



Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Proving Temporal Properties by Abstract Interpretation

Master's Thesis Report

Samuel Ueltschi

September 6, 2017

Supervised by Dr. Caterina Urban

Prof. Dr. Peter Müller
Chair of Programming Methodology
Department of Computer Science, ETH Zurich

Abstract

The intended behavior of a program can be specified with temporal logic. Computation-Tree-Logic (CTL) is a temporal logic that quantifies over possible program executions in the face of non-determinism. It can state properties about all possible program executions and about the existence of some possible program execution. In this thesis we present a static analysis framework based on abstract interpretation for verifying if a program satisfies a given CTL property. In case of liveness properties, properties that state that something good happens eventually, the result of the analysis indicates how long the program will take at most until the good outcome is reached. The analysis is implemented in the research prototype static analyzer FuncTion. An initial evaluation suggests that FuncTion is on par with other existing comparable static analyzers for temporal properties.

Acknowledgements

I would like to thank Caterina Urban for the valuable discussions and inputs in our weekly meetings, her feedback while writing this report and her overall great support during my work on this thesis. Furthermore, I am grateful to Prof. Peter Müller for giving me the opportunity to work on this project.

Contents

1	Introduction	2
2	State Transition Systems	4
3	Computation Tree Logic (CTL)	5
3.1	Syntax	6
3.2	Semantic	6
4	Ranking Functions	7
5	Concrete Semantics for CTL	9
5.1	Path Independent Operators	10
5.2	Path Dependent Operators	11
6	Imperative Language	16
7	Decision Tree Abstract Domain	17
7.1	Domain	18
7.2	Join	20
7.3	Meet	21
7.4	Widening and Dual Widening	21
7.5	Filter	21
7.6	Backward Assign	22
8	Abstract Semantics for CTL	22
8.1	Path Independent Operators	22
8.2	Path Dependent Operators	24
9	Implementation	32
9.1	Previous Work	32
9.2	CTL Analysis	32
9.3	Improved Front-end	33
9.4	Examples	34
10	Evaluation	37
11	Conclusion	42
11.1	Future Work	42

1 Introduction

Temporal logic is a widely accepted language for the specification of the intended behavior of various systems (operating systems, embedded systems, network communication protocols, etc.). Powerful tools have emerged over the years for proving temporal logic properties of programs. Useful properties that temporal logic allows to express are *safety properties* such as partial correctness (i.e., the program does not produce the wrong answer), mutual exclusion (i.e., two processes are not in their critical section at the same time), and deadlock-freedom (i.e., the program does not reach a deadlock state), and *liveness properties* such as termination (i.e., the program eventually terminates), and starvation freedom (i.e., a process eventually receives service).

Abstract interpretation [1] is a general theory for approximating the behavior of programs that was developed by Patrick Cousot and Radhia Cousot in the late 1970s. Recently, Cousot and Cousot developed an abstract interpretation framework for proving termination [2]. This framework uses abstract interpretation to derive *ranking functions*. These functions assign an upper bound on the number of steps until termination to program states. Caterina Urban and Antoine Miné extended this framework to the analysis of *guarantee properties* and *recurrence properties* [3]. Termination, *guarantee* and *recurrence* properties are liveness properties. Termination is a liveness property that states that a program terminates eventually, *guarantee properties* state that some general desired program state will be reached eventually during program execution and *recurrence properties* state that some general desired program state will be reached infinitely often.

This thesis extends the existing framework for analyzing guarantee and recurrence properties to general temporal properties specified in Computation-Tree-Logic (CTL). CTL is a temporal logic for specifying the behavior of state transition systems [4]. It can be used to express various temporal properties. For instance, the previously mentioned liveness properties can all be expressed in CTL. In addition to liveness properties, CTL can also express safety properties.

When dealing with non-deterministic programs, it is often of interest to quantify over the set of possible program executions. Either one wants to reason about all possible program executions, or, one is interested in the existence of at least one possible program execution that satisfies a given property. These two kinds of properties are called universal and existential temporal properties and both can be expressed in CTL.

We now show two example CTL properties for the program `SAMPLE` in Figure 1. We assume $x \geq 0$ at the start of the program. `SAMPLE` decrements x until it reaches the value 0 if either $y = 1$ or if a random non-deterministic coin toss succeeds. Otherwise the program sets x to -1 . Because of non-determinism, this program has two possible execution traces, one where the if-block is entered and one where it is not. Consider

```

assume( $x \geq 0$ )
if ( $y = 1 \vee ?$ ) then
  while ( $x > 0$ ) do
     $x := x - 1$ 
  od
else
   $x := -1$ 
fi

```

Figure 1: Program SAMPLE

the CTL property $y = 1 \Rightarrow \forall(x \geq 0 \ U \ x = 0)$. Informally, the property states that if initially $y = 1$, i.e., the if-block is entered, then $x \geq 0$ holds among all possible execution traces until a state is reached where $x = 0$ holds (a formal definition of CTL can be found in Section 3). Due non-determinism, only one of the two program traces modifies x as stated in the property. This fact can be expressed with the CTL property $\exists(x \geq 0 \ U \ x = 0)$. It states that there exists at least one execution trace where $x \geq 0$ holds until a state is reached where $x = 0$ holds. This example demonstrates the difference between universal and existential CTL properties and it shows the expressiveness of CTL. Note that there are temporal logic systems like CTL* that are strictly more expressive than CTL [4]. We chose CTL because it is a natural extension of the previous work on the analysis of liveness properties with abstract interpretation.

Our framework proves CTL properties by inferring *counting functions*. A counting function is a partial function that assigns natural numbers to program states. A program satisfies a given CTL property if the initial state is in the domain of this function. If the CTL property is a liveness property, i.e., a property stating that some goal state will be reached eventually, then the value of the counting function is an upper bound on the number of steps until the goal is reached. Otherwise, the value of the function has no particular meaning. Counting functions are an adaptation of ranking functions as defined in [2] and [3].

Consider the property $\exists(x \geq 0 \ U \ x = 0)$ from the previous example. This property is a liveness property that states that $x = 0$ holds eventually for at least one possible program execution. Our framework derives the following counting function for the initial program state through backward interpretation of the control-flow-graph of the program.

$$f(x) = \begin{cases} 2x + 1 & \text{if } x \geq 0 \\ 0 & \text{if } x = 0 \\ \text{undefined} & \text{otherwise} \end{cases}$$

If the initial value of x is greater than 0 then the goal of the liveness property is reached in exactly $2x + 1$ program execution steps ($2x$ steps inside the while loop plus one step to

enter the if-block). It takes 0 steps if the goal is already reached. Otherwise the function is undefined, stating that the analysis does not know if the property is satisfied if $x < 0$. In this case it is clear that the property is not satisfied if $x < 0$, however, the analysis is not able to prove or disprove all CTL properties in general.

Deciding if a CTL property is satisfied is in general undecidable. We achieve a decidable approximating of counting functions by performing the analysis on the decision tree abstract domain introduced in [5]. This abstract domain approximates counting functions in terms of piecewise-defined linear functions. We extend this domain with new operators for analyzing CTL properties. Existential CTL properties require the addition of under-approximating operators for the abstract domain.

An implementation of the CTL analysis presented in this thesis was added to the FuncTion static analyzer. We conducted an evaluation of the FuncTion CTL analysis capabilities which suggests that our implementation is on par with other comparable static analyzers for temporal properties.

The remainder of this document is structured as follows. Sections 2 and 3 formally define state transition systems and CTL. Section 4 introduces ranking functions and the existing techniques for proving liveness, guarantee and recurrence properties on which this thesis is based on. Section 5 defines how to verify if a state transition systems satisfies a CTL property. In Section 6 we introduce a simple imperative programming language. This language is used as an example for the definition of the CTL analysis. Section 7 introduces the decision tree abstract domain which is then used in Section 8 where the CTL analysis is defined for the previously introduced imperative language. Section 9 presents the implementation for FuncTion and Section 10 presents the experimental evaluation. Finally Section 11 concludes and gives an outlook for future work.

We assume that the reader is familiar with the theory of abstract interpretation.

2 State Transition Systems

To be able to analyze the behavior of a program, it is necessary to express said behavior through a mathematical model. We model the operational semantics of programs using transition systems. This is based on the definitions presented in [3].

Definition 2.1. A transition system is a tuple $\langle \Sigma, \tau \rangle$ where Σ is the set of all states in the system and $\tau \in \Sigma \times \Sigma$ is the relation that defines if one can transition from one state to the other (called *transition relation*).

Transition systems allow us to model the semantics of a program independently of the programming language in which it was written. By expressing the possible transitions between states in terms of a relation, it is also possible to capture nondeterminism. Figure 2 shows a simple transition system represented as a directed graph. States are




Figure 2: A basic state transition system

represented as nodes and state transitions as directed edges.

We introduce the following auxiliary functions over states of a transition systems which will become useful in Section 5 where we define the semantics of CTL operators in terms of transition systems.

Definition 2.2. Let $\langle \Sigma, \tau \rangle$ be a transition system. The function $\text{pre}: \mathcal{P}(\Sigma) \rightarrow \mathcal{P}(\Sigma)$ maps a set of states $X \in \mathcal{P}(\Sigma)$ to the set of their predecessors with respect to the program transition relation τ :

$$\text{pre}(X) \stackrel{\text{def}}{=} \{s \in \Sigma \mid \exists s' \in X: \langle s, s' \rangle \in \tau\} \quad (1)$$

Definition 2.3. Let $\langle \Sigma, \tau \rangle$ be a transition system. The function $\widetilde{\text{pre}}: \mathcal{P}(\Sigma) \rightarrow \mathcal{P}(\Sigma)$ maps a set of states $X \in \mathcal{P}(\Sigma)$ to the set of their predecessors with respect to the program transition relation τ with the limitation that only those predecessor states are selected which exclusively transition to states in X :

$$\widetilde{\text{pre}}(X) \stackrel{\text{def}}{=} \{s \in \Sigma \mid \forall s' \in X: \langle s, s' \rangle \in \tau \Rightarrow s' \in X\} \quad (2)$$

To get an intuition for the difference between $\widetilde{\text{pre}}$ and pre , consider the state transition system depicted in Figure 2. There it holds that $\text{pre}(\{b, d\}) = \{a, c\}$ because a is the predecessor of b and c the predecessor of d . However note that $\widetilde{\text{pre}}(\{b, d\}) = \{c\}$ since only c has transitions that exclusively end up in either b or d . Consequently it holds that $\widetilde{\text{pre}}(\{b, c\}) = \{a\}$ because a transitions exclusively to either b or c .

The next section introduces Computation Tree Logic, a logic for stating properties about transition systems.

3 Computation Tree Logic (CTL)

Computation Tree Logic (CTL) is a logic which allows stating properties about execution traces of state transition systems. In the context of this thesis, CTL is used to express temporal properties about the runtime behavior of programs. This section gives a brief introduction into the syntax and semantic of CTL. Further information about

CTL can be found in [4].

We assume the existence of some atomic proposition logic over the set of states. Atomic propositions $a \in \mathbf{AP}$ state properties of states $\sigma \in \Sigma$. The satisfaction relation $\models \subseteq \Sigma \times \mathbf{AP}$ determines if a state satisfies an atomic proposition.

3.1 Syntax

The syntax of a CTL formula is given by the following grammar.

$$\begin{aligned} \Phi ::= & \\ & a \mid \\ & \neg\Phi \mid \Phi \wedge \Phi \mid \Phi \vee \Phi \mid \\ & \forall \bigcirc \Phi \mid \exists \bigcirc \Phi \mid \\ & \forall \diamond \Phi \mid \exists \diamond \Phi \mid \\ & \forall \square \Phi \mid \exists \square \Phi \mid \\ & \forall (\Phi \ U \ \Phi) \mid \exists (\Phi \ U \ \Phi) \end{aligned}$$

The term a is an atomic propositions. Formulae with quantifiers \exists or \forall are called path-dependent, formulae without are called path-independent.

3.2 Semantic

We now define the satisfaction relation \models between states $\sigma \in \Sigma$ and CTL formulae. The satisfaction relation for atomic propositions $a \in \mathbf{AP}$ is assumed to be given by the underlying logic for atomic propositions.

$$\begin{aligned} \sigma \models a & \iff \sigma \models a \\ \sigma \models \neg\Phi & \iff \text{not } \sigma \models \Phi \\ \sigma \models \Phi_1 \wedge \Phi_2 & \iff (\sigma \models \Phi_1) \text{ and } (\sigma \models \Phi_2) \\ \sigma \models \Phi_1 \vee \Phi_2 & \iff (\sigma \models \Phi_1) \text{ or } (\sigma \models \Phi_2) \\ \sigma \models \forall \bigcirc \Phi & \iff \forall \pi \in \text{Paths}(\sigma): (\pi[1] \models \Phi) \\ \sigma \models \exists \bigcirc \Phi & \iff \exists \pi \in \text{Paths}(\sigma): (\pi[1] \models \Phi) \\ \sigma \models \forall (\Phi_1 \ U \ \Phi_2) & \iff \forall \pi \in \text{Paths}(\sigma): \\ & \quad (\exists j \geq 0: \pi[j] \models \Phi_2 \wedge (\forall 0 \leq k < j: \pi[k] \models \Phi_1)) \\ \sigma \models \exists (\Phi_1 \ U \ \Phi_2) & \iff \exists \pi \in \text{Paths}(\sigma): \\ & \quad (\exists j \geq 0: \pi[j] \models \Phi_2 \wedge (\forall 0 \leq k < j: \pi[k] \models \Phi_1)) \\ \sigma \models \forall \square \Phi & \iff \forall \pi \in \text{Paths}(\sigma): (\forall j \geq 0: \pi[j] \models \Phi) \\ \sigma \models \exists \square \Phi & \iff \exists \pi \in \text{Paths}(\sigma): (\forall j \geq 0: \pi[j] \models \Phi) \end{aligned}$$

The states $\sigma \in \Sigma$ are part of a state transition system $\langle \Sigma, \tau \rangle$ and $Paths(\sigma_0)$ is the set of all paths $\pi = \sigma_0 \sigma_1 \sigma_2 \dots$ starting from σ_0 with $\pi[j] = \sigma_j$. The CTL formulae $\forall \diamond \Phi$ and $\exists \diamond \Phi$ are not explicitly defined. They are equivalent to $\forall (true \ U \ \Phi)$ and $\exists (true \ U \ \Phi)$.

Properties $\forall (\Phi_1 \ U \ \Phi_2)$ and $\exists (\Phi_1 \ U \ \Phi_2)$ are called *until* properties. These properties require that a state satisfying Φ_2 is reached eventually among all (some, respectively) paths, and until then Φ_1 holds.

Properties $\forall \square \Phi$ and $\exists \square \Phi$ are called *global* properties. They state that all states among all (some, respectively) reachable paths satisfy the property Φ .

Properties $\forall \bigcirc \Phi$ and $\exists \bigcirc \Phi$ are called *next* properties. They state that all (some, respectively) immediate next states satisfy the property Φ .

The following useful equivalence relations exists which can be used to relate existential to universal CTL formulae.

$$\begin{aligned} \exists \bigcirc \Phi &\equiv \neg \forall \bigcirc (\neg \Phi) \\ \exists \diamond \Phi &\equiv \neg \forall \square (\neg \Phi) \\ \exists \square \Phi &\equiv \neg \forall \diamond (\neg \Phi) \end{aligned}$$

The next section introduces the concept of ranking functions, a proof method for liveness properties. This proof method will then be extended to CTL in Section 5.

4 Ranking Functions

The traditional method for proving termination are *ranking functions* [6] [7]. A ranking function is a partial function from program states to a well-ordered set like natural numbers or ordinals. To prove termination, the value of the ranking function must decrease during program execution. Cousot and Cousot prove the existence of a *most precise ranking function* and that it can be derived by abstract interpretation [2]. We will call this *most precise ranking function*, as defined by Cousot and Cousot, the *termination semantics* from now on.

Definition 4.1. The *termination semantics* is a ranking function $\tau^t \in \Sigma \rightarrow \mathbb{O}$. A program starting from some initial state $\sigma \in \Sigma$ terminates if and only if $\sigma \in \text{dom}(\tau^t)$.

The *termination semantics* assigns an upper bound on the number of steps until termination to each state. Therefore, a program terminates if its initial state is in the domain of the *termination semantics*.

Based on the work of Cousot and Cousot [2], Urban and Miné [3] extended the *termination semantics* to other liveness properties. More precisely to guarantee and recurrence properties. A guarantee property states that some state satisfying a given property is guaranteed to be reached eventually. Termination is therefore just a guarantee property stating that some final state will be eventually reached. As with termination, the *guarantee semantics* is a ranking function that assigns each state an upper bound on the number of steps until a state satisfying said property is reached. Guarantee properties can be expressed using the CTL formula $\forall\Diamond(a)$

Definition 4.2. The *guarantee semantics* is a ranking function $\tau_{[S]}^g \in \Sigma \rightarrow \mathbb{O}$ where $S \subseteq \Sigma$ is a set of states satisfying a desired property. A program starting from some state $\sigma \in \Sigma$ will reach a state $s \in S$ if and only if $\sigma \in \text{dom}(\tau_{[S]}^g)$.

In addition to guarantee properties, Urban and Miné [3] also introduced the *recurrence semantics*. A recurrence property guarantees that a program starting from some state $\sigma \in \Sigma$ will reach some state satisfying a given property infinitely often. The value assigned to a state by the *recurrence semantics* is an upper bound on the number of executions steps until the next state satisfying the property is reached. Recurrence properties can be expressed using the CTL formula $\forall\Box\forall\Diamond(a)$.

Definition 4.3. The *recurrence semantics* is a ranking function $\tau_{[S]}^r \in \Sigma \rightarrow \mathbb{N}$ where $S \subseteq \Sigma$ is a set of states satisfying a desired property. A program starting from some state $\sigma \in \Sigma$ will reach a state $s \in S$ infinitely often if and only if $\sigma \in \text{dom}(\tau_{[S]}^r)$.

Figure 3 shows an example for the semantics discussed in this section. We illustrate ranking functions by labeling the states with the corresponding value assigned to them by the function. The first example (a) shows the *termination semantics* for a state transition system that always terminates. Therefore the initial state has the value 2 assigned to it stating that this program terminates in at most two steps.

The second example (b) shows the *guarantee semantics* for the guarantee property that states that a gray state will be reached eventually ($\forall\Diamond(\text{gray})$). This holds for example (b), therefore the initial state has the value 2 assigned to it. The program reaches a gray state in at most two steps.

The last example (c) shows the *recurrence semantics* for the recurrence property that states that a gray state will be reached infinitely often ($\forall\Box\forall\Diamond(\text{gray})$). As one can see in the transition system, this is not true when starting from the initial state. Therefore the *recurrence semantics* is undefined for the initial state. However the property holds when starting from state c or d . Accordingly these two states have the values 1 and 0 assigned to them.

We refer to [5] for a detailed discussion of the various semantics presented in this section. The next section extends on the concepts of this section and presents a proof method for CTL properties.




Figure 3: Example *termination semantics* (a), *guarantee semantics* (b) and *recurrence semantics* (c)

5 Concrete Semantics for CTL

In Section 4 we introduced the concept of ranking functions and explained how they express the semantics of termination and also liveness properties in general. We recall that ranking functions are a proof method for liveness properties. The ranking function assigns well-ordered values to program states. Liveness properties state that some goal state will be reached eventually. The ranking function assigns 0 to those goal states and increased the values for states leading up to said goal state through backtracking. One can determine if a state satisfies a liveness property by checking if the corresponding ranking function assigns a value to said state. The assigned value is an upper bound on the number of steps until a goal state is reached. In this section we extend this proof method presented in [5] to CTL.

CTL can define both liveness and safety properties. For liveness properties, ranking functions can be used as proof method. However, ranking functions are not suitable in the case of safety properties since there is no goal state to be reached eventually. Because a CTL formula can arbitrarily combine liveness and safety properties we will use a *counting function* $\tau_\Phi \in \Sigma \rightarrow \mathbb{O}$ as proof method. The function τ_Φ is similar to a ranking function, it assigns well-ordered values to states. A state satisfies a CTL property Φ if it is in the domain of τ_Φ . However, as opposed to ranking functions, the value of the *counting function* τ_Φ is not guaranteed to decrease. It will only decrease if Φ is a liveness property. In that case, the value of the function is an upper bound on the number of steps until some goal states is reached. Otherwise, in the case of safety properties, the value may be constant and is irrelevant. We call τ_Φ the *CTL semantics* from now on.

Theorem 5.1. The *CTL semantics* for a given CTL formula Φ is a counting function $\tau_\Phi \in \Sigma \rightarrow \mathbb{N}$. It encodes the semantics of Φ for a given state transition system $\langle \Sigma, \tau \rangle$ such that $\sigma \models \Phi \iff \sigma \in \text{dom}(\tau_\Phi)$

We will define the *CTL semantics* inductively for each CTL operator such that arbitrary combinations of CTL properties can be expressed. Furthermore we split the definition into path-independent and path-dependent CTL operators.

5.1 Path Independent Operators

We start by defining the *CTL semantics* for atomic propositions and logic operators. These CTL properties are path independent and can be defined individually for each state $\sigma \in \Sigma$. The definitions are formed according to Theorem 5.1 and the satisfiability relation for CTL properties defined in Section 3.

Definition 5.1. Equations for path independent CTL operators

$$\tau_a \stackrel{\text{def}}{=} \lambda\sigma. \begin{cases} 0 & \text{if } \sigma \models a \\ \text{undefined} & \text{otherwise} \end{cases} \quad (3)$$

$$\tau_{\neg\Phi} \stackrel{\text{def}}{=} \lambda\sigma. \begin{cases} 0 & \text{if } \sigma \notin \text{dom}(\tau_\Phi) \\ \text{undefined} & \text{otherwise} \end{cases} \quad (4)$$

$$\tau_{\Phi_1 \wedge \Phi_2} \stackrel{\text{def}}{=} \lambda\sigma. \begin{cases} \sup\{\tau_{\Phi_1}(\sigma), \tau_{\Phi_2}(\sigma)\} & \text{if } \sigma \in \text{dom}(\tau_{\Phi_1}) \cap \text{dom}(\tau_{\Phi_2}) \\ \text{undefined} & \text{otherwise} \end{cases} \quad (5)$$

$$\tau_{\Phi_1 \vee \Phi_2} \stackrel{\text{def}}{=} \lambda\sigma. \begin{cases} \sup\{\tau_{\Phi_1}(\sigma), \tau_{\Phi_2}(\sigma)\} & \text{if } \sigma \in \text{dom}(\tau_{\Phi_1}) \cap \text{dom}(\tau_{\Phi_2}) \\ \tau_{\Phi_1}(\sigma) & \text{if } \sigma \in \text{dom}(\tau_{\Phi_1}) \setminus \text{dom}(\tau_{\Phi_2}) \\ \tau_{\Phi_2}(\sigma) & \text{if } \sigma \in \text{dom}(\tau_{\Phi_2}) \setminus \text{dom}(\tau_{\Phi_1}) \\ \text{undefined} & \text{otherwise} \end{cases} \quad (6)$$

Equation 3 assigns 0 to all states that satisfy the atomic proposition a . In case of liveness properties, such as $\forall\Diamond a$, the value 0 marks that the goal of the property has been reached. For safety properties, the value 0 simply states that the states satisfies the property.

The logic \neg operator in equation 4 follows the same approach as for atomic propositions. It interprets $\neg\Phi$ as an atomic proposition and uses the fact that $\sigma \notin \text{dom}(\tau_\Phi) \implies \sigma \models \neg\Phi$ which follows from Theorem 5.1.

The logical \wedge and \vee connectives in equations 5 and 6 reuse the values of the underlying functions τ_{Φ_1} and τ_{Φ_2} according to the semantics of these operators. If a state is assigned a value by both functions, then the supremum of the two values is used. That way, if the two underlying properties Φ_1 and Φ_2 are liveness-properties, we preserve the notion of increasing the values of the function when backtracking from the goal states of underlying properties.

Lemma 5.2. The *CTL semantics* for path-independent CTL operators are sound and

complete. Let $\sigma \in \Sigma$ and Φ, Φ_1 and Φ_2 be arbitrary CTL properties.

$$\sigma \models a \iff \sigma \in \text{dom}(\tau_a) \quad (7)$$

$$\sigma \models \neg\Phi \iff \sigma \in \text{dom}(\tau_{\neg\Phi}) \quad (8)$$

$$\sigma \models \Phi_1 \wedge \Phi_2 \iff \sigma \in \text{dom}(\tau_{\Phi_1 \wedge \Phi_2}) \quad (9)$$

$$\sigma \models \Phi_1 \vee \Phi_2 \iff \sigma \in \text{dom}(\tau_{\Phi_1 \vee \Phi_2}) \quad (10)$$

5.2 Path Dependent Operators

CTL semantics for path-dependent operators `until` and `global` are defined in terms of fixed-points. These fixed-points are defined over the partially ordered set of functions $\langle \Sigma \rightarrow \mathbb{N}, \sqsubseteq \rangle$. Fixed-point iterates are related to each other using the *computational order* \sqsubseteq . This partial order relates functions in terms of expressiveness, i.e., for how many states can a function prove that the CTL property holds.

Definition 5.2. Let $f, g \in \Sigma \rightarrow \mathbb{O}$. The *computational order* \sqsubseteq is defined as follows.

$$f \sqsubseteq g \iff \text{dom}(f) \subseteq \text{dom}(g) \wedge \forall x \in \text{dom}(f) : f(x) \leq g(x)$$

The *CTL semantics* are not computable in general. In Section 8 we will present a sound decidable abstraction of the *CTL semantics*. Soundness is related to the *approximation order* \preceq .

Definition 5.3. Let $f, g \in \Sigma \rightarrow \mathbb{O}$. The *approximation order* \preceq is defined as follows.

$$f \preceq g \iff \text{dom}(f) \supseteq \text{dom}(g) \wedge \forall x \in \text{dom}(g) : f(x) \leq g(x)$$

The *approximation order* ranks counting functions in terms of precision. A function f is more precise than a function g if it is defined over more states than g and if the value is smaller for all states in the domain of g . Intuitively, a more precise counting function is able to prove for more states that a CTL Φ holds or does not hold.

Until

Recall that for the CTL property $\forall(\Phi_1 U \Phi_2)$ to hold for some state $\sigma \in \Sigma$, all paths starting from said state must form a chain of states satisfying Φ_1 ending in a state satisfying Φ_2 . In case of $\exists(\Phi_1 U \Phi_2)$ at least on such path must exist. The *CTL semantics* for universal and existential until properties are defined as least fixed-points of the abstract transformers

$$\phi_{\forall(\Phi_1 U \Phi_2)} \in (\Sigma \rightarrow \mathbb{N}) \rightarrow (\Sigma \rightarrow \mathbb{N})$$

$$\phi_{\exists(\Phi_1 U \Phi_2)} \in (\Sigma \rightarrow \mathbb{N}) \rightarrow (\Sigma \rightarrow \mathbb{N})$$

starting from the totally undefined counting function $\dot{\emptyset}$.

Definition 5.4. *CTL semantics* for existential and universal until properties

$$\tau_{\forall(\Phi_1 U \Phi_2)} \stackrel{\text{def}}{=} \text{lfp}_{\emptyset}^{\sqsubseteq} \phi_{\forall(\Phi_1 U \Phi_2)} \quad (11)$$

$$\phi_{\forall(\Phi_1 U \Phi_2)} f \stackrel{\text{def}}{=} \lambda \sigma. \begin{cases} 0 & \text{if } \sigma \in \text{dom}(\tau_{\Phi_2}) \\ \sup\{f(\sigma') + 1 \mid \langle \sigma, \sigma' \rangle \in \tau\} & \text{if } \sigma \notin \text{dom}(\tau_{\Phi_2}) \wedge \\ & \sigma \in \text{dom}(\tau_{\Phi_1}) \wedge \\ & \sigma \in \widetilde{\text{pre}}(\text{dom}(f)) \\ \text{undefined} & \text{otherwise} \end{cases} \quad (12)$$

$$\tau_{\exists(\Phi_1 U \Phi_2)} \stackrel{\text{def}}{=} \text{lfp}_{\emptyset}^{\sqsubseteq} \phi_{\exists(\Phi_1 U \Phi_2)} \quad (13)$$

$$\phi_{\exists(\Phi_1 U \Phi_2)} f \stackrel{\text{def}}{=} \lambda \sigma. \begin{cases} 0 & \text{if } \sigma \in \text{dom}(\tau_{\Phi_2}) \\ \sup\{f(\sigma') + 1 \mid \langle \sigma, \sigma' \rangle \in \tau\} & \text{if } \sigma \notin \text{dom}(\tau_{\Phi_2}) \wedge \\ & \sigma \in \text{dom}(\tau_{\Phi_1}) \wedge \\ & \sigma \in \text{pre}(\text{dom}(f)) \\ \text{undefined} & \text{otherwise} \end{cases} \quad (14)$$

This definition is a generalization of the *guarantee semantics* presented in [3]. The fixed-point iteration starts by assigning the value 0 to all states that satisfy Φ_2 . In subsequent iterations we consider all states that satisfy Φ_1 and from which one can only transition to states that already satisfy $\forall(\Phi_1 U \Phi_2)$. These states are then assigned the largest ranking value of all reachable states plus one. By performing iterations this way, we backtrack paths in the state transition systems that end in a state satisfying Φ_2 and which are preceded by an unbroken chain of states satisfying Φ_1 . Every state on such a path is guaranteed to satisfy $\forall(\Phi_1 U \Phi_2)$. Furthermore, by starting from 0 at states that satisfy Φ_2 and incrementing the value of the function while backtracking, we construct a counting function such that the value assigned to each state is an upper bound on the number of steps until a state satisfying Φ_2 is reached.

The $\widetilde{\text{pre}}$ relation guarantees that during the backtracking, only those states that exclusively transition to states satisfying $\forall(\Phi_1 U \Phi_2)$, are considered. This condition can be relaxed for existential until properties by using the *pre* relation instead. That way, states that have at least one reachable state satisfying $\exists(\Phi_1 U \Phi_2)$ are also considered during the backtracking (see Definitions 2.2 and 2.3).

Lemma 5.3. The *CTL semantics* for until properties are sound and complete. Let $\sigma \in \Sigma$ and Φ_1, Φ_2 be arbitrary CTL properties.

$$\sigma \models \forall(\Phi_1 U \Phi_2) \iff \sigma \in \text{dom}(\tau_{\forall(\Phi_1 U \Phi_2)}) \quad (15)$$

$$\sigma \models \exists(\Phi_1 U \Phi_2) \iff \sigma \in \text{dom}(\tau_{\exists(\Phi_1 U \Phi_2)}) \quad (16)$$




Figure 4: Iterative computation of $\tau_{\forall(\text{gray } U \text{ black})}$.

Figures 4 and 5 give an example on how the iterative computation for until properties works for universal and existential quantifiers. Note how Figure 5 has one additional iteration because of the existential quantifier. The initial state is added to the function in the last iteration because there exists one edge that leads to a state satisfying the property. For the universal property, the iteration stops after three iterations because not all successor states of the initial state satisfy the property.

Global

Recall that the CTL global operator states that some property must hold globally, i.e., indefinitely for *all* paths starting from some state in the case of the universal quantifier ($\forall \square \Phi$) or for *some* paths in case of the existential quantifier ($\exists \square \Phi$). The CTL semantics for global properties are defined as greatest fixed-point of the abstract transformers

$$\phi_{\forall(\Phi_1 U \Phi_2)} \in (\Sigma \rightarrow \mathbb{N}) \rightarrow (\Sigma \rightarrow \mathbb{N})$$

$$\phi_{\exists(\Phi_1 U \Phi_2)} \in (\Sigma \rightarrow \mathbb{N}) \rightarrow (\Sigma \rightarrow \mathbb{N})$$

starting from the CTL semantics τ_{Φ} of the inner CTL property.




Figure 5: Iterative computation of $\tau_{\exists(\text{gray} \cup \text{black})}$.

Definition 5.5. Equations for CTL global operator

$$\tau_{\forall\Box\Phi} \stackrel{\text{def}}{=} \text{gfp}_{\tau_{\Phi}}^{\Xi} \phi_{\forall\Box\Phi} \quad (17)$$

$$\phi_{\forall\Box\Phi} f \stackrel{\text{def}}{=} \lambda\sigma. \begin{cases} f(x) & \text{if } \sigma \in \widetilde{\text{pre}}(\text{dom}(f)) \\ \text{undefined} & \text{otherwise} \end{cases} \quad (18)$$

$$\tau_{\exists\Box\Phi} \stackrel{\text{def}}{=} \text{gfp}_{\tau_{\Phi}}^{\Xi} \phi_{\exists\Box\Phi} \quad (19)$$

$$\phi_{\exists\Box\Phi} f \stackrel{\text{def}}{=} \lambda\sigma. \begin{cases} f(x) & \text{if } \sigma \in \text{pre}(\text{dom}(f)) \\ \text{undefined} & \text{otherwise} \end{cases} \quad (20)$$

This definition is based on the *recurrence semantics* presented in [3]. As with the *until* operator, we distinguish between universal and existential properties by using either $\widetilde{\text{pre}}$ or pre . The fixed-point iteration starts with the counting function τ_{Φ} of the inner property Φ . At each iteration, every state that is still part of the domain of the function is inspected. The inspected state is kept in the domain of the function if *all* its successor states (or *some* for the existential case) are also part of the domain of the function, otherwise it is removed. That way, only states which are part of an infinite path consisting exclusively of states satisfying Φ are kept in the domain of the function.

Lemma 5.4. The *CTL semantics* for global properties are sound and complete. Let




Figure 6: Iterative computation of $\tau_{\forall\Box\Phi}$.



Figure 7: Iterative computation of $\tau_{\exists\Box\Phi}$.

$\sigma \in \Sigma$ and Φ be an arbitrary CTL property.

$$\sigma \models \forall\Box\Phi \iff \sigma \in \text{dom}(\tau_{\forall\Box\Phi}) \quad (21)$$

$$\sigma \models \exists\Box\Phi \iff \sigma \in \text{dom}(\tau_{\exists\Box\Phi}) \quad (22)$$

Figures 6 and 7 show this for $\tau_{\exists\Box\Phi}$ and $\tau_{\forall\Box\Phi}$. Both iterations start with some initial counting function τ_Φ . In the first iteration state b is removed because it has no outgoing edges. For the existential case, the iteration stops here because all remaining states a , c and d have at least one edge to a node that is part of the function. In the universal case we get an additional iteration that removes state a because not all of its successor nodes (namely b) are part of the function. Note that only infinite paths are considered to hold globally.

Next

The next operator is path dependent but does not require fixed-point iterations. A state satisfies $\forall\bigcirc\Phi$ if all its immediate successors satisfy the property Φ , correspondingly $\exists\bigcirc\Phi$ is satisfied if at least one immediate successor satisfies the property Φ . This corresponds to the definition of the \widetilde{pre} and pre relations. Zero is assigned to each state that satisfies the property to construct a valid counting function according to Theorem 5.1.

Definition 5.6. Equations for CTL next operator

$$\tau_{\forall\bigcirc\Phi} \stackrel{\text{def}}{=} \lambda\sigma. \begin{cases} 0 & \text{if } \sigma \in \widetilde{\text{pre}}(\text{dom}(\tau_{\Phi})) \\ \text{undefined} & \text{otherwise} \end{cases} \quad (23)$$

$$\tau_{\exists\bigcirc\Phi} \stackrel{\text{def}}{=} \lambda\sigma. \begin{cases} 0 & \text{if } \sigma \in \text{pre}(\text{dom}(\tau_{\Phi})) \\ \text{undefined} & \text{otherwise} \end{cases} \quad (24)$$

Lemma 5.5. The *CTL semantics* for next properties are sound and complete. Let $\sigma \in \Sigma$ and Φ be an arbitrary CTL property.

$$\sigma \models \forall\bigcirc\Phi \iff \sigma \in \text{dom}(\tau_{\forall\bigcirc\Phi}) \quad (25)$$

$$\sigma \models \exists\bigcirc\Phi \iff \sigma \in \text{dom}(\tau_{\exists\bigcirc\Phi}) \quad (26)$$

6 Imperative Language

In this section we briefly introduce a minimal imperative programming language. It will be used in Section 8 to define the abstract CTL semantics. The language has no procedures, pointers or recursion and is non-deterministic. Variables are integer valued and statically allocated. We chose a simple language to keep the definitions of the abstract CTL semantics simple and concise. However, the implementation of our static analyzer will support a more feature-rich programming language (see Section 9).

First we define the syntax for arithmetic and boolean expressions. The syntax definitions are based on Chapter 3 of [5].

Definition 6.1. Syntax for arithmetic and boolean expressions.

Arithmetic expressions are defined over a set of variables \mathcal{X} .

$$\begin{aligned} aexp ::= & X && X \in \mathcal{X} \\ & | [i_1, i_2] && i_1 \in \mathbb{Z} \cup \{-\infty\}, i_2 \in \mathbb{Z} \cup \{\infty\}, i_1 \leq i_2 \\ & | -aexp \\ & | aexp \diamond aexp && \diamond \in \{+, -, *, /\} \end{aligned}$$

$$\begin{aligned} bexp ::= & ? && \text{non-deterministic choice} \\ & | \neg bexp \\ & | bexp \wedge bexp \\ & | bexp \vee bexp \\ & | aexp \diamond aexp && \diamond \in \{<, \leq, >, \geq\} \end{aligned}$$

The semantics for expressions are defined as expected. Please refer to [5] for a formal definition. Note that the symbol ? stands for non-deterministic choice.

Programs are defined in terms of control-flow-graphs. The control-flow-graph of a program models all possible program executions as paths in the graph.

Definition 6.2. Program representation as control-flow-graph. A control-flow-graph is a tuple (\mathcal{L}, E) . \mathcal{L} is the set of program labels, also called nodes. $E \subseteq (\mathcal{L} \times stmt \times \mathcal{L})$ is the set of edges in the control-flow-graph where $stmt$ is the set of all program statements.

$$\begin{aligned}
 stmt ::= & \text{ skip} \\
 & | bexp \\
 & | X := aexp \qquad X \in \mathcal{X}
 \end{aligned}$$

Every control point of a program is assigned a label $l \in \mathcal{L}$. The nodes in the control-flow-graph correspond to those labels. An edge $(u, s, v) \in E$ states that one can transition from node u to v by executing statement s . The `skip` statement transitions from one node to another without doing anything, the boolean expression `bexp` limits the set of states that are allowed to transition to the next node to the one satisfying `bexp` and the assignment $X := aexp$ assigns the value of the arithmetic expression `aexp` to the variable X . A program consists of a control-flow-graph and two special labels l_{entry} and l_{exit} that denote the program entry and exit point.

We introduce the following auxiliary functions on nodes of a control-flow-graph to refer to the incoming and outgoing edges of a node.

Definition 6.3. Given a control-flow-graph (\mathcal{L}, E) and some node $l \in \mathcal{L}$

$$\begin{aligned}
 in(l) &\stackrel{\text{def}}{=} \{(u, s, v) \in E \mid v = l\} \\
 out(l) &\stackrel{\text{def}}{=} \{(u, s, v) \in E \mid u = l\}
 \end{aligned}$$

7 Decision Tree Abstract Domain

This section briefly recaps the decision tree abstract domain [5]. Decision trees encode piecewise-defined linear functions which are used as an abstraction of the counting functions introduced in Section 4. First we give a description of the decision tree abstract domain. Then we introduce ordering relations between the elements of the domain and relevant operations on the elements of the domain. An in-depth description of the topics covered in this section can be found in [5].




Figure 8: Example for decision tree

7.1 Domain

The elements of the abstract domain are binary decision trees. The nodes of the trees are linear constraints and the leafs are linear functions of the program variables. Decision trees partition the state space, given by a set of variables \mathcal{X} , into partitions. Each partition is defined through the conjunction of linear constraints on the path from root to leaf in the decision tree. The linear function at the leaf determines the value of the function for the corresponding partition of the state space.

Figure 8 gives an example for such a decision tree. It consists of two nodes with linear constraints $x \leq 0$ and $y \leq 0$. The left most leaf is the function $\lambda x y. x + y$. It is defined for all states satisfying $x \leq 0 \wedge y \leq 0$ according to the constraints from root to leaf. The right most leaf is the function $\lambda x. 2x$, it is defined for all states satisfying $\neg(x \leq 0)$ (following the right child of a node negates the linear constraint). The leaf in the middle is a bottom node, signifying that the function for the corresponding partition is undefined. In summary the decision tree in Figure 8 is equivalent to the following partial function:

$$f(x, y) = \begin{cases} x + y & \text{if } x \leq 0 \wedge y \leq 0 \\ 2x & \text{if } x > 0 \\ \text{undefined} & \text{otherwise} \end{cases}$$

We will now formalize the decision tree abstract domain.

Constraints

The constrains at the inner nodes of the decision tree are elements of the *linear constraints auxiliary abstract domain* \mathcal{C} .

$$\mathcal{C} \stackrel{\text{def}}{=} \left\{ c_1 X_1 + \dots + c_n X_n + c_{n+1} \geq 0 \left| \begin{array}{l} \mathcal{X} = \{X_1, \dots, X_n\} \\ c_1, \dots, c_n, c_{n+1} \in \mathbb{Z} \\ \text{gcd}(|c_1|, \dots, |c_n|, |c_{n+1}|) = 1 \end{array} \right. \right\}$$

Elements of \mathcal{C} can be instances of the *interval abstract domain* [8], the *octagon abstract domain* [9] or the *polyhedra abstract domain* [10].

Functions

Leaves of the decision trees are elements of the *functions auxiliary abstract domain* \mathcal{F} . Elements of \mathcal{F} are either natural valued linear functions or one of the two special elements $\top_{\mathcal{F}}$ or $\perp_{\mathcal{F}}$. The element $\perp_{\mathcal{F}}$ indicates that the function is undefined on the given partition. The element $\top_{\mathcal{F}}$ indicates that the value of the function is unknown for the given partition.

$$\mathcal{F} \stackrel{\text{def}}{=} \{\mathbb{Z}^{|\mathcal{X}|} \rightarrow \mathbb{N}\} \cup \{\top_{\mathcal{F}}, \perp_{\mathcal{F}}\}$$

We will write constant functions that return a value $n \in \mathbb{N}$ by just stating the constant value, e.g., $0 \in \mathcal{F}$ denotes the constant function that returns 0 for every state.

In the following sections we will distinguish between so called *defined* and *undefined* leaves. A leaf $f \in \mathcal{F}$ is called *defined* if f is neither $\top_{\mathcal{F}}$ nor $\perp_{\mathcal{F}}$ and *undefined* otherwise. Defined leaves assign an actual value to its partition, therefore the function that the decision tree represents is defined for that partition. Otherwise the function is undefined.

Decision Trees

We now define the decision tree abstract domain \mathcal{T} . An element $t \in \mathcal{T}$ is either a *leaf node* **LEAF**: f consisting of a function $f \in \mathcal{F}$ (denoted $t.f$), or a *decision node* **NODE**{ c } : $l; r$ consists of a linear constraint $c \in \mathcal{C}$ (denoted $t.c$) and a left and a right sub tree $l, r \in \mathcal{T}$ (denoted $t.l$ and $t.r$).

$$\mathcal{T} \stackrel{\text{def}}{=} \{\text{LEAF}: f \mid f \in \mathcal{F}\} \cup \{\text{NODE}\{c\} : l; r \mid c \in \mathcal{C}, l, r \in \mathcal{T}\}$$

For algorithmic purposes we also define \mathcal{T}_{NIL} . This adds an additional leaf element **NIL** to \mathcal{T} to represent the absence of information about a partition. **NIL** leaves usually appear if a partition in a decision tree can be excluded because it is infeasible w.r.t. the program execution.

$$\mathcal{T}_{\text{NIL}} \stackrel{\text{def}}{=} \{\text{NIL}\} \cup \{\text{LEAF}: f \mid f \in \mathcal{F}\} \cup \{\text{NODE}\{c\} : l; r \mid c \in \mathcal{C}, l, r \in \mathcal{T}_{\text{NIL}}\}$$

Sound Abstractions

Decision trees are an abstraction of counting functions. A sound abstraction is defined in terms of the *approximation order* \preceq (see Definition 5.3).

Definition 7.1. Let $\gamma \in \mathcal{T} \rightarrow (\Sigma \rightarrow \mathbb{O})$ be the concretization function from decision trees to counting functions. A decision tree $t \in \mathcal{T}$ is a sound abstraction of a counting function $f \in \Sigma \rightarrow \mathbb{N}$ if $f \preceq \gamma(t)$.

This definitions allows us to soundly define computable abstractions of counting functions. Every decision tree computed in such a way captures the behavior of the concrete counting function for all defined partitions of the tree.




Figure 9: Hasse diagrams for $\sqsubseteq_{\mathcal{F}}$ and $\preceq_{\mathcal{F}}$

Orders

We now define the *computational order* $\sqsubseteq_{\mathcal{T}}$ and *approximation order* $\preceq_{\mathcal{T}}$ over elements of \mathcal{T} . These orders are an approximation of the corresponding orders presented in Section 5 (see Definitions 5.2 and 5.3).

Both orders are defined by leaf-wise comparison of two decision trees. To that end we define the *computational order* $\sqsubseteq_{\mathcal{F}}$ and *approximation order* $\preceq_{\mathcal{F}}$ for elements of \mathcal{F} . Two trees are related to each other if all leafs are pairwise related w.r.t. $\sqsubseteq_{\mathcal{F}}$ ($\preceq_{\mathcal{F}}$, respectively). Please refer to [5] for a detailed explanation.

The two orders $\sqsubseteq_{\mathcal{F}}$ and $\preceq_{\mathcal{F}}$ are identical for function values $f \in \mathbb{Z}^{|\mathcal{X}|} \rightarrow \mathbb{N}$. Pairings with the special elements $\top_{\mathcal{F}}$ and $\perp_{\mathcal{F}}$ are related to each other according to the Hasse diagrams in Figure 9.

Definition 7.2. The *computation order* $\sqsubseteq_{\mathcal{F}}$ and *approximation order* $\preceq_{\mathcal{F}}$ for elements $f_1, f_2 \in \mathcal{F} \setminus \{\top_{\mathcal{F}}, \perp_{\mathcal{F}}\}$ are defined as follows.

$$f_1 \sqsubseteq_{\mathcal{F}} f_2 \iff f_1 \preceq_{\mathcal{F}} f_2 \iff \forall x \in \mathbb{Z}^{|\mathcal{X}|}: f_1(x) \leq f_2(x)$$

7.2 Join

Two trees can be joined to form the union of all partitions represented by the two trees. When joining two trees, they are first reshaped such that both trees consist of the same partitions. They only differ in the values of the leafs. Then the two trees can be joined leaf-wise. There are two join variations. The *computational join* $\sqcup_{\mathcal{T}}: (\mathcal{T}_{NIL} \times \mathcal{T}_{NIL}) \rightarrow \mathcal{T}_{NIL}$ and the *approximation join* $\vee_{\mathcal{T}}: (\mathcal{T}_{NIL} \times \mathcal{T}_{NIL}) \rightarrow \mathcal{T}_{NIL}$. Two leafs l_1 and l_2 are joined by taking the least upper bound of the two functions $l_1.f$ and $l_2.f$ w.r.t. $\sqsubseteq_{\mathcal{F}}$ ($\preceq_{\mathcal{F}}$, respectively). Figure 10 demonstrates the difference between the two join types. When joining two trees where one leaf is defined and one is undefined (see left leaf in t_1 and t_2), the *computational join* will preserve the defined leaf and the *approximation join* will make the leaf undefined. If one of the two leafs is NIL then they are joined by taking the one that is not NIL. When both are NIL then the result of joining the leafs is NIL.




Figure 10: Decision Tree Join Example

7.3 Meet

The meet operator intersects the partitions of two decision trees. As with the join, both trees are first brought into the same shape such that they can be combined leaf-wise. There are two meet variations. The *computational meet* $\sqcap_{\mathcal{T}}: (\mathcal{T}_{\text{NIL}} \times \mathcal{T}_{\text{NIL}}) \rightarrow \mathcal{T}_{\text{NIL}}$ and the *approximation meet* $\wedge_{\mathcal{T}}: (\mathcal{T}_{\text{NIL}} \times \mathcal{T}_{\text{NIL}}) \rightarrow \mathcal{T}_{\text{NIL}}$. Both combine defined leaves using the least upper bound w.r.t. $\preceq_{\mathcal{F}}$. If at least one of the two leaves is NIL, then the result is $\perp_{\mathcal{F}}$ in case of the *computational meet* and NIL in case of the *approximation meet*.

7.4 Widening and Dual Widening

The *widening* operator $\nabla_{\mathcal{T}}: (\mathcal{T}_{\text{NIL}} \times \mathcal{T}_{\text{NIL}}) \rightarrow \mathcal{T}_{\text{NIL}}$ is used to enforce convergence when computing increasing sequences of values in the decision tree abstract domain. Once the sequence is stable for the *computational order* $\sqsubseteq_{\mathcal{T}}$, the limit of this sequence is guaranteed to be a sound approximation of the corresponding concrete counting function w.r.t. the *approximation order* $\preceq_{\mathcal{F}}$. The dual of this concept for decreasing sequences is called *dual widening* $\bar{\nabla}_{\mathcal{T}}: (\mathcal{T}_{\text{NIL}} \times \mathcal{T}_{\text{NIL}}) \rightarrow \mathcal{T}_{\text{NIL}}$. Intuitively, the widening operator tries to extrapolate the function values of decision trees to linear functions. This extrapolation is initially a guess. If the guess proves to be wrong, it is corrected by setting the corresponding partition to $\top_{\mathcal{T}}$ in future iterations. A thorough discussion about widening for elements of the decision tree abstract domain can be found in Section 5.2.4 of [5].

7.5 Filter

The filter operator $\text{FILTER}_o[[bexp]]: \mathcal{T}_{\text{NIL}} \rightarrow \mathcal{T}_{\text{NIL}}$ prunes all partitions of a decision tree that do not satisfy the given boolean expression $bexp$. Leafs are pruned by replacing them with NIL. Partitions are pruned using over-approximation. That means that the resulting tree can still be defined for states that do not satisfy the boolean expression. It is however guaranteed that all states that satisfy the boolean expression remain defined in the tree. In addition to the over-approximating version there is also the under-approximating version $\text{FILTER}_u[[bexp]]: \mathcal{T}_{\text{NIL}} \rightarrow \mathcal{T}_{\text{NIL}}$. Here the resulting decision tree is

guaranteed to not contain any partitions that do not satisfy the boolean expressions. However states that do satisfy it might be removed if the underlying numerical domain is not expressive enough.

7.6 Backward Assign

The operator $\text{ASSIGN}_o[[X := \text{aexp}]]: \mathcal{T}_{\text{NIL}} \rightarrow \mathcal{T}_{\text{NIL}}$ handles the backward assignment of the arithmetic expression aexp to variable X . The linear constraints of the decision tree nodes and the functions at the leafs are adjusted accordingly. ASSIGN_o uses over-approximation on the underlying numerical domains. As with the filter operator, there also exists an under-approximating version $\text{ASSIGN}_u[[X := \text{aexp}]]: \mathcal{T}_{\text{NIL}} \rightarrow \mathcal{T}_{\text{NIL}}$.

8 Abstract Semantics for CTL

The counting function τ_Φ is in general not computable. In this section we present a sound and computable approximation of the *CTL semantics* τ_Φ defined in Section 5. We approximate τ_Φ by using the decision tree abstract domain (Section 7) to approximate counting functions in terms of piecewise defined ranking functions.

Theorem 8.1. The abstract CTL semantics $\tau_\Phi^\sharp \in \mathcal{L} \rightarrow \mathcal{T}$ is a sound approximation of the CTL semantics τ_Φ with regards to the *approximation order* $\preceq_{\mathcal{T}}$ (see Definition 7.1).

Recall that the CTL semantics $\tau_\Phi \in \Sigma \rightarrow \mathbb{N}$ is a partial function that assigns numerical values to program states $\sigma \in \Sigma$. In the abstract version, program states are grouped by program labels $l \in \mathcal{L}$ and partitioned by decision trees $t \in \mathcal{T}$. A program satisfies a given CTL property Φ if the decision tree of the initial program label $\tau_\Phi^\sharp(t_{init})$ is defined over all partitions of program states, i.e., all leafs of the decision tree are defined.

The following sections present how to compute τ_Φ^\sharp for each CTL operator. We start with the basic operators \wedge, \vee, \neg and atomic propositions. These can be computed directly for each program label. Then we present how to compute the universal $\forall(\cdot U \cdot), \forall \bigcirc$ and $\forall \square$ operators through fixed-point iteration. Followed by a discussion on how to adapt the universal operators to their existential version. Note that the abstract CTL semantics are computed recursively. The recursion stops at atomic propositions.

8.1 Path Independent Operators

Atomic propositions are path independent, therefore τ_a^\sharp assigns the same decision tree to each program label $l \in \mathcal{L}$. This decision tree assigns the constant function 0 to all partitions that satisfy the atomic proposition a . We compute this tree by using the $\text{RESET}[[a]]$ operator on the totally undefined decision tree $\perp_{\mathcal{T}}$. The $\text{RESET}[[a]]: \mathcal{T} \rightarrow \mathcal{T}$ operator takes as input a decision tree and returns a copy of said tree where every partition that satisfies a is replaced with the constant function 0. We refer the reader to [5] for a detailed description of the RESET operator. Note however, that the implementation

of `RESET` in [5] has a small error that can lead to unsoundness. This is discussed in the excursion below titled ‘`RESET` and over-approximation’.

Definition 8.1. Abstract CTL semantics for atomic propositions

$$\tau_a^\# \stackrel{\text{def}}{=} \lambda l. \text{RESET}[[a]] \sqcup_{\mathcal{T}} \quad (27)$$

Lemma 8.2. The abstract CTL semantics $\tau_a^\#$ is a sound approximation of the CTL semantics τ_a with regards to the *approximation order* $\preceq_{\mathcal{T}}$ (see Definition 7.1).

RESET and over-approximation

The `RESET`[[a]] operator was originally introduced in [5] in the context of abstract guarantee semantics and would over-approximate the set of partitions that satisfy the atomic proposition a . During the work on this thesis, we discovered that this original definition is actually unsound and leads to incorrect analysis results.

The problem is best described using an example. Consider the abstract CTL semantics $\tau_{x^2 < y^3 + 1}^\#$. The non-linear constraint $x^2 < y^3 + 1$ can usually not be represented by any of the commonly used numerical domains. An over-approximating implementation of `RESET`[[$x^2 < y^3 + 1$]] will therefore reset some pairs (x, y) for which $x^2 < y^3 + 1$ does not hold which is unsound as to the definition of the CTL semantics $\tau_{x^2 < y^3 + 1}$. Note that this problem propagates to more complex temporal properties that depend on atomic propositions.

We resolve this problem by using under-approximating on the underlying numerical domains to soundly determine which partitions satisfy the atomic proposition a .

Now we define the abstract CTL semantics for the logical operators \wedge , \vee and \neg .

Definition 8.2. Abstract CTL semantics for logic operators

$$\tau_{\Phi_1 \wedge \Phi_2}^\# \stackrel{\text{def}}{=} \lambda l. (\tau_{\neg \Phi_1}^\# l) \sqcup_{\mathcal{T}} (\tau_{\neg \Phi_2}^\# l) \quad (28)$$

$$\tau_{\Phi_1 \vee \Phi_2}^\# \stackrel{\text{def}}{=} \lambda l. (\tau_{\neg \Phi_1}^\# l) \sqcap_{\mathcal{T}} (\tau_{\neg \Phi_2}^\# l) \quad (29)$$

$$\tau_{\neg \Phi}^\# \stackrel{\text{def}}{=} \lambda l. \text{COMPLEMENT} (\tau_{\Phi}^\# l) \quad (30)$$

The abstract CTL semantics for \wedge and \vee (Equations 28 and 29) combine the decision trees of the nested properties piecewise for each program label.

The *computational join* $\sqcup_{\mathcal{T}}$ is used to combine the two trees (see Section 7.2) in case of the logical \vee . This operator forms the union of the two decision trees. Note that we use the computational version of the join operator to include partitions that are defined in at least one of the two trees. If a partition is defined in both trees, the least upper

Algorithm 1 Tree Complement

```
function COMPLEMENT( $t$ )  $\triangleright t \in \mathcal{T}_{NIL}$ 
  if ( $isNode(t) \wedge t.f = \top$ )  $\vee isNil(t)$  then
    return  $t$   $\triangleright$  ignore  $\top$  and  $NIL$ 
  else if  $isLeaf(t) \wedge t.f = \perp$  then
    return  $LEAF : 0$   $\triangleright$  undefined becomes defined
  else if  $isLeaf(t)$  then
    return  $LEAF : \perp$   $\triangleright$  defined becomes undefined
  else
     $l \leftarrow COMPLEMENT(t.l)$ 
     $r \leftarrow COMPLEMENT(t.r)$ 
    return  $NODE\{t.c\} : l; r$ 
  end if
end function
```

bound of the two functions assigned to that partition is used w.r.t. to the partial order $\sqsubseteq_{\mathcal{F}}$ given in Definition 7.2.

The corresponding definition for the logical \wedge operator forms the intersection of the two trees by using the *computational meet* $\sqcap_{\mathcal{T}}$ (see Section 7.3). By using the computational version of the meet, we ensure that no NIL leaves are introduced when forming the intersection. As with the logical \vee , the least upper bound w.r.t. to $\sqsubseteq_{\mathcal{F}}$ is used if a partition is defined in both trees.

For the logical \neg operator we introduce the $COMPLEMENT: \mathcal{T}_{NIL} \rightarrow \mathcal{T}_{NIL}$ operator. This operator replaces all defined leaves with a \perp -leaf and all \perp -leaves with the constant function 0. By doing so, all states that originally satisfied the property do not satisfy it any more and vice versa. However one has to be careful when changing a partition from undefined to defined. Decision trees are an approximation of the concrete CTL semantics. Therefore not all states that are undefined in the abstract decision tree are actually undefined in the concrete counting function. Partitions that are undefined because of this uncertainty are marked with a \top -leaf. To ensure soundness, these leaves have to be ignored when forming the complement of a decision tree. The $COMPLEMENT$ operator is implemented in Algorithm 1.

Lemma 8.3. The abstract CTL semantics $\tau_{\Phi_1 \wedge \Phi_2}^{\#}$, $\tau_{\Phi_1 \vee \Phi_2}^{\#}$ and $\tau_{\neg \Phi}^{\#}$ are a sound approximation of the CTL semantics $\tau_{\Phi_1 \wedge \Phi_2}$, $\tau_{\Phi_1 \vee \Phi_2}$ and $\tau_{\neg \Phi}$ with regards to the *approximation order* $\preceq_{\mathcal{T}}$ (see Definition 7.1).

8.2 Path Dependent Operators

In this section we describe how the abstract CTL semantics for the path-dependent operators ($\forall \bigcirc \Phi$, $\exists \bigcirc \Phi$, $\forall(\Phi_1 U \Phi_2)$, $\exists(\Phi_1 U \Phi_2)$, $\forall \square \Phi$, $\exists \square \Phi$) are defined.

First, we define the two functions $\llbracket stmt \rrbracket_o \in \mathcal{T}_{NIL} \rightarrow \mathcal{T}_{NIL}$ and $\llbracket stmt \rrbracket_u \in \mathcal{T}_{NIL} \rightarrow \mathcal{T}_{NIL}$. The first one uses over-approximation on the underlying numerical domains, the second one under-approximation.

Both functions implement the effect of backward propagating an edge in the control-flow-graph, i.e., the effect of executing a statement. Assume that we have computed a decision tree for the target node of some edge. This decision tree represents the value of the counting function for this node. By applying $\llbracket \cdot \rrbracket$ to this tree, we compute the decision tree that holds before executing the statement, i.e. the value of the function at the source node of this edge.

Definition 8.3. Abstract semantics for basic statements

$$\begin{aligned} \llbracket \text{skip} \rrbracket_o &\stackrel{\text{def}}{=} \lambda t. \text{STEP}(t) \\ \llbracket bexp \rrbracket_o &\stackrel{\text{def}}{=} \lambda t. \text{ASSIGN}_o(t) \\ \llbracket X := aexp \rrbracket_o &\stackrel{\text{def}}{=} \lambda t. \text{FILTER}_o(t) \\ \\ \llbracket \text{skip} \rrbracket_u &\stackrel{\text{def}}{=} \lambda t. \text{STEP}(t) \\ \llbracket bexp \rrbracket_u &\stackrel{\text{def}}{=} \lambda t. \text{ASSIGN}_u(t) \\ \llbracket X := aexp \rrbracket_u &\stackrel{\text{def}}{=} \lambda t. \text{FILTER}_u(t) \end{aligned}$$

The `skip` statement is handled by the `STEP` operator. This operator increases the value of all defined partitions in the decision tree by one. Recall that a defined partition in a decision tree represents a set of states that satisfies some CTL property. The associated value is an upper bound on the number of steps until some condition is reached. By executing `skip` this number is incremented by one. For assignments and boolean conditions we use the corresponding *assign* and *filter* operators that were introduced in Sections 7.6 and 7.5. The definitions for the remaining path dependent operators all depend on these two functions.

Until

The abstract CTL semantics for universal and existential until properties are defined as the least fixed-point of the abstract transformers

$$\begin{aligned} \phi_{\forall(\Phi_1 U \Phi_2)}^\# &\in (\mathcal{L} \rightarrow \mathcal{T}_{NIL}) \rightarrow (\mathcal{L} \rightarrow \mathcal{T}_{NIL}) \\ \phi_{\exists(\Phi_1 U \Phi_2)}^\# &\in (\mathcal{L} \rightarrow \mathcal{T}_{NIL}) \rightarrow (\mathcal{L} \rightarrow \mathcal{T}_{NIL}) \end{aligned}$$

starting from the function τ_\perp that assigns the decision tree $\perp_{\mathcal{T}}$ to every node in the control-flow-graph. $\forall l \in \mathcal{L}: \tau_\perp(l) = \perp_{\mathcal{T}}$.

Definition 8.4. Abstract semantics for until operator.

$$\begin{aligned}
\tau_{\forall}^{\sharp}(\Phi_1 U \Phi_2) &\stackrel{\text{def}}{=} \text{lfp}_{\perp_{\mathcal{T}}}^{\sqsubseteq_{\mathcal{T}}} \phi_{\forall}^{\sharp}(\Phi_1 U \Phi_2) \\
t_{\gamma}(m, l) &\stackrel{\text{def}}{=} \bigvee_{(l, stmt, l') \in \text{out}(l)} \llbracket stmt \rrbracket_o(m(l')) \\
\phi_{\forall}^{\sharp}(\Phi_1 U \Phi_2)(m)l &\stackrel{\text{def}}{=} \text{UNTIL}[\tau_{\Phi_1}^{\sharp}(l), \tau_{\Phi_2}^{\sharp}(l)](t_{\gamma}(m, l)) \\
\\
\tau_{\exists}^{\sharp}(\Phi_1 U \Phi_2) &\stackrel{\text{def}}{=} \text{lfp}_{\perp_{\mathcal{T}}}^{\sqsubseteq_{\mathcal{T}}} \phi_{\exists}^{\sharp}(\Phi_1 U \Phi_2) \\
t_{\sqcup}(m, l) &\stackrel{\text{def}}{=} \bigsqcup_{(l, stmt, l') \in \text{out}(l)} \llbracket stmt \rrbracket_u(m(l')) \\
\phi_{\exists}^{\sharp}(\Phi_1 U \Phi_2)(m)l &\stackrel{\text{def}}{=} \text{UNTIL}[\tau_{\Phi_1}^{\sharp}(l), \tau_{\Phi_2}^{\sharp}(l)](t_{\sqcup}(m, l))
\end{aligned}$$

We will first discuss the universal version and then explain what changes for the existential case. The value $m \in \mathcal{L} \rightarrow \mathcal{T}_{NIL}$ is the current iteration value of the fixed-point iteration, $\phi_{\forall}^{\sharp}(\Phi_1 U \Phi_2)(m)$ describes the value of the next iteration. Recall that $\text{out}(l)$ denotes all outgoing edges of node l leading to its immediate successor nodes. Every edge is labeled with a statement. The abstract transformer $\phi_{\forall}^{\sharp}(\Phi_1 U \Phi_2)$ computes decision trees point-wise for each node l in the control-flow-graph, based on the decision trees of its successor nodes.

First, the decision tree of each successor node l' is passed to the $\llbracket stmt \rrbracket_o$ function. This approximates the effect of transitioning from l to l' . The resulting decision tree approximates the value of the ranking function before executing the statement.

If a node has multiple successor nodes then the resulting decision trees are combined using the *approximation join* \vee . The *approximation join* discards all partitions (i.e., makes them undefined) of decision trees that are not defined for all successor nodes. By doing so, we approximate the semantic of the universal path quantifier \forall . Note that if a node has no successors then $t_{\gamma}(m, l) = \perp_{\mathcal{T}}$ since the least upper bound (join) of the empty set is $\perp_{\mathcal{T}}$.

We use the over-approximating version of the $\llbracket \cdot \rrbracket_o$ function. This might temporarily lead to unsound decision trees due to over-approximation. Decision trees produced by

$\llbracket \cdot \rrbracket_o$ can contain defined partitions for states that are unfeasible among that path in the control-flow-graph. For the universal case however, this is not a problem since the *approximation join* only keeps those partitions which are feasible among all paths. Partitions that are unfeasible among some paths are discarded.

Finally the result of joining the decision trees of the immediate predecessors are applied to the $\text{UNTIL}[\tau_{\Phi_1}^\#, \tau_{\Phi_2}^\#]$ operator. The purpose of this operator is to implement the semantics of the *until* CTL operator. All partitions that satisfy Φ_1 are set to zero and all partitions that neither satisfy Φ_1 nor Φ_2 are discarded (see Algorithm 3). That way we end up with a decision tree that is only defined for the partitions that satisfy $\forall(\Phi_1 U \Phi_2)$.

The abstract transformer for the existential case follows the same structure as in the universal case. However instead of using the *approximation join* it uses the *computational join* $\sqcup_{\mathcal{T}}$ to approximate the semantics of the \exists path quantifier. The *computational join* preserves all partitions that are defined for at least one decision tree. Note however, that all decision trees passed to the *computation join* must be sound since we can no longer rely on the join operator to discard unfeasible partitions. Therefore we apply the under-approximating $\llbracket \text{stmt} \rrbracket_u$ function when processing statements to guarantee soundness.

Lemma 8.4. The abstract CTL semantics $\tau_{\forall(\Phi_1 U \Phi_2)}^\#$ and $\tau_{\exists(\Phi_1 U \Phi_2)}^\#$ are a sound approximation of the CTL semantics $\tau_{\forall(\Phi_1 U \Phi_2)}$ and $\tau_{\exists(\Phi_1 U \Phi_2)}$ with regards to the *approximation order* $\preceq_{\mathcal{T}}$ (see Definition 7.1).

Convergence of the least fixed-point iteration is guaranteed after a finite amount of iterations by using the widening operator $\nabla_{\mathcal{T}}$ (see Section 7.4). The decision tree computed for every node $l \in \mathcal{L}$ is guaranteed to converge by applying the following widening scheme (Φ_U is a placeholder for $\forall(\Phi_1 U \Phi_2)$ and $\exists(\Phi_1 U \Phi_2)$):

$$y_0 \stackrel{\text{def}}{=} \perp_{\mathcal{T}}$$

$$y_{n+1} \stackrel{\text{def}}{=} \begin{cases} y_n & \text{if } \phi_{\Phi_U}^\#(y_n) \sqsubseteq_{\mathcal{T}} y_n \wedge \phi_{\Phi_U}^\#(y_n) \preceq_{\mathcal{T}} y_n \\ y_n \nabla_{\mathcal{T}} \phi_{\Phi_U}^\#(y_n) & \text{otherwise} \end{cases}$$

Global

The abstract CTL semantics for universal and existential *global* properties are defined as the greatest fixed-point of the abstract transformers

$$\phi_{\forall \square \Phi}^\# \in (\mathcal{L} \rightarrow \mathcal{T}_{NIL}) \rightarrow (\mathcal{L} \rightarrow \mathcal{T}_{NIL})$$

$$\phi_{\exists \square \Phi}^\# \in (\mathcal{L} \rightarrow \mathcal{T}_{NIL}) \rightarrow (\mathcal{L} \rightarrow \mathcal{T}_{NIL})$$

starting from the abstract CTL semantics $\tau_{\Phi}^\#$ of the inner CTL property Φ .

Algorithm 2 Tree Until Filter

```
function FILTER_UNTIL( $t, t_{\text{valid}}$ )  
  if  $isNil(t) \vee isNil(t_{\text{valid}})$  then  
     $\triangleright$  ignore NIL nodes  
    return  $t$   
  else if  $isLeaf(t) \wedge isLeaf(t_{\text{valid}}) \wedge isDefined(t)$  then  
     $\triangleright$   $t$  is defined in  $t_{\text{valid}}$   
    return  $t$   
  else if  $isLeaf(t) \wedge isLeaf(t_{\text{valid}}) \wedge \neg isDefined(t)$  then  
     $\triangleright$   $t$  is not defined in  $t_{\text{valid}}$ , make undefined  
    return LEAF :  $\perp$   
  else  
     $l \leftarrow$  FILTER_UNTIL( $t.l, t_{\text{valid}}.l$ )  
     $r \leftarrow$  FILTER_UNTIL( $t.r, t_{\text{valid}}.r$ )  
    return NODE{ $t.c$ } :  $l; r$   
  end if  
end function
```

Definition 8.5. Abstract semantics for global operator.

$$\begin{aligned}\tau_{\forall \square \Phi}^{\#} &\stackrel{\text{def}}{=} \text{gfp}_{\tau_{\Phi}^{\#}}^{\sqsubseteq \tau} \phi_{\forall \square \Phi}^{\#} \\ t_{\gamma}(m, l) &\stackrel{\text{def}}{=} \bigvee_{(l, stmt, l') \in \text{out}(l)} \llbracket stmt \rrbracket_o(m(l')) \\ \phi_{\forall \square \Phi}^{\#}(m)l &\stackrel{\text{def}}{=} \text{MASK}[\llbracket t_{\gamma}(l, m) \rrbracket](m(l)) \\ \tau_{\exists \square \Phi}^{\#} &\stackrel{\text{def}}{=} \text{gfp}_{\tau_{\Phi}^{\#}}^{\sqsubseteq \tau} \phi_{\exists \square \Phi}^{\#} \\ t_{\sqcup}(m, l) &\stackrel{\text{def}}{=} \bigsqcup_{(l, stmt, l') \in \text{out}(l)} \llbracket stmt \rrbracket_u(m(l')) \\ \phi_{\exists \square \Phi}^{\#}(m)l &\stackrel{\text{def}}{=} \text{MASK}[\llbracket t_{\sqcup}(l, m) \rrbracket](m(l))\end{aligned}$$

The value $m \in \mathcal{L} \rightarrow \mathcal{T}_{NIL}$ is the current iteration value of the fixed-point iteration. The

Algorithm 3 Tree Until

```
function RESET_UNTIL( $t, t_{\text{reset}}$ )
  if  $isNil(t) \vee isNil(t_{\text{reset}})$  then
     $\triangleright$  ignore NIL nodes
    return  $t$ 
  else if  $isLeaf(t) \wedge isLeaf(t_{\text{reset}}) \wedge isDefined(t)$  then
     $\triangleright t$  is defined in  $t_{\text{valid}}$ , reset leaf
    return LEAF : 0
  else if  $isLeaf(t) \wedge isLeaf(t_{\text{valid}}) \wedge \neg isDefined(t)$  then
     $\triangleright t$  is undefined in  $t_{\text{valid}}$ , keep as is
    return  $t$ 
  else
     $l \leftarrow \text{RESET\_UNTIL}(t.l, t_{\text{reset}}.l)$ 
     $r \leftarrow \text{RESET\_UNTIL}(t.r, t_{\text{reset}}.r)$ 
    return NODE{ $t.c$ } :  $l; r$ 
  end if
end function

function UNTIL $[[t_{\Phi_1}, t_{\Phi_2}]](t)$   $\triangleright t, t_{\Phi_1}, t_{\Phi_2} \in \mathcal{T}_{NIL}$ 
   $(t_1, t_2) \leftarrow \text{TREE\_UNIFICATION}(t, t_{\Phi_1} \sqcup t_{\Phi_2})$ 
   $t_{\text{filtered}} \leftarrow \text{FILTER\_UNTIL}(t_1, t_2)$ 
   $(t_1, t_2) \leftarrow \text{TREE\_UNIFICATION}(t_{\text{filtered}}, t_{\Phi_2})$ 
  return  $\text{RESET\_UNTIL}(t_1, t_2)$ 
end function
```

definition for $\phi_{\forall\Box\Phi}^\sharp(m)$ ($\phi_{\exists\Box\Phi}^\sharp(m)$, respectively) describes the value of the next iteration. We use the same approach as the `until` operator to join decision trees of successor nodes (see $t_\gamma(m, l)$ and $t_\sqcup(m, l)$).

In the final step however, the decision tree $m(l)$ of each node l is *masked* with the result of $t_\gamma(l, m)$ ($t_\sqcup(l, m)$ respectively). Masking is implemented by the `MASK` operator. This operator sets all partitions of $m(l)$ to $\perp_{\mathcal{T}}$, that are not defined in $t_\gamma(l, m)$ ($t_\sqcup(l, m)$, respectively). That way all states that do not satisfy Φ indefinitely among all (some, respectively) paths are iteratively removed from the decision tree until a fixed-point is reached. The `MASK` operator is defined in Algorithm 4.

Lemma 8.5. The abstract CTL semantics $\tau_{\forall\Box\Phi}^\sharp$ and $\tau_{\exists\Box\Phi}^\sharp$ are a sound approximation of the CTL semantics $\tau_{\forall\Box\Phi}$ and $\tau_{\exists\Box\Phi}$ with regards to the *approximation order* $\preceq_{\mathcal{T}}$ (see Definition 7.1).

Convergence of the greatest fixed-point iteration is guaranteed after a finite amount of iterations by using the dual widening operator $\bar{\nabla}_{\mathcal{T}}$ (see Section 7.4). The decision tree computed for every node $l \in \mathcal{L}$ is guaranteed to converge by applying the following widening scheme (Φ_G is a placeholder for $\forall\Box\Phi$ and $\exists\Box\Phi$):

Algorithm 4 Tree Mask

```
function MASK $\llbracket t_{\text{MASK}} \rrbracket(t)$ 
  function MASK_AUX( $t, t_{\text{mask}}$ )
    if  $isNil(t) \vee isNil(t_{\text{reset}})$  then
       $\triangleright$  ignore NIL nodes
      return  $t$ 
    else if  $isLeaf(t) \wedge isDefined(t) \wedge isLeaf(t_{\text{mask}})$  then
      if  $isDefined(t) \wedge \neg isDefined(t_{\text{mask}})$  then
         $\triangleright$   $t$  is defined and  $t_{\text{mask}}$  is undefined, discard leaf
        return LEAF :  $\perp$ 
      else
        return  $t$ 
      end if
    else
       $l \leftarrow \text{MASK}(t.l, t_{\text{mask}}.l)$ 
       $r \leftarrow \text{MASK}(t.r, t_{\text{mask}}.r)$ 
      return NODE{ $t.c$ } :  $l; r$ 
    end if
  end function
  ( $t_1, t_2$ )  $\leftarrow$  TREE_UNIFICATION( $t, t_{\text{MASK}}$ )
  return MASK_AUX( $t_1, t_2$ )
end function
```

$$y_0 \stackrel{\text{def}}{=} \tau_{\Phi_U}^\#(l)$$
$$y_{n+1} \stackrel{\text{def}}{=} \begin{cases} y_n & \text{if } y_n \sqsubseteq_{\mathcal{T}} \phi_{\Phi_G}^\#(y_n) \wedge y_n \preceq_{\mathcal{T}} \phi_{\Phi_G}^\#(y_n) \\ y_n \nabla_{\mathcal{T}} \phi_{\Phi_G}^\#(y_n) & \text{otherwise} \end{cases}$$

Next

The abstract CTL semantics for the next operator are given in Definition 8.6:

Definition 8.6. Abstract semantics for next operator.

$$t_{\gamma}(l) \stackrel{\text{def}}{=} \bigvee_{(l, \text{stmt}, l') \in \text{out}(l)} \llbracket \text{stmt} \rrbracket_o(\tau_{\Phi}^{\#}(l))$$

$$\tau_{\forall \circ \Phi}^{\#} \stackrel{\text{def}}{=} \lambda l. \text{ZERO}(t_{\gamma}(l))$$

$$t_{\sqcup}(l) \stackrel{\text{def}}{=} \bigsqcup_{(l, \text{stmt}, l') \in \text{out}(l)} \llbracket \text{stmt} \rrbracket_u(\tau_{\Phi}^{\#}(l))$$

$$\tau_{\exists \circ \Phi}^{\#} \stackrel{\text{def}}{=} \lambda l. \text{ZERO}(t_{\sqcup}(l))$$

As opposed to the until and global operator, the decision trees for each label only depends on the immediate successor nodes. Therefore no fixed-point iteration is necessary. Each node is computed in one step based on the immediate successor nodes. Outgoing edges are joined as describe for the until and global operators. The resulting value is then passed to the ZERO operator which sets all defined partitions to zero (see Algorithm 5).

Lemma 8.6. The abstract CTL semantics $\tau_{\forall \circ \Phi}^{\#}$ and $\tau_{\exists \circ \Phi}^{\#}$ are a sound approximation of the CTL semantics $\tau_{\forall \circ \Phi}$ and $\tau_{\exists \circ \Phi}$ with regards to the *approximation order* $\preceq_{\mathcal{T}}$ (see Definition 7.1).

Algorithm 5 Tree Zero

```

function ZERO( $t$ )
  if  $\text{isLeaf}(t) \wedge \text{isDefined}(t)$  then
    return LEAF : 0
  else if  $\text{isNode}(t)$  then
     $l \leftarrow \text{ZERO}(t.l)$ 
     $r \leftarrow \text{ZERO}(t.r)$ 
    return NODE $\{t.c\} : l; r$ 
  else
    return  $t$ 
  end if
end function

```

Next we present how the abstract CTL semantics were implemented.

9 Implementation

This section describes how the CTL analysis was implemented as part of the static analyzer FuncTion¹. First, we briefly discuss the initial feature set of FuncTion, followed by a description of the various extensions that were added to the analyzer. Then we present a number of examples that showcase the capabilities of FuncTion.

9.1 Previous Work

FuncTion is an abstract interpretation-based static analyzer that was written by Caterina Urban in the context of her PhD thesis [5]. When starting this thesis, FuncTion had support for analyzing termination, guarantee properties and recurrence properties (see Section 4). To this end, it featured an implementation of the decision tree abstract domain (see Section 7) with the corresponding *join*, *meet*, *widen* (but not *dual-widen*), *filter*, *backward assign* and *reset* operators. Note that the *reset* operator was over-approximating. The *backward assign* operator over-approximates the assignment of each partition, the resulting decision tree however is an under-approximation in case of termination, recurrence and guarantee analysis due to the use of the *approximation join*. Note that something similar happens with *filter* when merging control flows. The Apron² library was used as basis for the numerical domains. FuncTion can analyze programs using the *interval*, *octagon* and *polyhedra* abstract domain. In addition to natural-value ranking functions, the decision tree abstract domain can also handle ordinal-valued ranking functions [11]. This makes it possible to infer ranking functions for programs with unbounded non-determinism. Finally, FuncTion was equipped with a front-end for parsing programs written in a C-like language. The termination analysis could only handle tail-recursive function calls. The analysis of guarantee and recurrence properties could not handle any function call.

9.2 CTL Analysis

The CTL analysis was implemented based on the existing components of FuncTion. We reused the implementation of the decision tree abstract domain and added implementations for the *complement*, *until* and *mask* operators presented in Section 8. We also wrote an implementation of the *dual-widen* operator which was defined in [5] but not implemented in FuncTion.

The analysis of existential properties is performed using the *computational join* (see Section 8), therefore the original *backward assign* operator cannot be reused due to the fact that it over-approximates the assignment of individual partitions. The same holds for the original *filter* operator. Therefore, we required under-approximating implementations of both *backward assign* and *filter*. Such an implementation for the polyhedra numerical abstract domain was proposed by Antoine Miné [12] and implemented in the

¹<https://github.com/caterinaurban/function>

²<http://apron.cri.enscm.fr/library>

Banal Static Analyzer³. We were able to reuse his implementation in `FuncTion` which allowed us to provide under-approximating version of *filter* and *backward assign* with minimal effort. In addition to that, the under-approximating version of *filter* also allowed us to fix the unsoundness problem of the *reset* operator which we discussed in Section 8. This resolves the problem for the new CTL analysis and also the existing guarantee and recurrence analysis. Note that every analysis that requires an under-approximating operator can only be performed with the polyhedra numerical abstract domain at the moment⁴. This includes the existing recurrence and guarantee analysis due to the use of the *reset* operator.

The initial implementation of the CTL analysis in `FuncTion` reuses the existing front-end and performs the analysis through a backward analysis on the abstract-syntax-tree. This had the drawback that the analyzer did not support control-flow statements such as `goto`, `continue` or `break`. Furthermore there was no support for function calls. These problems were later resolved by moving the analysis to a control-flow-graph based implementation. This will be discussed in the next subsection.

9.3 Improved Front-end

As mentioned the analysis based on abstract-syntax-trees (ASTs) had several drawbacks. Performing the analysis on the AST made it difficult to support `goto` statements and function calls. Therefore we changed the analysis to a control-flow-graph based analysis. To that end we reused an existing front-end⁵ which features an abstract-syntax-tree to control-flow-graph conversion.

The new implementation of the backward analysis is based on a standard worklist algorithm [13]. Nodes added to the worklist are processed in FIFO order. To improve performance, nodes are never added twice.

Widening is only performed at loop heads. To that end, we implemented a simple loop-detection algorithm. First we check if the control-flow-graph is reducible. Then we detect loop heads by finding edges in the control-flow-graph where the head dominates its tail. The dominator tree is computed with a simple $O(mn)$ algorithm. The performance of this algorithm proved to be sufficient for all analyzed test cases. Better algorithms exists [14] should the need for analyzing very large control-flow-graphs arise. More information about loop-detection can be found in the corresponding literature [15].

Non-recursive function calls are handled by inlining. Recursive calls are resolved by adding edges to the control-flow-graph from every call site to the entry node of the function and from the exit node of the function back to the call sites. This approach leads

³<https://www-apr.lip6.fr/~mine/banal/>

⁴Banal supports the *interval* and *octagon* domain in theory but the implementation is currently broken due to a bug.

⁵<https://www-apr.lip6.fr/~mine/enseignement/13/2015-2016/project.html>

```

while (x > y) do
  x := x - 1
od

```

Figure 11: Program `test_until`

to infeasible paths in the control-flow-graph. However the result of the analysis is still sound. We did not pursue more involved intra-procedural analysis approaches since this was not the focus of our work.

Existential CTL properties can be analyzed using the abstract semantics presented in Section 8. In addition to that, it is also possible to convert existential properties to equivalent universal properties (see Section 3). This can yield more precise results in some cases (see Example 5 in the next subsection).

9.4 Examples

We now demonstrate the capabilities and limitations of `FuncTion` with a number of example programs.

Example 1

Consider the program `test_until` in Figure 11. One can show that the program satisfies the property $x \geq y \Rightarrow \forall(x \geq y U x = y)$. `FuncTion` is able to prove this with the following counting function:

$$f(x, y) = \begin{cases} 2(x - y) & \text{if } x \geq y \\ \text{undefined} & \text{otherwise} \end{cases}$$

This states that the program satisfies the CTL properties if $x \geq y$ holds initially. Furthermore, the goal of the *until* property ($x = y$) is reached in at most $2(x - y)$ program execution steps.

Example 2

Consider program `test_global` in Figure 12. One can show that the property $\forall \diamond \forall \square (y > 0)$ is satisfied if $x < 10$ holds initially. In that case, the condition of the inner loop is reached eventually and we end up in a loop that infinitely increments y . `FuncTion` can prove this by deriving the following counting function:

$$f(x) = \begin{cases} -3x + 32 & \text{if } x < 10 \\ \text{undefined} & \text{otherwise} \end{cases}$$

```

y := 0
while (true) do
  x := x + 1
  while (x = 10) do y := y + 1
od

```

Figure 12: Program `test_global`

```

r := 0
while (?) do
  x := x + 1
  if (x = 200) then r := 1
od

```

Figure 13: Program `test_existential_1`

Given the precondition $x < 10$, the program will reach a state that satisfies $\forall\Box(y > 10)$ in at most $-3x + 32$ steps.

Example 3

The program `test_existential_1` in Figure 13 increments x an arbitrary amount of time. If x reaches 200 then r is set to 1. This program satisfies the existential property $x < 200 \Rightarrow \exists\Diamond(r = 1)$. FuncTion proves this with the following counting function:

$$f(x) = \begin{cases} -3x + 601 & \text{if } x < 200 \\ \text{undefined} & \text{otherwise} \end{cases}$$

If $x < 200$ then the program will set r to 1 in at most $-3x + 601$ steps.

Example 4

The program `test_existential_2` in Figure 14 is a variation of `test_existential_1`. Instead of performing a non-deterministic check at the loop-head, the program assigns a random number to variable i and then enters the loop if the choice for i was positive. Unfortunately, FuncTion is not able to show that `test_existential_2` satisfies $x < 200 \Rightarrow \exists\Diamond(r = 1)$. This is because the decision tree abstract domain is not able to capture the effect of non-deterministic variable assignments in case of existential CTL properties.

For program `test_existential_2`, FuncTion infers the following counting function for the loop-head:

$$f(x) = \begin{cases} -3x + 600 & \text{if } x < 200 \wedge i \geq 0 \\ \text{undefined} & \text{otherwise} \end{cases}$$

```

r := 0
i := ?
while (i ≥ 0) do
  x := x + 1
  if (x = 200) then r := 1
od

```

Figure 14: Program test_existential_2

```

r := 0
while (x > 0) do
  x := x - 1
  if (?) then r := 1
od

```

Figure 15: Program test_existential_3

After applying the $\text{ASSIGN}_u[[i := ?]]$ operator, all function values that depend on i are lost. That is because the decision tree abstract domain is not expressive enough to capture the fact that there exists a choice for i that satisfies the property. The reason for that is inherent in the current implementation of the decision tree abstract domain. Assignments are applied to each partition individually. It is not possible to capture the effect of the non-deterministic assignment over multiple partitions. Currently there is no straight forward solution to this problem.

Example 5

The duality between existential and universal CTL properties (see Section 3) can improve the precision of the analysis in some cases. Consider program test_existential_3 in Figure 15. This program satisfies the existential property $x > 0 \Rightarrow \exists \diamond (r = 1)$. Unfortunately, FuncTion is not able to prove this directly. While FuncTion is able to show $x = 2 \Rightarrow \exists \diamond (r = 1)$, it fails for the stronger property due to imprecision. However FuncTion successfully proves the equivalent universal property $x > 0 \Rightarrow \neg \forall \square (r \neq 1)$. FuncTion can be configured to automatically convert all existential properties to their universal counterpart when possible.

Example 6

The program acqrel in Figure 16 was taken from a set of example programs⁶ written by Eric Koskinen for his PhD thesis [16]. We assume that initially $a = r = 0$. The condition of the first loop is a non-deterministic choice, it is executed an arbitrary amount of times. Whenever the loop is entered a is set to 1. Eventually r is set to 1 after an arbitrary

⁶<https://github.com/ultimate-pa/ultimate/tree/dev/trunk/examples/LTL/koskinen>

```

a := 0
r := 0
while (?) do
  a := 1
  a := 0
  n := ?
  while (n > 0) do n := n - 1
  r := 1
  r := 0
od
while (true) do skip

```

Figure 16: Program `acqrel`

amount of steps, depending on the choice for n . One can see that this program satisfies the property $\forall \square((a = 1) \Rightarrow \forall \diamond(r = 1))$. `FuncTion` is able to prove this property using ordinal-valued ranking functions. Note that the property is a safety property, therefore the value of the inferred counting function at the program entry point is 0 and holds no meaningful information.

Next we present an evaluation of the `FuncTion` CTL analysis capabilities.

10 Evaluation

This section presents the result of our evaluation of the analysis capabilities of `FuncTion` for CTL properties. We evaluated the capabilities of `FuncTion` based on series of test cases. To compare how `FuncTion` fares against other static analyzers, we ran the same test cases (where possible) on these other analyzers. To following two static analyzers where chosen for a comparative analysis:

T2 Temporal Prover

The T2 Temporal Prover [17] is a static analyzer with support (among other things) for CTL properties. The analysis of CTL properties is based on the work of Cook et. al. [18]. T2 uses a control-flow-graph based program representation. Programs need to be convert to the T2 format by hand in most cases which makes analyzing larger C programs rather tedious. There exists an LLVM [19] backend `llvm2KITTeL`⁷ which converts LLVM bitcode to T2 format. However, `llvm2KITTeL` removes dead assignments and renames variables due to the static-single-assignment format of LLVM. This makes `llvm2KITTeL` unsuitable for this evaluation, since all CTL properties in our test cases depend on variable writes and the original variable names defined in the C source code.

⁷<https://github.com/s-falke/llvm2kittel>

Ultimate LTL Automizer

Ultimate LTL Automizer [20] is part of the Ultimate program analysis framework⁸. It can analyze LTL properties of C programs. LTL and CTL share a common fragment but are not comparable in general [21]. Therefore not all of our test cases can be run by Ultimate.

Table 1 lists all test cases and the corresponding CTL properties that were analyzed for the evaluation⁹. The first set of test cases is a series of toy examples that demonstrate the basic capabilities of FuncTion (1.x). The second set of test cases (2.x) were taken from Eric Kosinens PhD thesis [16] and are publicly available on GitHub¹⁰. Note that we only took a subset of these test cases since not all of them were compatible with FuncTion due to arrays, pointers, structs etc. The third set of test cases (3.x) are sample programs from Ultimate LTL Automizer¹¹. The final set of test cases (4.x) are a selection of non-terminating programs from SV-Comp¹². The corresponding CTL properties were chosen by us.

Table 2 compares FuncTion against T2 and Ultimate.

The first set of test cases were converted to the T2 format by hand. All test cases with non-existential properties were also executed on Ultimate since they are compatible with LTL. Note that Ultimate cannot handle existential CTL properties. T2 struggles with some of the examples of the first set. It even crashed for test cases 1.4 and 1.9¹³. Test cases 1.13 and 1.14 can only be verified by FuncTion. Note that, as opposed to FuncTion, T2 does not support existential `until` properties (see 1.14). Ultimate can handle all non-existential CTL properties of the first test set. FuncTion fails test case 1.8 because it cannot handle non-deterministic assignments with existential properties.

In the second test set, FuncTion fails to verify test cases 2.3 and 2.6 due to imprecision of the decision tree abstract domain. Test case 2.4 is a slightly modified version of 2.3 that FuncTion is able to verify. Ultimate was able to verify all examples of the second test set. T2 verified the first four, the last two were too large to convert to the T2 format.

All test cases of the third test set safe 3.9, 3.15 and 3.16 could be successfully verified by FuncTion. The three failed test cases could not be verified due to imprecision of the decision tree abstract domain. As in the second test set, test case 3.10 is a slightly modified version of 3.9 that FuncTion is able to verify. All test cases of the third set are rather large and contain function calls. Therefore we could not afford to translate

⁸<https://ultimate.informatik.uni-freiburg.de>

⁹Test system: Lenovo Thinkpad T460s, Intel i7-6600U, 20GB Ram, Arch Linux with Linux 4.11, OCaml 4.04.1, APRON 20160125

¹⁰<https://github.com/ultimate-pa/ultimate/tree/dev/trunk/examples/LTL/koskinen>

¹¹<https://github.com/ultimate-pa/ultimate/tree/dev/trunk/examples/LTL/coolant>

¹²<https://sv-comp.sosy-lab.org/2017>

¹³This was reported to the T2 team

No	Program	CTL Property	Result	Time
1.1	and_test.c	$\forall \square \forall \diamond (n = 1) \wedge \forall \diamond (n = 0)$	✓	0.05s
1.2	and_test.c	$\forall \square \forall \diamond (n = 1)$	✓	0.05s
1.3	and_test.c	$\forall \diamond (n = 0)$	✓	0.05s
1.4	global_test.c	$\forall \square \forall \diamond (x \leq -10)$	✓	0.15s
1.5	multi_choice.c	$\forall \diamond (x = 4 \vee x = -4)$	✓	0.04s
1.6	multi_choice.c	$\exists \diamond (x = -4)$	✓	0.04s
1.7	or_test.c	$\forall \diamond \forall \square (x < -100) \vee \forall \diamond (x = 20)$	✓	0.05s
1.8	may_term...	$\exists \diamond (\text{exit} : \text{true})$	✗	-
1.9	until_test.c	$\forall (x \geq y \text{ U } x = y)$	✓	0.03s
1.10	exists_test1.c	$\exists \diamond (r = 1)$	✓	0.01s
1.11	exists_test3.c	$\exists \diamond (r = 1)$	✓	0.02s
1.12	exists_test4.c	$\exists \diamond (r = 1)$	✓	0.01s
1.13	fin_ex.c	$\exists \square \exists \diamond (n = 1)$	✓	0.04s
1.14	until_ex.c	$\exists (x \geq y \text{ U } x = y)$	✓	0.03s
2.1	acqrel.c	$\forall \square (a = 1 \Rightarrow \forall \diamond (r = 1))$	✓*	0.75s
2.2	acqrel.c*	$\forall \square (a = 1 \Rightarrow \forall \diamond (r = 1))$	✓	8s
2.3	fig8-2007.c	$\text{set} = 1 \Rightarrow \forall \diamond (\text{unset} = 1)$	✗	-
2.4	fig8-2007.c*	$\text{set} = 1 \Rightarrow \forall \diamond (\text{unset} = 1)$	✓	1s
2.5	win4.c	$\forall \diamond \forall \square (\text{WItemsNum} \geq 1)$	✓	0.15s
2.6	toylin.c	$(c \leq 5 \wedge c > 0) \vee \forall \diamond (\text{resp} > 5)$	✗	-
3.1	cb1_safe.c	$\forall \square (\text{cb} = 1 \Rightarrow \forall \square (\text{cb} = 1))$	✓	0.75s
3.2	cb1_unsafe.c	$\neg \forall \square (\text{cb} = 1 \Rightarrow \forall \square (\text{cb} = 1))$	✓	3s
3.3	cb2_safe.c	$\forall \square \forall \diamond (\text{ot} < t)$	✓	3.2s
3.4	cb2_unsafe.c	$\neg \forall \square \forall \diamond (\text{ot} < t)$	✓	0.8s
3.5	cb3_safe.c	$\forall \square (i = 3 \Rightarrow \forall \square \forall \diamond (t > \text{ot}))$	✓	1.7s
3.6	cb3_unsafe.c	$\neg \forall \square (i = 3 \Rightarrow \forall \square \forall \diamond (t > \text{ot}))$	✓	1.25s
3.7	cb4_safe.c	$\forall \square (i \neq 3 \vee \text{tem} \leq \text{lim} \vee \forall \diamond \forall \square (\text{cb} = 1))$	✓	5.5s
3.8	cb4_unsafe.c	$\neg \forall \square (i \neq 3 \vee \text{tem} \leq \text{lim} \vee \forall \diamond \forall \square (\text{cb} = 1))$	✓	85s
3.9	cb5_safe.c	$\forall (i = 0 \text{ U } (\forall (i = 1 \text{ U } \forall \square (i = 3)) \vee \forall \square (i = 1)))$	✗	-
3.10	cb5_safe.c*	$\forall (i = 0 \text{ U } (\forall (i = 1 \text{ U } \forall \square (i = 3)) \vee \forall \square (i = 1)))$	✓	14s
3.11	cb5_unsafe.c	$\neg \forall (i = 0 \text{ U } (\forall (i = 1 \text{ U } \forall \square (i = 3)) \vee \forall \square (i = 1)))$	✓	14s
3.12	cb6_safe.c	$\forall \square (\text{lim} \leq -273 \vee \text{lim} \geq 10 \vee \text{ti} \geq 0 \vee \forall \diamond (\text{wl} = 1))$	✓	6.5s
3.13	cb6_unsafe.c	$\neg \forall \square (\text{lim} \leq -273 \vee \text{lim} \geq 10 \vee \text{ti} \geq 0 \vee \forall \diamond (\text{wl} = 1))$	✓	6.5s
3.14	nestedR...	$\forall \square (i \geq n)$	✓	0.01s
3.15	timer...	$\neg \forall \square (\text{timer} = 0 \Rightarrow \forall \diamond (\text{output} = 1))$	✗	-
3.16	togglec...	$\forall \square (\forall \diamond (t = 1) \wedge \forall \diamond (t = 0))$	✗	-
3.17	togglet...	$\forall \square (\forall \diamond (t = 1) \wedge \forall \diamond (t = 0))$	✓	0.02s
3.18	Pot... Min...	$\forall \square (x > 0 \Rightarrow \forall \diamond (y = 0))$	✓*	0.15s
3.19	cav2015.c	$\forall \square (x > 0 \Rightarrow \forall \diamond (y = 0))$	✓*	0.14s
3.20	simple-1.c	$\forall \diamond (x > 10000)$	✓	0.07s
3.21	simple-2.c	$\forall \diamond (x > 100)$	✓	0.05s
3.22	Bug_No...	$\neg \forall \diamond (\text{ap} > 2)$	✓	0.12s
4.1	Bangalore...	$\exists \diamond (x < 0)$	✗	-
4.2	Ex02...	$i < 5 \Rightarrow \forall \diamond (\text{exit} : \text{true})$	✓	0.04s
4.3	Ex07...	$\forall \diamond \forall \square (i = 0)$	✓	0.1s
4.4	java.Seq...	$\forall \diamond (\forall \diamond (j \geq 21) \wedge i = 100)$	✓	0.3s
4.5	Madrid...	$\forall \diamond (x = 7 \wedge \forall \diamond \forall \square (x = 2))$	✓	0.02s
6.6	NO_02...	$\forall \diamond \forall \square (j = 0)$	✓	0.05s

Table 1: Table of test cases. Lists the number, program name, CTL property to be analyzed, result of the analysis and runtime. Test cases with a ✓* use ordinal-valued counting functions. Program names marked with * are a slightly adjusted version of the original.

them to T2 due to time limitations. Ultimate LTL Automizer successfully verified all test cases of the third set. This was to be expected since the test cases were taken from the Ultimate LTL Automizer repository.

In the fourth test set we again compare FuncTion against T2 and Ultimate where possible. This yields a similar result as the first test set. FuncTion fails 4.1 due to a non-deterministic assignment. Ultimate is able to verify all LTL compatible test cases. The CTL analysis capabilities of T2 seem to be rather error prone at this point in time which makes it hard to draw a fair comparison. Furthermore, the T2 specific input format makes it difficult to run larger examples with T2.

FuncTion comes close to the capabilities of Ultimate LTL Automizer when analyzing compatible LTL properties. Ultimate was able to verify all LTL compatible test cases. FuncTion, as opposed to Ultimate, struggles in some cases when dealing with non-determinism. It is easy to pinpoint these problems and introduce adjustments that would lead to a successful analysis. This gives important insights for future improvements to FuncTion. The benefit of FuncTion over Ultimate is that it can analyze existential properties and that it provides more information when analyzing liveness properties due to the derived counting functions. Furthermore, only FuncTion is able to analyze existential until properties.

In conclusion, we believe that our CTL analysis framework is able to keep up with other available static analyzers. The next section concludes the results of this thesis and discusses possible future work.

No	Program	FuncTion	T2	Ultimate LTL Automizer
1.1	and_test.c	✓	✗	✓
1.2	and_test.c	✓	✗	✓
1.3	and_test.c	✓	✓	✓
1.4	global_test.c	✓	error	✓
1.5	multi_choice.c	✓	✓	✓
1.6	multi_choice.c	✓	✓	-
1.7	or_test.c	✓	error	✓
1.8	may_term...	✗	✓	-
1.9	until_test.c	✓	✗	✓
1.10	exists_test1.c	✓	✓	-
1.11	exists_test3.c	✓	✓	-
1.12	exists_test4.c	✓	✓	-
1.13	fin_ex.c	✓	✗	-
1.14	until_ex.c	✓	-	-
2.1	acqrel.c	✓*	✓	-
2.2	acqrel.c*	✓	✓	-
2.3	fig8-2007.c	✗	✓	-
2.4	fig8-2007.c*	✓	✓	-
2.5	win4.c	✓	-	-
2.6	toylin.c	✗	-	-
3.1	cb1_safe.c	✓	-	✓
3.2	cb1_unsafe.c	✓	-	✓
3.3	cb2_safe.c	✓	-	✓
3.4	cb2_unsafe.c	✓	-	✓
3.5	cb3_safe.c	✓	-	✓
3.6	cb3_unsafe.c	✓	-	✓
3.7	cb4_safe.c	✓	-	✓
3.8	cb4_unsafe.c	✓	-	✓
3.9	cb5_safe.c	✗	-	✓
3.10	cb5_safe.c*	✓	-	✓
3.11	cb5_unsafe.c	✓	-	✓
3.12	cb6_safe.c	✓	-	✓
3.13	cb6_unsafe.c	✓	-	✓
3.14	nestedR...	✓	-	✓
3.15	timer...	✗	-	✓
3.16	togglec...	✗	-	✓
3.17	togglet...	✓	-	✓
3.18	Pot...Min...	✓*	-	✓
3.19	cav2015.c	✓*	-	✓
3.20	simple-1.c	✓	-	✓
3.21	simple-2.c	✓	-	✓
3.22	Bug_No...	✓	-	✓
4.1	Bangalore...	✗	✓	-
4.2	Ex02...	✓	out of memory	✓
4.3	Ex07...	✓	✗	✓
4.4	java_Seq...	✓	error	✓
4.5	Madrid...	✓	error	✓
4.6	NO_02...	✓	✓	✓

Table 2: Table with results of test cases for FuncTion, T2 and Ultimate LTL Automizer. Test cases with a ✓* use ordinal-valued counting functions. Program names marked with * are a slightly adjusted version of the original.

11 Conclusion

This section concludes the work conducted for this thesis and gives an outline for future work.

Our main contribution is an abstract interpretation based static analysis framework for CTL. The framework inductively defines the concrete semantics of arbitrary combinations of CTL operators for state transition system. We present a sound, computable approximation of the concrete semantics that is based on the decision-tree abstract domain. Our approach for analyzing existential CTL properties is based on under-approximation operators for the domain of piecewise-defined linear functions.

We provide an experimental implementation of our framework as part of the `FuncTion` static analyzer. A series of test-cases demonstrates that our framework is capable of analyzing a large variety of sample programs and properties. Furthermore, we compared our implementation against existing static analysis frameworks, the T2 Temporal Prover and Ultimate LTL Automizer. The results of our evaluation show that our implementation is on par with the two existing analyzer on the majority of the test-cases.

A weakness of our implementation is the handling of non-deterministic assignments in case of existential properties. Information about variables is lost during the analysis due to imprecision of the decision-tree abstract domain.

Another contribution is a new frontend for `FuncTion` that supports arbitrary control-flow jumps and branches. In addition to that, the CTL analysis, as opposed to the previous versions of `FuncTion`, now supports function inlining and function calls. This greatly extends the variety of programs that can be analyzed by `FuncTion`.

11.1 Future Work

We now propose possible future work based on the results of this thesis.

Improved Precision for Non-determinism

`FuncTion` suffers from imprecision when handling non-deterministic assignments in case of existential properties. The decision-tree abstract domain could be improved to offer preciser results when dealing with non-determinism. This would greatly improve the existing CTL analysis capabilities of `FuncTion`.

Proving Soundness

Currently, there is no formal proof for the soundness of the presented abstract semantics. A formal proof would improve the confidence in the analysis results of our framework.

Supporting LTL Properties

Future work could investigate how the current analysis framework could be extended to support linear temporal logic (LTL). This would require a more substantial adaptation of the existing work because CTL and LTL are not comparable in general. LTL is not as easily composable as CTL. The semantic of CTL is defined in terms of states which makes it possible to take a set of states that satisfy one CTL formula as input to another atomic CTL operator. In LTL however, the semantic is defined in terms of traces starting from an initial state. If one can show that a property holds for a set of traces, it is not obvious how that information can be reused to show that a more complex LTL formula holds.

References

- [1] Patrick Cousot and Radhia Cousot, “Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints,” in *Proceedings of the 4th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, New York, NY, USA, 1977, POPL ’77, pp. 238–252, ACM.
- [2] P. Cousot and R. Cousot, “An abstract interpretation framework for termination,” in *Conference Record of the 39th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, Philadelphia, PA, Jan. 25-27 2012, pp. 245–258, ACM Press, New York.
- [3] Caterina Urban and Antoine Miné, “Proving Guarantee and Recurrence Temporal Properties by Abstract Interpretation,” in *VMCAI*, 2015, pp. 190–208.
- [4] Christel Baier, Joost-Pieter Katoen, and Kim Guldstrand Larsen, *Principles of model checking*, MIT press, 2008.
- [5] Caterina Urban, *Static Analysis by Abstract Interpretation of Functional Temporal Properties of Programs.*, Ph.D. thesis, École Normale Supérieure, Paris, France, 2015.
- [6] Alan Turing, “Checking a large routing,” *Report of a Conference on High Speed Automatic Calculating Machines*, pp. 67–69, 1949.
- [7] Robert W. Floyd, “Assigning meanings to programs,” *Proceedings of Symposium on Applied Mathematics*, pp. 19:19–32, 1967.
- [8] P. Cousot and R. Cousot, “Static determination of dynamic properties of programs,” in *Proceedings of the Second International Symposium on Programming*. 1976, pp. 106–130, Dunod, Paris, France.
- [9] Antoine Miné, “The octagon abstract domain,” *Higher Order Symbol. Comput.*, vol. 19, no. 1, pp. 31–100, Mar. 2006.
- [10] Patrick Cousot and Nicolas Halbwachs, “Automatic discovery of linear restraints among variables of a program,” in *Proceedings of the 5th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, New York, NY, USA, 1978, POPL ’78, pp. 84–96, ACM.
- [11] Caterina Urban and Antoine Miné, “An Abstract Domain to Infer Ordinal-Valued Ranking Functions,” in *ESOP*, 2014, pp. 412–431.
- [12] Antoine Miné, “Inferring sufficient conditions with backward polyhedral under-approximations,” *Electr. Notes Theor. Comput. Sci.*, vol. 287, pp. 89–100, 2012.
- [13] Flemming Nielson, Hanne Riis Nielson, and Chris Hankin, *Principles of program analysis*, Springer, 1999.

- [14] Thomas Lengauer and Robert Endre Tarjan, “A fast algorithm for finding dominators in a flowgraph,” *ACM Trans. Program. Lang. Syst.*, vol. 1, no. 1, pp. 121–141, Jan. 1979.
- [15] Alfred V. Aho, Ravi Sethi, and Jeffrey D. Ullman, *Compilers: Principles, Techniques, and Tools*, Addison-Wesley series in computer science / World student series edition. Addison-Wesley, 1986.
- [16] Eric Koskinen, *Temporal verification of programs*, Ph.D. thesis, University of Cambridge, 2012.
- [17] Marc Brockschmidt, Byron Cook, Samin Ishtiaq, Heidy Khlaaf, and Nir Piterman, “T2: temporal property verification,” in *Tools and Algorithms for the Construction and Analysis of Systems - 22nd International Conference, TACAS 2016, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2016, Eindhoven, The Netherlands, April 2-8, 2016, Proceedings*, 2016, pp. 387–393.
- [18] Byron Cook, Eric Koskinen, and Moshe Y. Vardi, “Temporal property verification as a program analysis task - extended version,” *Formal Methods in System Design*, vol. 41, no. 1, pp. 66–82, 2012.
- [19] Chris Lattner and Vikram Adve, “Llvm: A compilation framework for lifelong program analysis & transformation,” in *Proceedings of the International Symposium on Code Generation and Optimization: Feedback-directed and Runtime Optimization*, Washington, DC, USA, 2004, CGO '04, pp. 75–, IEEE Computer Society.
- [20] Daniel Dietsch, Matthias Heizmann, Vincent Langenfeld, and Andreas Podelski, “Fairness modulo theory: A new approach to LTL software model checking,” in *Computer Aided Verification - 27th International Conference, CAV 2015, San Francisco, CA, USA, July 18-24, 2015, Proceedings, Part I*, 2015, pp. 49–66.
- [21] Christel Baier and Joost-Pieter Katoen, *Principles of Model Checking (Representation and Mind Series)*, The MIT Press, 2008.



Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Declaration of originality

The signed declaration of originality is a component of every semester paper, Bachelor's thesis, Master's thesis and any other degree paper undertaken during the course of studies, including the respective electronic versions.

Lecturers may also require a declaration of originality for other written papers compiled for their courses.

I hereby confirm that I am the sole author of the written work here enclosed and that I have compiled it in my own words. Parts excepted are corrections of form and content by the supervisor.

Title of work (in block letters):

Proving Temporal Properties by Abstract Interpretation

Authored by (in block letters):

For papers written by groups the names of all authors are required.

Name(s):

Samuel Ueltschi

First name(s):

Samuel

With my signature I confirm that

- I have committed none of the forms of plagiarism described in the 'Citation etiquette' information sheet.
- I have documented all methods, data and processes truthfully.
- I have not manipulated any data.
- I have mentioned all persons who were significant facilitators of the work.

I am aware that the work may be screened electronically for plagiarism.

Place, date

Zürich, 6.9.2017

Signature(s)

Samuel Ueltschi

For papers written by groups the names of all authors are required. Their signatures collectively guarantee the entire content of the written paper.