case, the set S_1 is limited to states with non-negative values for n, i.e., $S_1 = \{\mathrm{x}_1\} \cup \{\mathrm{n}_i \mid 0 \leq i\}$. Thus, the neighborhood also contains sequences in which at least one state satisfies neither $\mathrm{x} = 1 \wedge 0 \leq \mathrm{n}$ nor $\mathrm{x} = 0$ (e.g., the sequence $\mathrm{x}_1\mathrm{n}_{-1}\ldots$). Hence $\mathsf{A}(\mathrm{x} = 1 \wedge 0 \leq \mathrm{n}\ \mathsf{U}\ \mathrm{x} = 0)$ is not satisfied at program point 2 since $\mathrm{nbhd}(\sigma, \mathrm{sf}(T) \cap Z) \neq \emptyset$. Instead, the existential until-formula $\mathsf{E}(\mathrm{x} = 1 \wedge 0 \leq \mathrm{n}\ \mathsf{U}\ \mathrm{x} = 0)$ is satisfied since, for instance, the subsequence σs where $\sigma = \mathrm{x}_1\mathrm{n}_1$ and $s = \mathrm{x}_0$ satisfies $(\mathrm{x} = 1 \wedge 0 \leq \mathrm{n}\ \mathsf{U}\ \mathrm{x} = 0)$. ■

Until Program Semantics. We now have all the ingredients that define the program semantics $\Lambda_{\mathsf{Q}(\phi_1 \mathsf{U} \phi_2)} \colon \Sigma \rightharpoonup \mathbb{O}$ for an until-formula $\mathsf{Q}(\phi_1\ \mathsf{U}\ \phi_2)$ (cf. Definition 1). Let $\langle \Sigma \rightharpoonup \mathbb{O}, \sqsubseteq \rangle$ be the partially ordered set for the computational order $f_1 \sqsubseteq f_2 \Leftrightarrow \mathrm{dom}(f_1) \subseteq \mathrm{dom}(f_2) \wedge \forall x \in \mathrm{dom}(f_1) : f_1(x) \leq f_2(x)$. The program semantics $\Lambda_{\mathsf{Q}(\phi_1 \mathsf{U} \phi_2)}$ can be expressed as a least fixpoint in $\langle \Sigma \rightharpoonup \mathbb{O}, \sqsubseteq \rangle$ as:

$$\Lambda_{\mathsf{Q}(\phi_1 \mathsf{U} \phi_2)} = \mathrm{lfp}^{\sqsubseteq}_{\dot{\emptyset}}\ \Theta_{\mathsf{Q}(\phi_1 \mathsf{U} \phi_2)}[\mathrm{dom}(\Lambda_{\phi_1})][\mathrm{dom}(\Lambda_{\phi_2})]$$

$$\Theta_{\mathsf{Q}(\phi_1 \mathsf{U} \phi_2)}[S_1][S_2]f \stackrel{\text{def}}{=} \lambda s. \begin{cases} 0 & s \in S_2 \\ \mathrm{bnd}_{\mathsf{Q}}\,\{f(s') + 1 \mid \langle s, s' \rangle \in \tau\} & s \in S_1 \wedge s \notin S_2 \wedge \\ & s \in \mathrm{trans}_{\mathsf{Q}}(\mathrm{dom}(f)) \\ \text{undefined} & \text{otherwise} \end{cases} \tag{4.4}$$

where $\dot{\emptyset}$ is the totally undefined function. The program semantics $\Lambda_{\mathsf{A}(\phi_1 \mathsf{U} \phi_2)}$ (resp. $\Lambda_{\mathsf{E}(\phi_1 \mathsf{U} \phi_2)}$) is a well-founded function mapping each program state in $\mathrm{dom}(\Lambda_{\phi_1})$

inevitably (resp. possibly) leading to a desirable state in $\mathrm{dom}(\Lambda_{\phi_2})$ to an ordinal, which represents an upper bound (resp. lower bound) on the number of program execution steps needed until a desirable state is reached.

Globally-Formulas. Globally-formulas $\mathsf{QG}\phi$ express that ϕ holds indefinitely in all traces (if Q is A) or at least one trace (if Q is E) starting in a state.

The definition of the CTL abstraction $\alpha_{\mathsf{QG}\phi}\colon \mathcal{P}(\Sigma^{+\infty}) \to (\Sigma \rightharpoonup \mathbb{O})$ for $\mathsf{QG}\phi$ given in Figure 2 retains the value of the CTL abstraction corresponding to the sub-formula ϕ. Intuitively, each iteration discards the states that satisfy ϕ (i.e., the states in $\mathrm{dom}(\alpha_\phi(T))$) but branch to (sub)sequences of T that do not satisfy $\mathsf{QG}\phi$. Preserving the value of α_ϕ provides more information than just mapping each state to the constant value zero. For instance, the CTL abstraction $\alpha_{\mathsf{AGAF}\phi}$ for a globally-formula $\mathsf{AGAF}\phi$ provides an upper bound on the number of program execution steps needed until the *next occurrence* of ϕ is satisfied, for all executions starting in the corresponding state.

The corresponding program semantics $\Lambda_{\mathsf{QG}\phi}\colon \Sigma \rightharpoonup \mathbb{O}$ (cf. Definition 1) preserves the value of Λ_ϕ for each state satisfying the sub-formula ϕ and inevitably (if Q is A) or possibly (if Q is E) leading only to other states that also satisfy ϕ.

Other Formulas. We are left with describing the CTL abstraction of $\neg\phi$, $\phi \wedge \phi$, and $\phi \vee \phi$ defined in Figure 2. For a negation $\neg\phi$, the CTL abstraction $\alpha_{\neg\phi}$ maps each program state that does not satisfy ϕ to the value zero. The CTL abstraction for a binary formula $\phi_1 \wedge \phi_2$ or $\phi_1 \vee \phi_2$ retains the largest value of the functions Λ_{ϕ_1} and Λ_{ϕ_2} for each program state satisfying both ϕ_1 and ϕ_2; for a disjunction $\phi_1 \vee \phi_2$, it also retains the value of the function for each program state satisfying either sub-formula.

Theorem 1. *A program satisfies a CTL formula ϕ for all traces starting from a given set of states $\mathcal{I}$ if and only if $\mathcal{I} \subseteq \mathrm{dom}(\Lambda_\phi)$.*

Proof. The proof proceeds by induction over the structure of the CTL formula ϕ. The base case are atomic propositions a for which the proof is immediate.

For a next-formulas $\mathsf{QX}\phi$, by induction hypothesis, $\mathrm{dom}(\Lambda_\phi)$ coincides with the set of states that satisfy ϕ. By Definition 1 and the definition of $\alpha_{\mathsf{QX}\phi}$ in Figure 2, the domain of $\Lambda_{\mathsf{QX}\phi}$ coincides with $\mathrm{trans}_{\mathsf{Q}}(\mathrm{dom}(\alpha_\phi(T)))$. Thus, by definition of $\mathrm{trans}_{\mathsf{Q}}$, we have that $\mathrm{dom}(\Lambda_{\mathsf{QX}\phi})$ coincides with the set of states that satisfy $\mathsf{QX}\phi$ (cf. Equations 3.1 and 3.2).

For an until-formula $\mathsf{Q}(\phi_1\ \mathsf{U}\ \phi_2)$, by induction hypothesis, $\mathrm{dom}(\Lambda_{\phi_1})$ and $\mathrm{dom}(\Lambda_{\phi_2})$ coincide with the set of states that satisfy ϕ_1 and ϕ_2, respectively. We have $\Lambda_{\mathsf{Q}}(\phi_1\ \mathsf{U}\ \phi_2) = \Theta_{\mathsf{Q}(\phi_1\mathsf{U}\phi_2)}[\mathrm{dom}(\Lambda_{\phi_1})][\mathrm{dom}(\Lambda_{\phi_2})](\Lambda_{\mathsf{Q}}(\phi_1\ \mathsf{U}\ \phi_2))$ from Equation 4.4. Therefore, by definition of $\Theta_{\mathsf{Q}(\phi_1\mathsf{U}\phi_2)}$, $\mathrm{dom}(\Lambda_{\mathsf{Q}}(\phi_1\ \mathsf{U}\ \phi_2))$ coincides with the states that satisfy ϕ_2 and all states that satisfy ϕ_1 and inevitably (if Q is A) or possibly (if Q is E) lead to states that satisfy ϕ_2. So $\mathrm{dom}(\Lambda_{\mathsf{Q}}(\phi_1\ \mathsf{U}\ \phi_2))$ coincides with the states that satisfy $\mathsf{Q}(\phi_1\ \mathsf{U}\ \phi_2)$ (cf. Equations 3.1 and 3.2).

For a globally-formula $\mathsf{QG}\phi$, by induction hypothesis, $\mathrm{dom}(\Lambda_\phi)$ coincides with the set of states that satisfy ϕ. By Definition 1 and the definition of $\alpha_{\mathsf{QG}\phi}$ in

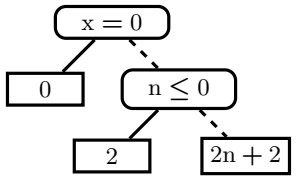

Fig. 3: Simplified decision tree representation of the piecewise-defined function inferred at program point 4 of the program of Figure 1 (cf. Equation 1.1). Each constraint is satisfied by the left subtree of the decision node, while the right subtree satisfies its negation. The leaves represent partial functions whose domain is determined by the constraints satisfied along the path to the leaf.

Figure 2, we have that $\Lambda_{\mathsf{QG}\phi} = \Theta_{\mathsf{QG}\phi}(\Lambda_{\mathsf{QG}\phi})$. Therefore, by definition of $\Theta_{\mathsf{QG}\phi}$, we have that $\mathrm{dom}(\Lambda_{\mathsf{QG}\phi})$ coincides with the states that satisfy ϕ inevitably (if Q is A) or possibly (if Q is E) lead to other states that satisfy ϕ. So $\mathrm{dom}(\Lambda_{\mathsf{QG}\phi})$ coincides with the states that satisfy $\mathsf{QG}\phi$ (cf. Equations 3.1 and 3.2).

Finally, all other cases ($\neg\phi$, $\phi_1 \wedge \phi_2$, and $\phi_1 \vee \phi_2$) follow immediately from the induction hypothesis, the semantics of the CTL formulas (cf. Equation 3.1) and the definition of the corresponding program semantics (cf. Definition 1 and the CTL abstractions in Figure 2). □

The program semantics for a CTL formula is not computable when the program state space is infinite. In the next section, we present decidable abstractions by means of piecewise-defined functions [38,39].

5 Static Analysis for CTL Properties

We recall here the features of the abstract domain of piecewise-defined functions [39] that are relevant for our purposes, and describe the new elements that we need to introduce to obtain a static analysis for proving CTL properties. We refer to [38] for an exhaustive presentation of the original abstract domain.

For illustration, we model a program using a control flow graph $\langle \mathcal{L}, E \rangle$, where $\mathcal{L}$ is the set of program points and $E \subseteq \mathcal{L} \times A \times \mathcal{L}$ is the set of edges in the control flow graph. Each edge is labeled by an action $s \in A$; possible actions are skip, a boolean condition b, or an assignment $x := e$. In the following, we write $out(l)$ to denote the set of outgoing edges from a program point l.

Piecewise-Defined Functions Abstract Domain. An element $t \in \mathcal{T}$ of the abstract domain is a piecewise-defined partial function represented by a *decision tree*, where the decision nodes are labeled by linear constraints over the program variables, and the leaf nodes are labeled by functions of the program variables. The decision nodes recursively partition the space of possible values

of the program variables, and the leaf nodes represent the value of the function corresponding to each partition. An example of (simplified) decision tree representation of a piecewise-defined function is shown in Figure 3.

Specifically, the decision nodes are labeled by linear constraints supported by an existing underlying numerical domain, i.e., interval [15] constraints (of the form $\pm \mathrm{x} \leq \mathrm{c}$), octagonal [30] constraints (of the form $\pm \mathrm{x}_i \pm \mathrm{x}_j \leq c$), or polyhedral [19] constraints (of the form $\mathrm{c}_1 \cdot \mathrm{x}_1 + \cdots + \mathrm{c}_k \cdot \mathrm{x}_k \leq \mathrm{c}_{k+1}$). The leaf nodes are labeled by *affine functions* of the program variables (of the form $\mathrm{m}_1 \cdot \mathrm{x}_1 + \cdots + \mathrm{m}_k \cdot \mathrm{x}_k + \mathrm{q}$), or the special elements $\bot$ and $\top$, which explicitly represent undefined functions. The element $\top$ is introduced to manifest an irrecoverable precision loss of the analysis. We also support *lexicographic* affine *functions* $(\mathrm{f}_k, \ldots, \mathrm{f}_1, \mathrm{f}_0)$ in the isomorphic form of ordinals $\omega^k \cdot \mathrm{f}_k + \cdots + \omega \cdot \mathrm{f}_1 + \mathrm{f}_0$ [29,40].

The partitioning is dynamic: during the analysis of a control flow graph, partitions (i.e. decision nodes and constraints) are modified by assignments and split (i.e., added) by boolean conditions and when merging control flows. More specifically, for each action $s \in A$, we define sound over-approximating abstract transformers $[\![s]\!]_o : \mathcal{T} \to \mathcal{T}$ as well as *new under-approximating abstract transformers* $[\![s]\!]_u : \mathcal{T} \to \mathcal{T}$. These transformers always increase by one the value of the functions labeling the leaves of a given decision tree to count the number of executed program steps (i.e., actions in the control flow graph). The transformers for boolean conditions and assignments additionally modify the decision nodes by building upon the underlying numerical abstract domain. For instance, the abstract transformer $[\![b]\!]_o$ (resp. $[\![b]\!]_u$) for a boolean condition b uses the underlying numerical domain to obtain an over-approximation (resp. an under-approximation) of b as a set of linear constraints; then it adds these constraints to the given decision tree and discards the paths that become unfeasible (because they do not satisfy the added constraints). Let $\{n \leq 0\}$ (resp. $\{n = 0\}$) be the set of constraints obtained by $[\![b]\!]_o$ (resp. $[\![b]\!]_u$) for the boolean condition $b \equiv n \leq 0 \wedge n\%2 = 0$; then, given the right subtree in Figure 3, $[\![b]\!]_o$ (resp. $[\![b]\!]_u$) would discard the path leading to the leaf with value $2n+2$ by replacing it with a leaf with undefined value $\bot$ (resp. replace $n \leq 0$ with $n = 0$ and replace $2n+2$ with $\bot$). Decision trees are merged using either the approximation join $\curlyvee$ or the computational join $\sqcup$. Both join operators add missing decision nodes from either of the given trees; $\curlyvee$ retains the leaves that are labeled with an undefined function in at least one of the given trees, while $\sqcup$ preserves the leaves that are labeled with a defined function over the leaves labeled with $\bot$ (but preserves the leaves labeled with $\top$ over all other leaves). To minimize the cost of the analysis and to enforce termination, a (dual) widening operator limits the height of the decision trees and the number of maintained partitions.

Abstract Program Semantics for CTL Properties. The abstract program semantics $\Lambda^\natural_\phi : \mathcal{L} \to \mathcal{T}$ for a CTL formula ϕ maps each program point $l \in \mathcal{L}$ to an element $t \in \mathcal{T}$ of the piecewise-defined functions abstract domain. The definition of $\Lambda^\natural_\phi$ for each CTL formula ϕ is summarized in Figure 4 and explained in some detail below. More details and formal definitions can be found in [37,38].

$$\Lambda^\natural_a \stackrel{\text{def}}{=} \lambda l.\ \textsc{reset}\ [\![a]\!]\ (\bot) \tag{5.1}$$

$$\Lambda^\natural_{\mathsf{Q}\mathsf{X}\phi} \stackrel{\text{def}}{=} \lambda l.\ \textsc{zero}\left(\bigsqcup_{(l,s,l')\in out(l)}^{\mathsf{Q}} [\![s]\!]_{\mathsf{Q}}\ (\Lambda^\natural_\phi(l'))\right) \tag{5.2}$$

$$\Lambda^\natural_{\mathsf{Q}(\phi_1\mathsf{U}\phi_2)} \stackrel{\text{def}}{=} \mathrm{lfp}^\natural_{\lambda l.\bot}\ \lambda m.\ \lambda l.\ \textsc{until}\left[\!\!\left[\Lambda^\natural_{\phi_1}(l), \Lambda^\natural_{\phi_2}(l)\right]\!\!\right]\left(\bigsqcup_{(l,s,l')\in out(l)}^{\mathsf{Q}} [\![s]\!]_{\mathsf{Q}}\ (m(l'))\right) \tag{5.3}$$

$$\Lambda^\natural_{\mathsf{Q}\mathsf{G}\phi} \stackrel{\text{def}}{=} \mathrm{gfp}^\natural_{\Lambda^\natural_\phi}\ \lambda m.\ \lambda l.\ \textsc{mask}\left[\!\!\left[\bigsqcup_{(l,s,l')\in out(l)}^{\mathsf{Q}} [\![s]\!]_{\mathsf{Q}}\ (m(l'))\right]\!\!\right](m(l)) \tag{5.4}$$

$$\Lambda^\natural_{\neg\phi} \stackrel{\text{def}}{=} \lambda l.\ \textsc{complement}(\Lambda^\natural_\phi(l)) \tag{5.5}$$

$$\Lambda^\natural_{\phi_1\wedge\phi_2} \stackrel{\text{def}}{=} \lambda l.\ \Lambda^\natural_{\phi_1}(l) \curlyvee \Lambda^\natural_{\phi_2}(l) \tag{5.6}$$

$$\Lambda^\natural_{\phi_1\vee\phi_2} \stackrel{\text{def}}{=} \lambda l.\ \Lambda^\natural_{\phi_1}(l) \sqcup \Lambda^\natural_{\phi_2}(l) \tag{5.7}$$

Fig. 4: Abstract program semantics $\Lambda^\natural_\phi$ for each CTL formula ϕ. The join operator $\bigsqcup^{\mathsf{Q}}$ and the abstract transformer $[\![s]\!]_{\mathsf{Q}}$ respectively stand for $\sqcup$ and $[\![s]\!]_u$, if Q is E, or $\curlyvee$ and $[\![s]\!]_o$, if Q is A. With abuse of notation, we use $\bot$ to denote a decision tree with a single undefined leaf node.

The analysis starts with the totally undefined function (i.e., a decision tree that consists of a single leaf with undefined value $\bot$) at the final program points (i.e., nodes without outgoing edges in the control flow graph). Then it proceeds backwards through the control flow graph, taking the encountered actions into account, and joining decision trees when merging control flows. For existential CTL properties, the analysis uses the under-approximating abstract transformers $[\![s]\!]_u$ for each action s, to ensure that program states that do not satisfy the CTL property are discarded (i.e., removed from the domain of the current piecewise-defined function), and joins decision trees using the computational join $\sqcup$, to ensure that the current piecewise-defined function remains defined over states that satisfy the CTL property in at least one of the merged control flows. Dually, for universal CTL properties, the analysis uses the over-approximating abstract transformers $[\![s]\!]_o$ and joins decision trees using the approximation join $\curlyvee$, to ensure that the current piecewise-defined function remains defined only over states that satisfy the CTL property in all of the merged control flows.

At each program point, the analysis additionally performs operations that are specific to the considered CTL formula ϕ. For an atomic proposition a (cf. Equation 5.1), the analysis performs a $\textsc{reset}\ [\![a]\!]$ operation, which is analogous to the under-approximating transformer for boolean conditions but additionally replaces all the leaves that satisfy a with leaves labeled with the function with value zero. For example, given the atomic proposition $n = 0$ and the right subtree in Figure 3, $\textsc{reset}\ [\![n = 0]\!]$ would replace the constraint $n \leq 0$ with $n = 0$, the leaf $2n + 2$ with $\bot$ and the leaf 2 with 0. For a next-formula $\mathsf{Q}\mathsf{X}\phi$

(cf. Equation 5.2), the analysis approximates the effect of the transition from each program point l to each successor program point l' and performs a ZERO operation to replace all defined functions labeling the leaves of the so obtained decision tree with the function with value zero. For an until-formula $\mathsf{Q}(\phi_1 \,\mathsf{U}\, \phi_2)$ (cf. Equation 5.3), the analysis performs an ascending iteration with widening [13]. At each iteration, the analysis approximates the effect of the transition from each program point l to each successor program point l' and performs an UNTIL operation to model the until temporal operator: UNTIL replaces with the function with value zero all leaves that correspond to defined leaves in the decision tree $\Lambda^\natural_{\phi_2}(l)$ obtained for ϕ_2, and retains all leaves that are labeled with an undefined function in both $\Lambda^\natural_{\phi_1}(l)$ and $\Lambda^\natural_{\phi_1}(l)$. For a globally-formula $\mathsf{QG}\phi$ (cf. Equation 5.4), the analysis performs a descending iteration with dual widening [41], starting from the abstract semantics $\Lambda^\natural_\phi$ obtained for ϕ. At each iteration, the MASK operation models the globally temporal operator: it discards all defined partitions in $\Lambda^\natural_\phi(l)$ that become undefined as a result of the transition from each program point l to each successor program point l'; at the limit, the only defined partitions are those that remain defined across transitions and thus satisfy the globally-formula. For a negation formula $\neg\phi$ (cf. Equation 5.5), the analysis performs a COMPLEMENT operation on the decision tree $\Lambda^\natural_\phi(l)$ obtained for ϕ at each program point l; COMPLEMENT replaces all defined functions labeling the leaves of a decision tree with $\bot$, and all $\bot$ with the function with value zero. Note that $\Lambda^\natural_\phi$ is an abstraction of Λ_ϕ and thus not all undefined partitions in $\Lambda^\natural_\phi$ necessarily correspond to undefined partitions in Λ_ϕ. Leaves that are undefined in $\Lambda^\natural_\phi$ due to this uncertainty are labeled with $\top$, and are left unchanged by COMPLEMENT to guarantee the soundness of the analysis. Finally, for binary formulas $\phi_1 \wedge \phi_2$ and $\phi_1 \vee \phi_2$, the abstract semantics $\Lambda^\natural_{\phi_1 \wedge \phi_2}$ and $\Lambda^\natural_{\phi_1 \vee \phi_2}$ (cf. Equations 5.6 and 5.7) merge the decision trees obtained for ϕ_1 and ϕ_2 using the approximation join $\curlyvee$ and the computational join $\sqcup$, respectively.

The abstract program semantics $\Lambda^\natural_\phi$ for each CTL formula ϕ is *sound* with respect to the approximation order $f_1 \preccurlyeq f_2 \Leftrightarrow \mathrm{dom}(f_1) \supseteq \mathrm{dom}(f_2) \wedge \forall x \in \mathrm{dom}(f_1) : f_1(x) \leq f_2(x)$, which means that the abstract semantics $\Lambda^\natural_\phi$ *over-approximates* the value of the concrete semantics Λ_ϕ and *under-approximates* its domain of definition $\mathrm{dom}(\Lambda_\phi)$. In this way, the abstraction provides sufficient preconditions for the CTL property ϕ: if the abstraction is defined for a state then that state satisfies ϕ.

Theorem 2. *A program satisfies a CTL formula ϕ for all traces starting from a given set of states $\mathcal{I}$ if $\mathcal{I} \subseteq \mathrm{dom}(\gamma(\Lambda^\natural_\phi))$.*

Proof (Sketch). The proof proceeds by induction over the structure of the formula ϕ. The base case are atomic propositions for which the proof is immediate.

For a next-formula $\mathsf{QX}\phi$, by induction hypothesis, $\mathrm{dom}(\Lambda^\natural_\phi)$ is a subset of the set of states that satisfy ϕ. Using the over-approximating transformers $[\![s]\!]_o$ together with the approximation join $\curlyvee$ (resp. the under-approximating transformers $[\![s]\!]_u$ together with the computational join $\sqcup$) ensures that $\Lambda^\natural_{\mathsf{QX}\phi}$ soundly under-approximates the set of states that satisfy $\mathsf{QX}\phi$.

For an until-formula $\mathsf{Q}(\phi_1 \ \mathsf{U} \ \phi_2)$, by induction hypothesis, $\mathrm{dom}(\Lambda^\natural_{\phi_1})$ and $\mathrm{dom}(\Lambda^\natural_{\phi_2})$ are a subset of the set of states that satisfy ϕ_1 and ϕ_2, respectively. By definition, $\Lambda_{\mathsf{Q}(\phi_1 \mathsf{U} \phi_2)}$ is the limit of an ascending iteration sequence using widening. Again, using the over-approximating transformers $[\![s]\!]_o$ together with the approximation join $\curlyvee$ (resp. the under-approximating transformers $[\![s]\!]_u$ together with the computational join $\sqcup$) guarantees the soundness of the analysis with respect to each transition. The soundness of each iteration without widening is then guaranteed by the definition of the UNTIL operation. The iterations with widening are allowed to be unsound but the limit of the iterations (i.e., $\Lambda_{\mathsf{Q}(\phi_1 \mathsf{U} \phi_2)}$) is guaranteed to soundly under-approximate the set of states that satisfy $(\phi_1 \ \mathsf{U} \ \phi_2)$. We refer to [38] for a detailed proof for formulas of the form $(\mathit{true} \ \mathsf{U} \ \phi_2)$. The generalization to $(\phi_1 \ \mathsf{U} \ \phi_2)$ is trivial.

For a globally-formula $\mathsf{QG}\phi$, $\Lambda_{\mathsf{QG}\phi}$ is the limit of a descending iteration sequence with dual widening, starting from $\Lambda^\natural_\phi$, which soundly under-approximates the set of states that satisfy ϕ. The soundness of each iteration is guaranteed by the definition of the MASK operation and the dual widening operator (see [38]).

The case of a negation $\neg\phi$ is non-trivial since, by induction hypothesis, $\mathrm{dom}(\Lambda^\natural_\phi)$ is a subset of the set of states that satisfy ϕ. Specifically, $\Lambda^\natural_\phi$ maps each program point $l \in \mathcal{L}$ to a decision tree whose leaves determine this under-approximation: leaves labeled with $\bot$ represent states that do not satisfy ϕ while leaves labeled with $\top$ represent states that may or may not satisfy ϕ. The COMPLEMENT operation performed by $\Lambda^\natural_{\neg\phi}$ only considers leaves labeled by $\bot$ and ignores (i.e., leaves unchanged) leaves labeled by $\top$. Thus, $\Lambda^\natural_{\neg\phi}$ soundly under-approximates the set of states that satisfy $\neg\phi$.

Finally, for binary formulas $\phi_1 \wedge \phi_2$ and $\phi_1 \vee \phi_2$, the proof follows immediately from the definition of the approximation join $\curlyvee$ and the computational join $\sqcup$ used in the definition of $\Lambda^\natural_{\phi_1 \wedge \phi_2}$ and $\Lambda^\natural_{\phi_1 \vee \phi_2}$, respectively. □

6 Implementation

The proposed static analysis method for proving CTL properties is implemented in the prototype static analyzer FUNCTION [13].

The implementation is in OCAML and consists of around 9K lines of code. The current front-end of FUNCTION accepts non-deterministic programs written in a C-like syntax (without support for pointers, `struct` and `union` types). The only basic data type is mathematical integers. FUNCTION accepts CTL properties written using a syntax similar to the one used in the rest of this paper, with atomic propositions written as C-like pure expressions. The abstract domain of piecewise-defined functions builds on the numerical abstract domains provided by the APRON library [24], and the under-approximating numerical operators provided by the BANAL static analyzer [31].

The analysis is performed backwards on the control flow graph of a program with a standard worklist algorithm [32], using widening and dual widening at loop heads. Non-recursive function calls are inlined, while recursion is supported by augmenting the control flow graphs with call and return edges. The

No	Program	CTL Property	Result	Time
1.1	and_test.c	$\mathsf{AGAF}(\mathrm{n} = 1) \wedge \mathsf{AF}(\mathrm{n} = 0)$	✓	0.05s
1.2	and_test.c	$\mathsf{EGAF}(\mathrm{n} = 1)$	✓	0.05s
1.4	global_test.c	$\mathsf{AGEF}(\mathrm{x} \leq -10)$	✓	0.15s
1.7	or_test.c	$\mathsf{AFEG}(\mathrm{x} < -100) \vee \mathsf{AF}(\mathrm{x} = 20)$	✓	0.05s
1.8	may_term...	$\mathsf{EF}(\mathrm{exit} : \mathit{true})$	✗	-
1.9	until_test.c	$\mathsf{A}(\mathrm{x} \geq \mathrm{y}\ \mathsf{U}\ \mathrm{x} = \mathrm{y})$	✓	0.03s
1.11	fin_ex.c	$\mathsf{EGEF}(\mathrm{n} = 1)$	✓	0.04s
1.12	until_ex.c	$\mathsf{E}(\mathrm{x} \geq \mathrm{y}\ \mathsf{U}\ \mathrm{x} = \mathrm{y})$	✓	0.03s
2.3	win4.c	$\mathsf{AFAG}(\mathrm{WItemsNum} \geq 1)$	✓	0.15s
2.4	toylin.c	$(\mathrm{c} \leq 5 \wedge \mathrm{c} > 0) \vee \mathsf{AF}(\mathrm{resp} > 5)$	✗	-
3.9	cb5_safe.c	$\mathsf{A}(\mathrm{i} = 0\ \mathsf{U}\ (\mathsf{A}(\mathrm{i} = 1\ \mathsf{U}\ \mathsf{AG}(\mathrm{i} = 3)) \vee \mathsf{AG}(\mathrm{i} = 1)))$	✗	-
3.14	timer...	$\neg\mathsf{AG}(\mathrm{timer} = 0 \Rightarrow \mathsf{AF}(\mathrm{output} = 1))$	✗	-
3.15	togglec...	$\mathsf{AG}(\mathsf{AF}(\mathrm{t} = 1) \wedge \mathsf{AF}(\mathrm{t} = 0))$	✗	-
4.1	Bangalore...	$\mathsf{EF}(\mathrm{x} < 0)$	✗	-
4.2	Ex02...	$\mathrm{i} < 5 \Rightarrow \mathsf{AF}(\mathrm{exit} : \mathit{true})$	✓	0.04s
4.3	Ex07...	$\mathsf{AFEG}(\mathrm{i} = 0)$	✓	0.1s
4.4	java_Seq...	$\mathsf{EF}(\mathsf{AF}(\mathrm{j} \geq 21) \wedge \mathrm{i} = 100)$	✓	0.3s
4.5	Madrid...	$\mathsf{AF}(\mathrm{x} = 7 \wedge \mathsf{EFAG}(\mathrm{x} = 2))$	✓	0.02s

Fig. 5: Evaluation of FuncTion on selected test cases collected from various sources. All test cases were analyzed using polyhedral constraints for the decision nodes, and affine functions for the leaf nodes of the decision tree.

precision of the analysis can be tuned by choosing the underlying numerical abstract domain, by activating the extension to ordinal-value functions [40], and by adjusting the precision of the widening [13] and the widening delay. It is also possible to refine the analysis by considering only reachable states.

Experimental Evaluation. We evaluated our technique on a number of test cases obtained from various sources, and compared FuncTion against T2 [8] and Ultimate LTL Automizer [20] as well as E-HSF [4], and the prototype implementation from [10]. Figures 5 and 6 show an excerpt of the results, which demonstrates the differences between FuncTion, T2 [8] and Ultimate LTL Automizer. The first set of test cases are programs that we used to test our implementation. The second and third set were collected from [25] and the web interface of Ultimate LTL Automizer [20]. The fourth set are examples from the termination category of the 6th International Competition on Software Verification (SV-COMP 2017). The experiments were conducted on an Intel i7-6600U processor with 20GB of RAM on Arch Linux with Linux 4.11 and OCaml 4.04.1.

FuncTion passes all test cases with the exception of 2.4, 3.9, 3.14, and 3.15, which fail due to imprecisions introduced by the widening, and 1.8 and 4.1, which fail due to an unfortunate interaction of the under-approximations needed for existential properties and non-deterministic assignments in the programs. However, note that for these test cases we still get some useful information. For instance, for 3.15, FuncTion infers that the CTL property is satisfied if $\mathrm{x} < 0$.

No	FuncTion	T2 [8]	Ultimate LTL Automizer [20]
1.1	✓	✗	✓
1.2	✓	✗	-
1.4	✓	✗	-
1.7	✓	✗ (error)	-
1.8	✗	✓	-
1.9	✓	✗	✓
1.11	✓	✗	-
1.12	✓	✗ (no implementation)	-
2.3	✓	✗	✓
2.4	✗	✗	✓
3.9	✗	-	✓
3.14	✗	-	✓
3.15	✗	-	✓
4.1	✗	✓	-
4.2	✓	✗ (out of memory)	✓
4.3	✓	✗	-
4.4	✓	✗ (error)	-
4.5	✓	✗	-

Fig. 6: Differences between FuncTion, T2, and Ultimate LTL Automizer.

In Figure 6, the missing results for T2 are due to a missing conversion of the test cases to the T2 input format. The comparison with Ultimate LTL Automizer is limited to the test cases where the CTL property can be equivalently expressed in LTL (i.e., universal CTL properties). The results show that only FuncTion succeeds on numerous test cases (1.2, 1.4, 1.7, 1.11, 1.12, 4.3, 4.4, and 4.5). Ultimate LTL Automizer performs well on the supported test cases, but FuncTion still succeeds on most of the test cases provided by Ultimate LTL Automizer (not shown in Figure 6, since there are no differences between the results of FuncTion and Ultimate LTL Automizer). Overall, none of the tools subsumes the others. In fact, we observe that their combination is more powerful than any of the tools alone, as it would succeed on all test cases.

Finally, FuncTion only succeeds on two of the industrial benchmarks from [10], while T2, E-HSF and [10] fare much better (see [8, Figure 11]). The reason for the poor performance is that in these benchmarks the effect of function calls is modeled as a non-deterministic assignment and this heavily impacts the precision of FuncTion. We are confident that we would obtain better results on the original benchmarks, where function calls are not abstracted away.

7 Related Work

In the recent past, a large body of work has been devoted to proving CTL properties of programs. The problem has been extensively studied for finite-state programs [7,26, etc.], while most of the existing approaches for infinite-state

systems have limitations that restrict their applicability. For instance, they only support certain classes of programs [36], or they limit their scope to a subset of CTL [11], or to a single CTL property such as termination [27,34, etc.] or non-termination [2,5, etc.]. Our approach does not suffer from these limitations.

Some other approaches for proving CTL properties do not reliably support CTL formulas with arbitrary nesting of universal and existential path quantifiers [23], or support existential path quantifiers only indirectly by building upon recent work for proving non-termination [22], or by considering their universal dual [8]. In particular, the latter approach is problematic: since the universal dual of an existential until formula is non-trivial to define, the current implementation of T2 does not support such formulas (see Figure 6). Other indirect approaches [4,10] perform unnecessary computations that result in slower runtimes (see [8, Figure 12]). In comparison to all these approaches, our approach provides strictly more information in the form of a ranking function whose domain gives a precondition for a given CTL property and whose value estimates the number of program execution steps until the property is satisfied.

In [17], Cousot and Cousot define a trace-based semantics for a very general temporal language which subsumes LTL and CTL; this is subsequently abstracted to a state-based semantics. The abstraction has been later shown to be incomplete by Giacobazzi and Ranzato [21]. In contrast to the work of Cousot and Cousot, we do not define a trace-based semantics for CTL. The semantics that we propose is close to their state-based semantics in that their state-based semantics coincides with the domain of the functions that we define. Note that Theorem 1 is not in contrast with the result of Giacobazzi and Ranzato because completeness is proven with respect to the state-based semantics of CTL.

Finally, our abstract interpretation framework generalizes an existing framework [41] for proving guarantee and recurrence properties of programs [28]. Guarantee and recurrence properties are equivalently expressed in CTL as $\mathsf{A}(\mathit{true}\ \mathsf{U}\ \phi)$ and $\mathsf{AGA}(\mathit{true}\ \mathsf{U}\ \phi)$, respectively. In fact, we rediscover the guarantee and recurrence program semantics defined in [41] as instances of our framework: the guarantee semantics coincides with $\Lambda_{\mathsf{A}(\mathit{true}\mathsf{U}\phi)}$ (cf. Section 4) and the recurrence semantics coincides with $\Lambda_{\mathsf{AGA}(\mathit{true}\mathsf{U}\phi)}$ (cf. Section 4). The common insight with our work is the observation that CTL (sub)formulas are satisfied by finite subsequences (which can also be single states) of possibly infinite sequences. The program semantics for these (sub)formulas then counts the number of steps in these subsequences. Our work generalizes this idea to all CTL formulas and integrates the corresponding semantics in a uniform framework.

8 Conclusion and Future Work

In this paper, we have presented a new static analysis method for inferring preconditions for CTL properties of programs that overcomes the limitations of existing approaches. We have derived our static analysis within the framework of abstract interpretation by abstraction of the operational trace semantics of a program. Using experimental evidence, we have shown that our analysis is

effective and performs well on a wide variety of benchmarks, and is able to prove CTL properties that are out of reach for state-of-the-art tools.

It remains for future work to investigate and improve the precision of the analysis in the presence of non-deterministic program assignments. We also plan to support LTL properties [20] or, more generally, CTL* properties [9]. This requires some form of trace partitioning [35] as the interpretation of LTL formulas is defined in terms of program executions instead of program states as CTL.

References

1. C. Baier and J. P. Katoen. *Principles of Model Checking*. MIT Press, 2008.
2. A. Bakhirkin and N. Piterman. Finding Recurrent Sets with Backward Analysis and Trace Partitioning. In *TACAS*, pages 17–35, 2016.
3. J. Bertrane, P. Cousot, R. Cousot, J. Feret, L. Mauborgne, A. Miné, and X. Rival. Static Analysis and Verification of Aerospace Software by Abstract Interpretation. In *AIAA*, pages 1–38, 2010.
4. T. A. Beyene, C. Popeea, and A. Rybalchenko. Solving Existentially Quantified Horn Clauses. In *CAV*, pages 869–882, 2013.
5. H. Y. Chen, B. Cook, C. Fuhs, K. Nimkar, and P. W. O'Hearn. Proving Nontermination via Safety. In *TACAS*, pages 156–171, 2014.
6. E. M. Clarke and E. A. Emerson. Design and Synthesis of Synchronization Skeletons using Branching-Time Temporal Logic. In *Logic of Programs*, pages 52–71, 1981.
7. E. M. Clarke, E. A. Emerson, and A. P. Sistla. Automatic Verification of Finite-State Concurrent Systems Using Temporal Logic Specifications. *ACM Transactions on Programming Languages and Systems*, 8(2):244–263, 1986.
8. B. Cook, H. Khlaaf, and N. Piterman. Faster Temporal Reasoning for Infinite-State Programs. In *FMCAD*, page 5782, 2014.
9. B. Cook, H. Khlaaf, and N. Piterman. On Automation of CTL* Verification for Infinite-State Systems. In *CAV*, pages 13–29, 2015.
10. B. Cook and E. Koskinen. Reasoning about Nondeterminism in Programs. In *PLDI*, pages 219–230, 2013.
11. B. Cook, E. Koskinen, and M. Y. Vardi. Temporal Property Verification as a Program Analysis Task. In *CAV*, pages 333–348, 2011.
12. B. Cook, E. Koskinen, and M. Y. Vardi. Temporal Property Verification as a Program Analysis Task - Extended Version. *Formal Methods in Systems Design*, 41(1):66–82, 2012.
13. N. Courant and C. Urban. Precise Widening Operators for Proving Termination by Abstract Interpretation. In *TACAS*, pages 136–152, 2017.
14. P. Cousot. Constructive Design of a Hierarchy of Semantics of a Transition System by Abstract Interpretation. *Theoretical Computer Science*, 277(1-2):47–103, 2002.
15. P. Cousot and R. Cousot. Static Determination of Dynamic Properties of Programs. In *Symposium on Programming*, pages 106–130, 1976.
16. P. Cousot and R. Cousot. Abstract Interpretation: a Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fixpoints. In *POPL*, pages 238–252, 1977.
17. P. Cousot and R. Cousot. Temporal Abstract Interpretation. In *POPL*, pages 12–25, 2000.

18. P. Cousot and R. Cousot. An Abstract Interpretation Framework for Termination. In *POPL*, pages 245–258, 2012.
19. P. Cousot and N. Halbwachs. Automatic Discovery of Linear Restraints Among Variables of a Program. In *POPL*, pages 84–96, 1978.
20. D. Dietsch, M. Heizmann, V. Langenfeld, and A. Podelski. Fairness Modulo Theory: A New Approach to LTL Software Model Checking. In *CAV*, pages 49–66, 2015.
21. R. Giacobazzi and F. Ranzato. Incompleteness of states w.r.t. traces in model checking. *Information and Computation*, 204(3):376–407, 2006.
22. A. Gupta, T. A. Henzinger, R. Majumdar, A. Rybalchenko, and R. Xu. Proving Non-Termination. In *POPL*, pages 147–158, 2008.
23. A. Gurfinkel, O. Wei, and M. Chechik. Yasm: A Software Model-Checker for Verification and Refutation. In *CAV*, pages 170–174, 2006.
24. B. Jeannet and A. Miné. Apron: A Library of Numerical Abstract Domains for Static Analysis. In *CAV*, page 661667, 2009.
25. E. Koskinen. *Temporal Verification of Programs*. PhD thesis, University of Cambridge, November 2012.
26. O. Kupferman, M. Y. Vardi, and P. Wolper. An Automata-Theoretic Approach to Branching-Time Model Checking. *Journal of the ACM*, 47(2):312–360, 2000.
27. C. S. Lee, N. D. Jones, and A. M. Ben-Amram. The Size-Change Principle for Program Termination. In *POPL*, pages 81–92, 2001.
28. Z. Manna and A. Pnueli. A Hierarchy of Temporal Properties. In *PODC*, pages 377–410, 1990.
29. Z. Manna and A. Pnueli. The Temporal Verification of Reactive Systems: Progress, 1996.
30. A. Miné. The Octagon Abstract Domain. *Higher Order and Symbolic Computation*, 19(1):31–100, 2006.
31. A. Miné. Inferring Sufficient Conditions with Backward Polyhedral Under-Approximations. *Electronic Notes in Theoretical Computer Science*, 287:89–100, 2012.
32. F. Nielson, H. R. Nielson, and C. Hankin. *Principles of Program Analysis*. Springer, 1999.
33. A. Pnueli. The Temporal Logic of Programs. In *FOCS*, pages 46–57, 1977.
34. A. Podelski and A. Rybalchenko. Transition Invariants. In *LICS*, pages 32–41, 2004.
35. X. Rival and L. Mauborgne. The Trace Partitioning Abstract Domain. *ACM TOPLAS*, 29(5):26, 2007.
36. F. Song and T. Touili. Efficient CTL Model-Checking for Pushdown Systems. *Theoretical Computer Science*, 549:127–145, 2014.
37. S. Ueltschi. Proving Temporal Properties by Abstract Interpretation. Master's thesis, ETH Zurich, Zurich, Switzerland, 2017.
38. C. Urban. *Static Analysis by Abstract Interpretation of Functional Temporal Properties of Programs*. PhD thesis, École Normale Supérieure, Paris, France, July 2015.
39. C. Urban and A. Miné. A Decision Tree Abstract Domain for Proving Conditional Termination. In *SAS*, pages 302–318, 2014.
40. C. Urban and A. Miné. An Abstract Domain to Infer Ordinal-Valued Ranking Functions. In *ESOP*, pages 412–431, 2014.
41. C. Urban and A. Miné. Inference of Ranking Functions for Proving Temporal Properties by Abstract Interpretation. *Computer Languages, Systems and Structures*, 47:77–103, 2017.