FASTER TEMPORAL REASONING FOR INFINITE-STATE PROGRAMS

Heidy Khlaaf^I Byron Cook^{1,2} Nir Piterman³

¹University College London ²Microsoft Research ³University of Leicester

BRIEF OVERVIEW

- A new symbolic model checking procedure for CTL verification of infinite-state programs.
- Counterexample-guided precondition synthesis strategy to compute location(PC) specific preconditions.
- Existing tools not scalable. We propose:
 - Compositional strategy via exploiting the natural decomposition of the program's state space.
 - Performance improvement and scalability!

CTL - A REFRESHER

- Temporal logic reasoning about sets of states.
- Implemented via formula structure.
- Reasoning about non-deterministic (branching) programs.
- Termination AFAX false, Non-Termination EGEX true, AG Safety
- These operators can be nested to generate complex liveness and safety properties.
- Used to uncover bugs in device drivers, operating systems, servers, etc.

CTL - A REFRESHER

- A ϕ All: ϕ has to hold on all paths starting from all initial states.
- E ϕ Exists: there exists at least one path starting from all initial states where ϕ holds.
- G ϕ Globally: ϕ has to hold on the all states along a path.
- F ϕ Finally: ϕ eventually has to hold.
- Other operators include X ϕ Next, ϕ_1 U ϕ_2 Until, ϕ_1 W ϕ_2 Weak until.

INTUITION



• Verify EF y > z

INTUITION



 The set of states satisfying EF y > z before a program command is very often the same as the set of states respecting EF y > z after the command.

INTUITION



- We can infer whether a command is likely to affect the truth of EF y > z.
 - From one counterexample, we use precondition synthesis to infer the pre-image of all locations within a counterexample.

EXAMPLE (PT. I)



• Prove the CTL property AGEF y = 1.

EXAMPLE (PT. I)



- Recurse over the structure of the CTL formula.
- Find \wp such that AG \wp holds, and \wp |= EF y = 1.
- $\wp\langle\phi\rangle$ takes the form \bigwedge_{i} (pc = $i_i \Rightarrow \wp\langle i_i, \phi\rangle$).

EXAMPLE (PI.I) $\begin{array}{l} \rho_1: \mathbf{X}' = \ast \\ \mathbf{y}' = 0 \end{array}$ $\rho_4: \mathbf{X} > 0 \land$ $y \neq 1$ $\mathbf{y} \neq 1$ $\rho_2: \mathbf{X} \leq 0 \land$ $\mathbf{y} \neq 1$ ℓ_2 ℓ_1 $\rho_3: \mathbf{X} \leq 0 \land$ $\mathbf{y} \neq \mathbf{1}$ $\rho_6 : | \mathbf{y} = 1$ $\rho_7 : | \mathbf{y} = 1$ ERR

- Verify the universal dual of the existential property and seek a set of counterexamples to serve as witnesses.
- Transform the program for the property $\phi = EF y = I \text{ using its dual AG y} \neq I$.

EXAMPLE (PT. I) $\rho_1: \mathbf{X}' = \ast$ $\rho_4: \mathbf{X} > 0 \land$ $\mathbf{y}'=0$ $\mathbf{y} \neq \mathbf{1}$ $\mathbf{y} \neq 1$ ho_5 : $\rho_2: \mathbf{x} \leq 0 \land$ ℓ_2 ℓ_1 $\mathbf{x}' = \mathbf{x} + 1$ $\rho_3: \mathbf{X} \leq 0 \land$ $y \neq 1$ $\rho_6 : | \mathbf{y} = 1$ $\rho_7 : | \mathbf{y} = 1$ ERR

• Initially $\wp(\phi) =$ false as only failures to proving AG y \neq 1 imply that there exists a witness such that EF y = 1.



- Use a safety prover to check reachability of ERR, and begin from I_{1} .
- For every reachable location $I \in L$ in CEX₁, we compute a pre-image using the suffix of CEX₁ from I onwards.



- $pre(CEX_1) = y = 0.$
- When a new counterexample is discovered, we refine $\sqrt[6]{}\langle I, \phi \rangle$ resulting in $\sqrt[6]{}\langle I, \phi \rangle = V_{n \in N} \text{pre}(\text{CEX}_n)$



• $pre(CEX_{1'}) = x > 0.$

• Computed a refinement for I_2 from a counterexample generated for I_1 . No need to verify I_2 independently!

$\mathsf{EXAMPLE}(\mathsf{PT.I})$

- Ensure EF y = 1 satisfies all initial states: Rule out CEX by adding \neg pre(CEX₁) to each transition from I to the error state.
- Re-run the safety checker.
- No more counterexamples are generated and all locations covered:

$$\Im \langle \mathsf{EF} \mathsf{y} = \mathsf{I} \rangle = (\mathsf{pc} = \mathsf{I}_1 \Rightarrow \mathsf{y} = 0) \land (\mathsf{pc} = \mathsf{I}_2 \Rightarrow \mathsf{x} > 0).$$



- Modify $\phi = AGEF y = I$ by using $\Im(EF y = I)$:
 - $\varphi = AG ((pc = I_1 \Rightarrow y=0) \land (pc=I_2 \Rightarrow x>0)).$



• $\mathfrak{S}(I, \boldsymbol{\phi})$ resulting in $\mathfrak{S}(I, \boldsymbol{\phi}) = \bigwedge_{n \in \mathbb{N}} \neg \text{pre}(\text{CEX}_n)$

• Universal: the initial precondition $\wp\langle \varphi \rangle$ = true. No counterexamples are generated thus $\wp\langle AGEF \rangle = 1 \rangle$ = true!

RECAP

• Partition CTL formula preconditions by program location:

•
$$\mathfrak{O}\langle \varphi \rangle$$
 takes the form $\bigwedge_{i} (\text{pc} = i \Rightarrow \mathfrak{O}\langle i, \varphi \rangle).$

• Universal location preconditions:

•
$$\mathcal{O}(I, \phi) = \bigwedge_{n \in \mathbb{N}} \neg pre(CEX_n)$$

• Existential location preconditions:

•
$$\delta \langle I, \phi \rangle = V_{n \in N} \text{pre}(\text{CEX}_n)$$

EXPERIMENTS

- Built as an extension to the open source project T2
 - Source Code: http://research.microsoft.com/en-us/projects/t2/.
- Input: C files converted to t2 file format + CTL specification.
- Compared our tool to :

[1] T. A. Beyene, C. Popeea, and A. Rybalchenko, "Solving existentially quantified horn clauses," in *CAV* 13. Springer, 2013

[2] B. Cook and E. Koskinen, "Reasoning about nondeterminism in programs," in *PLDI'I* 3. ACM, 2013.

EXPERIMENTS

LoC	Property	T2	[1]	[2]
1050	$AG(b = I \longrightarrow AF(u = 0))$	67.3	T/O	T/O
1050	$EG(b = I \longrightarrow EF(u = 0))$	36.2	T/O	T/O
370	$AG(a = I \longrightarrow EF(r = I))$	6.8	35.45	90.0
370	EF(a = 1 && AG (r ≠1))	4.7	T/O	T/O
370	EG(io ≠ I) && EG(ret ≠ I)	13.5	T/O	7.6
370	AG(io ≠ I) AG(ret ≠ I)	8.0	0.1	T/O
90	AGEF w = 1	2.0	0.7	T/O
90	EFAG w ≠ I	2.0	0.1	T/O
90	EFEG w ≠ I	0.1	0.1	35.2

LIMITATIONS

- Divergence can occur due to infinitely many counterexamples.
 - Take pre-image α of a PC, quantify out all variables that are updated proceeding a program location.
 - Can lead to unsoundness due to over-approximation of the set of states for existential path quantifiers.
 - Check that the precondition is sound e.g. that $\wp_1 \Rightarrow EG \wp_2$, we can use SMT based strategies to double check the small lemma on initial locations.

SUMMARY

- A new symbolic model checking procedure for CTL verification of infinite-state programs.
- Use a counterexample-guided precondition synthesis strategy to compute location-specific preconditions.
- Reduces the amount of irrelevant reasoning traditionally performed as several preconditions for each location can be computed simultaneously.
- Performance improvement and scalability!

BACKGROUND

- $\mathbf{P} = (L, E, Vars),$
- Each edge τ = (I,ρ,I') in E, where I, I'∈ L and ρ is a condition, specifies possible transitions in the program.
- $\mathbf{T} = (S, R)$
 - $\mathbf{S} = \mathbf{L} \times (\text{Vars} \rightarrow \text{Vals})$
 - $\mathbf{R} \subseteq S \times S$
- A **cut-point** is a set C such that C ⊆ L and every cycle in the program's graph contains at least one cut-point.
- **Pre-images**: For a path $\pi = (l_0, \rho_0, l_0'), (l_1, \rho_1, l_1'), \dots, (l_n, \rho_n, l_n')$, we compute a pre-image for every possible suffix of π :

•
$$\operatorname{pre}_{n+1} = S$$
 and $\operatorname{pre}_{i} = \operatorname{pre}[(I_{i}, \rho_{i}, I'_{i}), \dots, (I_{n}, \rho_{n}, I'_{n})]$ as the set of states such that $\operatorname{pre}_{i} = \{s \mid \exists s' \in \operatorname{pre}_{i+1} \text{ s.t. } ((I_{i}, s), (I'_{i'}, s')) \mid = \rho_{i}\}.$

FINDING TEMPORAL PRECONDITIONS

- Recurse over the structure of the given CTL formula.
- For each CTL sub-formula ϕ we find a precondition $\wp\langle\phi\rangle$ that ensures its satisfaction.
- Each sub-formula ϕ is then replaced with $\wp\langle\phi\rangle$ within the original formula.
 - Note: It is only necessary to handle formulas of nesting depth 1.
- To utilize sequential locality of a counterexample's control-flow graph:

•
$$(\varphi \land \varphi)$$
 takes the form $\bigwedge_{i} (pc = i_i \Rightarrow \langle \varphi \land i_i, \phi \rangle).$

FINDING TEMPORAL PRECONDITIONS

• For a universal CTL sub-property:

• A precondition $\wp\langle I, \phi \rangle$ for a program location I is initialized to true.

• When a new counterexample is discovered, we refine $\log \langle I, \phi \rangle$ resulting in $\log \langle I, \phi \rangle = \Lambda_{n \in \mathbb{N}} \neg \text{pre}(\text{CEX}_n)$

• For an existential CTL sub-property:

- Verify the universal dual of the existential property and seek a set of counterexamples to serve as witnesses.
- A precondition $\wp\langle I, \phi \rangle$ for a program location I is initialized to false.
- When a new counterexample is discovered, we refine $\sqrt[6]{|, \phi\rangle}$ resulting in $\sqrt[6]{|, \phi\rangle} = V_{n \in \mathbb{N}} \text{pre}(\text{CEX}_n)$