

Probabilistic Hyperproperties with Nondeterminism

Erika **Ábrahám**, RWTH Aachen, Germany

Borzoo Bonakdarpour and **Oyendrilá Dobe**, Michigan State University

Ezio Bartocci, TU Wien, Austria

ATVA 2020

22 October 2020

- A **trace** $t = s_0, s_1, \dots$ is an infinite sequence of states $s_i \in S$.
- A **trace property** is a set of traces.

Example (LTL property)

I never start working before having a coffee.

$$(\mathcal{G} \neg work) \vee ((\neg work) \mathcal{U} coffee)$$

- A **trace** $t = s_0, s_1, \dots$ is an infinite sequence of states $s_i \in S$.
- A **trace property** is a set of traces.

Example (LTL property)

I never start working before having a coffee.

$$(\mathcal{G} \neg work) \vee ((\neg work) \mathcal{U} coffee)$$

Classical trace properties cannot express relations between traces.

- A **trace** $t = s_0, s_1, \dots$ is an infinite sequence of states $s_i \in S$.
- A **trace property** is a set of traces.

Example (LTL property)

I never start working before having a coffee.

$$(\mathcal{G} \neg work) \vee ((\neg work) \mathcal{U} coffee)$$

Classical trace properties cannot express relations between traces.

- A **hyperproperty** is a set of sets of traces.

- A **trace** $t = s_0, s_1, \dots$ is an infinite sequence of states $s_i \in S$.
- A **trace property** is a set of traces.

Example (LTL property)

I never start working before having a coffee.

$$(\mathcal{G} \neg work) \vee ((\neg work) \mathcal{U} coffee)$$

Classical trace properties cannot express relations between traces.

- A **hyperproperty** is a set of sets of traces.

Example (HyperLTL property)

I drink coffee every day at the same time.

$$\forall \pi. \forall \pi'. (\mathcal{G} (coffee_\pi \Leftrightarrow coffee_{\pi'}))$$

Consider a parallel program (h : high input / l : low output).

```
while  $h > 0$  do {  $h \leftarrow h - 1$  };  $l \leftarrow 2$  ||  $l \leftarrow 1$ 
```

Consider a parallel program (h : high input / l : low output).

```
while  $h > 0$  do {  $h \leftarrow h - 1$  };  $l \leftarrow 2$  ||  $l \leftarrow 1$ 
```

Assuming a **uniform probabilistic scheduler**:

Consider a parallel program (h : high input / l : low output).

```
while  $h > 0$  do {  $h \leftarrow h - 1$  };  $l \leftarrow 2$  ||  $l \leftarrow 1$ 
```

Assuming a **uniform probabilistic scheduler**:

- $h = 0 \rightarrow \mathbb{P}(l=1) = 1/4$ at termination.
- $h = 5 \rightarrow \mathbb{P}(l=1) = 1/4096$ at termination.

Consider a parallel program (h : high input / l : low output).

```
while  $h > 0$  do {  $h \leftarrow h - 1$  };  $l \leftarrow 2$  ||  $l \leftarrow 1$ 
```

Assuming a **uniform probabilistic scheduler**:

- $h = 0 \rightarrow \mathbb{P}(l=1) = 1/4$ at termination.
- $h = 5 \rightarrow \mathbb{P}(l=1) = 1/4096$ at termination.

Probabilistic noninterference stipulates that the **probability distribution** on the final values on publicly observable channels (**low outputs**) is independent of the initial values of secrets (**high inputs**).

Consider a parallel program (h : high input / l : low output).

```
while  $h > 0$  do {  $h \leftarrow h - 1$  };  $l \leftarrow 2$  ||  $l \leftarrow 1$ 
```

Assuming a **uniform probabilistic scheduler**:

- $h = 0 \rightarrow \mathbb{P}(l=1) = 1/4$ at termination.
- $h = 5 \rightarrow \mathbb{P}(l=1) = 1/4096$ at termination.

Probabilistic noninterference stipulates that the **probability distribution** on the final values on publicly observable channels (**low outputs**) is independent of the initial values of secrets (**high inputs**).

We need **probabilistic hyperproperties** to express probabilistic relations between independent executions of a system.

HyperPCTL: PCTL extended with quantification over initial states

Example (Probabilistic noninterference)

$$\forall \hat{s}. \forall \hat{s}'. \left(\text{init}_{\hat{s}} \wedge \text{init}_{\hat{s}'} \wedge h_{\hat{s}} \neq h_{\hat{s}'} \right) \Rightarrow$$

$$\left(\mathbb{P} \left(\mathcal{F}(\text{fin}_{\hat{s}} \wedge (l=1)_{\hat{s}}) \right) = \mathbb{P} \left(\mathcal{F}(\text{fin}_{\hat{s}'} \wedge (l=1)_{\hat{s}'} \right) \right)$$

HyperPCTL: PCTL extended with quantification over initial states

Example (Probabilistic noninterference)

$$\forall \hat{s}. \forall \hat{s}'. \left(\text{init}_{\hat{s}} \wedge \text{init}_{\hat{s}'} \wedge h_{\hat{s}} \neq h_{\hat{s}'} \right) \Rightarrow \\ \left(\mathbb{P} \left(\mathcal{F}(\text{fin}_{\hat{s}} \wedge (l=1)_{\hat{s}}) \right) = \mathbb{P} \left(\mathcal{F}(\text{fin}_{\hat{s}'} \wedge (l=1)_{\hat{s}'}) \right) \right)$$

But this argues about a fixed scheduler only...

HyperPCTL: PCTL extended with quantification over initial states

Example (Probabilistic noninterference)

$$\forall \hat{s}. \forall \hat{s}'. \left(\text{init}_{\hat{s}} \wedge \text{init}_{\hat{s}'} \wedge h_{\hat{s}} \neq h_{\hat{s}'} \right) \Rightarrow \\ \left(\mathbb{P} \left(\mathcal{F}(\text{fin}_{\hat{s}} \wedge (l=1)_{\hat{s}}) \right) = \mathbb{P} \left(\mathcal{F}(\text{fin}_{\hat{s}'} \wedge (l=1)_{\hat{s}'}) \right) \right)$$

But this argues about a fixed scheduler only...

Our contribution:

extend HyperPCTL to **non-deterministic probabilistic systems!**

HyperPCTL formulas are similar to PCTL formulas **BUT**

- they quantify ($Q_i \in \{\exists, \forall\}$) over schedulers and initial states:

$$Q_{\hat{\sigma}_1} \hat{\sigma}_1 \dots Q_{\hat{\sigma}_m} \hat{\sigma}_m \cdot Q_{\hat{s}_1} \hat{s}_1(\hat{\sigma}_{i_1}) \dots Q_{\hat{s}_n} \hat{s}_n(\hat{\sigma}_{i_n}) \cdot \psi$$


- they index atomic propositions:

$$\psi ::= a_s \mid \psi \wedge \psi \mid \neg \psi \mid p < c$$

- they support arithmetic computations with probability expressions:


$$p ::= \mathbb{P}(\mathcal{X}\psi) \mid \mathbb{P}(\psi \mathcal{U} \psi) \mid \mathbb{P}(\psi \mathcal{U}^{[k_1, k_2]} \psi) \mid \\ c \mid p + p \mid p - p \mid p \cdot p$$

Example




$$\psi = \exists \hat{\sigma}. \forall \hat{s}. \forall \hat{s}'. (\text{init}_{\hat{s}} \wedge \text{init}_{\hat{s}'}) \Rightarrow (\mathbb{P}(\mathcal{F}(l=1))_{\hat{s}} = \mathbb{P}(\mathcal{F}(l=1))_{\hat{s}'})$$

Example




Example




DTMCs and their self-composition


Example




two
quantifiers
 \rightsquigarrow




||





Example



two
quantifiers
 \rightsquigarrow



\parallel



$$\mathcal{M} \models \varphi \quad \text{iff} \quad \underbrace{\mathcal{M}}_{MDP}, \underbrace{()}_{\text{schedulers of initial states}}, \underbrace{()}_{\text{initial states}} \models \varphi$$

$$\mathcal{M}, \vec{\sigma}, \vec{r} \models \forall \hat{\sigma}. \varphi \quad \text{iff} \quad \forall \sigma \in \Sigma^{\mathcal{M}}. \mathcal{M}, \vec{\sigma}, \vec{r} \models \varphi[\hat{\sigma} \rightsquigarrow \sigma]$$

...

$$\mathcal{M}, \vec{\sigma}, \vec{r} \models \forall \hat{s}(\sigma). \varphi \quad \text{iff} \quad \forall s_{n+1} \in S. \mathcal{M}, \vec{\sigma} \circ \sigma, \vec{r} \circ (\text{init}(s_{n+1}), s_{n+1}) \models \varphi[\hat{s} \rightsquigarrow s_{n+1}]$$

...

$$\llbracket \mathbb{P}(\varphi_{path}) \rrbracket_{\mathcal{M}, \vec{\sigma}, \vec{r}} = Pr^{\mathcal{M}^{\vec{\sigma}}}(\{\pi \in Paths^{\vec{r}}(\mathcal{M}^{\vec{\sigma}}) \mid \mathcal{M}, \vec{\sigma}, \pi \models \varphi_{path}\})$$


...


$$\mathcal{M}, \vec{\sigma}, \pi \models \varphi_1 \mathcal{U} \varphi_2 \quad \text{iff} \quad \exists j \geq 0. (\mathcal{M}, \vec{\sigma}, \vec{r}_j \models \varphi_2 \wedge \forall i \in [0, j). \mathcal{M}, \vec{\sigma}, \vec{r}_i \models \varphi_1)$$

- For DTMCs, the two semantics are equivalent.
- For MDPs,
 - PCTL evaluates a probability constraint $\mathbb{P}(\varphi) < c$ to true if it holds **under all schedulers**,
 - HyperPCTL supports free quantifier type and positioning.

Theorem

The HyperPCTL model checking problem for MDPs is undecidable.





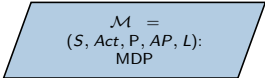
We propose an SMT-based technique for solving the model checking problem, such that

$$\mathcal{M} = (S, Act, P, AP, L) \text{ satisfies } Q\hat{\sigma}.Q_1\hat{\sigma}_1(\hat{\sigma}) \dots Q_n\hat{\sigma}_n(\hat{\sigma}).\varphi^{nq}$$

iff

SMT encoding is satisfied

We explain the **simplified case** of having a **single scheduler quantifier** for understanding of the basic ideas.




$\mathcal{M} =$
 $(S, Act, P, AP, L):$
MDP

SMT Encoding Algorithm for Existential scheduler quantifier


$\mathcal{M} =$
 $(S, Act, P, AP, L):$
MDP

$\exists \hat{\sigma}. Q_1 \hat{s}_1(\hat{\sigma}). \dots Q_n \hat{s}_n(\hat{\sigma}). \varphi^{nq}:$
HyperPCTL formula.


SMT Encoding Algorithm for Existential scheduler quantifier



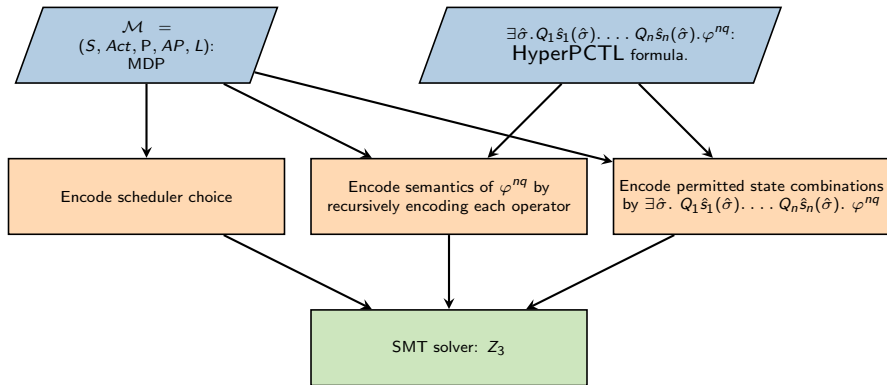
SMT Encoding Algorithm for Existential scheduler quantifier



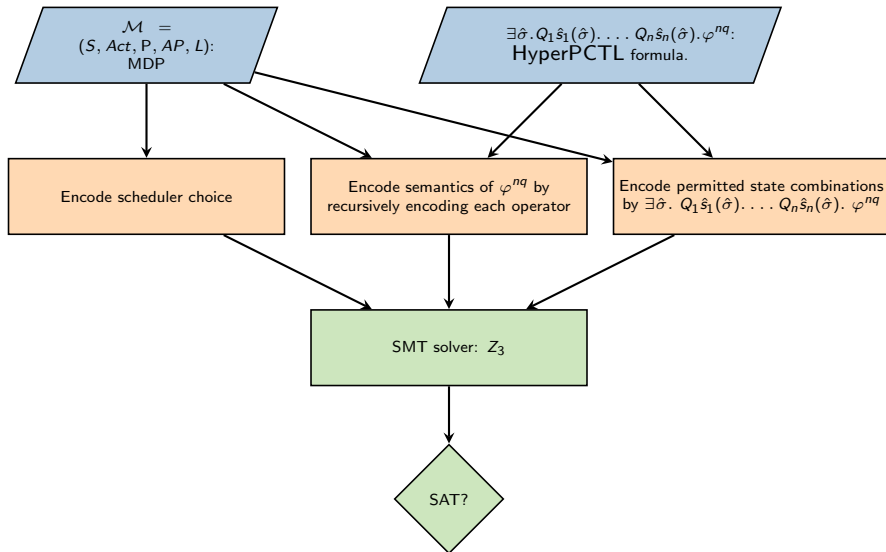
SMT Encoding Algorithm for Existential scheduler quantifier



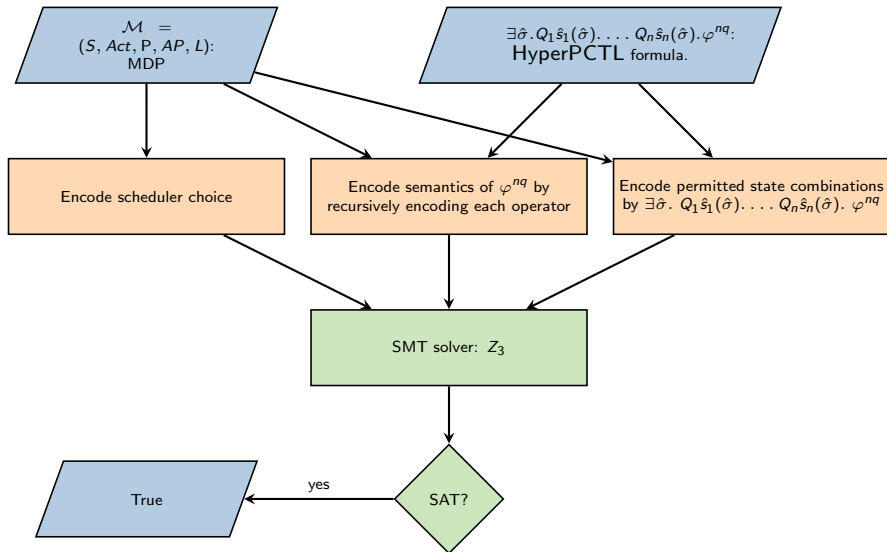
SMT Encoding Algorithm for Existential scheduler quantifier



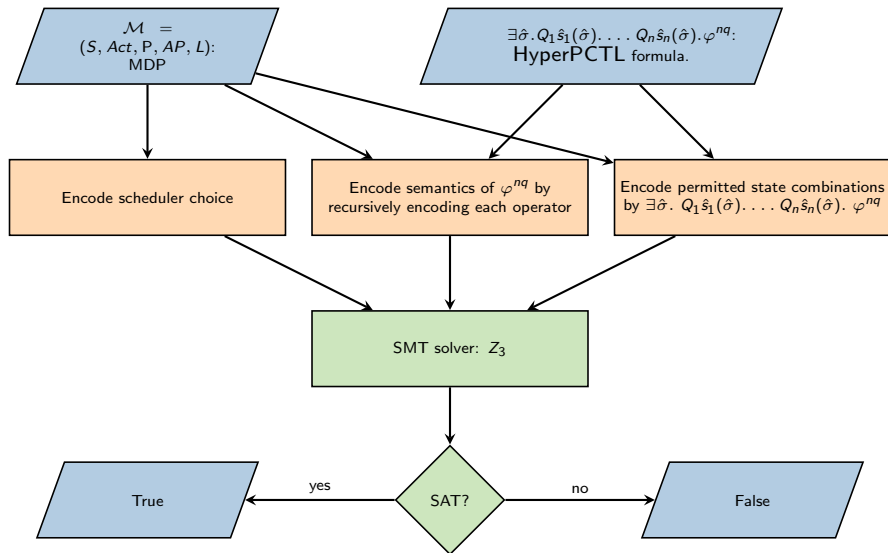
SMT Encoding Algorithm for Existential scheduler quantifier



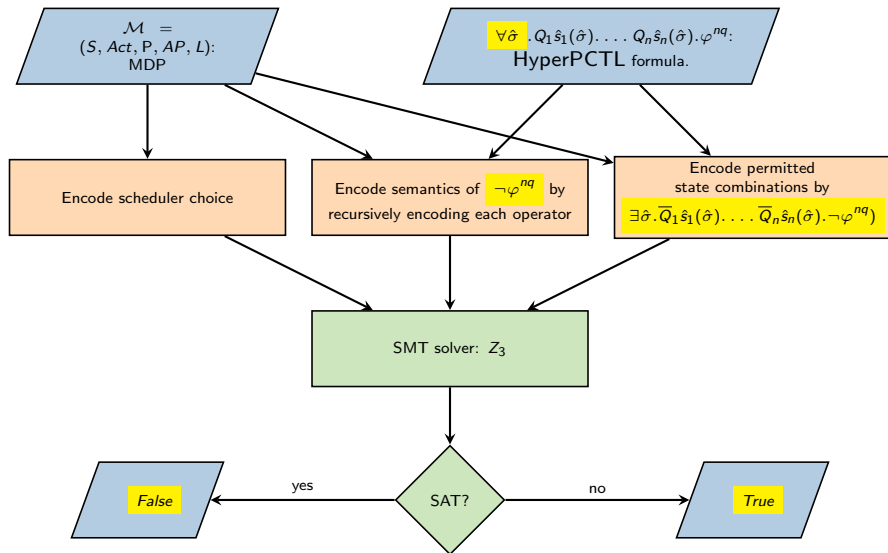
SMT Encoding Algorithm for Existential scheduler quantifier



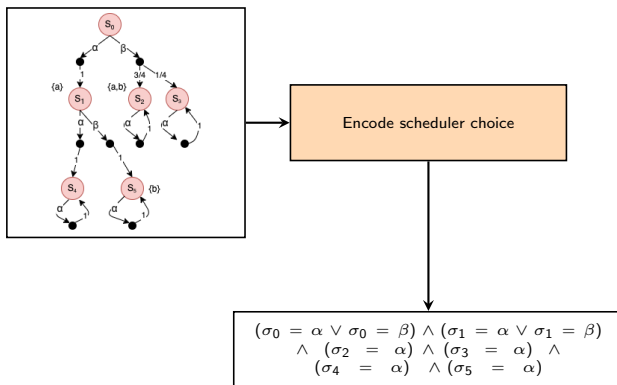
SMT Encoding Algorithm for Existential scheduler quantifier




SMT Encoding Algorithm for Universal scheduler quantifier



Elaborating on each sub-method used: scheduler choice



Elaborating on each sub-method used: State Quantifier encoding




$$\exists \hat{\sigma}. \forall \hat{s}(\hat{\sigma}). \exists \hat{s}'(\hat{\sigma}). (init_{\hat{s}} \wedge init_{\hat{s}'}) \Rightarrow (\mathbb{P}(\mathcal{F}a_{\hat{s}}) = \mathbb{P}(\mathcal{F}a_{\hat{s}'}))$$

Encode permitted state combinations

$$\begin{aligned} & (holds_{s_0, s_0, \varphi}^{nq} \vee holds_{s_0, s_1, \varphi}^{nq} \vee holds_{s_0, s_2, \varphi}^{nq} \vee holds_{s_0, s_3, \varphi}^{nq} \vee holds_{s_0, s_4, \varphi}^{nq} \vee holds_{s_0, s_5, \varphi}^{nq}) \wedge \\ & (holds_{s_1, s_0, \varphi}^{nq} \vee holds_{s_1, s_1, \varphi}^{nq} \vee holds_{s_1, s_2, \varphi}^{nq} \vee holds_{s_1, s_3, \varphi}^{nq} \vee holds_{s_1, s_4, \varphi}^{nq} \vee holds_{s_1, s_5, \varphi}^{nq}) \wedge \\ & (holds_{s_2, s_0, \varphi}^{nq} \vee \dots \vee holds_{s_2, s_5, \varphi}^{nq}) \wedge (holds_{s_3, s_0, \varphi}^{nq} \vee \dots \vee holds_{s_3, s_5, \varphi}^{nq}) \wedge \\ & (holds_{s_4, s_0, \varphi}^{nq} \vee \dots \vee holds_{s_4, s_5, \varphi}^{nq}) \wedge (holds_{s_5, s_0, \varphi}^{nq} \vee \dots \vee holds_{s_5, s_5, \varphi}^{nq}) \end{aligned}$$

Elaborating on each sub-method used: Semantics encoding




$$(init_{S_5} \wedge init_{S_1'}) \Rightarrow (\mathbb{P}(\mathcal{F}a_{S_5}) = \mathbb{P}(\mathcal{F}a_{S_1'}))$$

Encode semantics

$$Encoding(init_{S_5}) = (holds_{s_0,s_0,init_{S_5}} \wedge \dots \wedge holds_{s_0,s_5,init_{S_5}}) \wedge (\neg holds_{s_1,s_0,init_{S_5}} \wedge \neg holds_{s_1,s_1,init_{S_5}} \wedge \dots \wedge \neg holds_{s_5,s_5,init_{S_5}})$$

Elaborating on each sub-method used: Semantics encoding




$$(init_{S_5} \wedge init_{S_1'}) \Rightarrow (\mathbb{P}(\mathcal{F}a_{S_5}) = \mathbb{P}(\mathcal{F}a_{S_1'}))$$

Encode semantics

$$\begin{aligned}
 \text{Encoding}(init_{S_5}) &= (holds_{S_0, S_0, init_{S_5}} \wedge \dots \wedge holds_{S_0, S_5, init_{S_5}}) \wedge (\neg holds_{S_1, S_0, init_{S_5}} \wedge \neg holds_{S_1, S_1, init_{S_5}} \wedge \dots \wedge \neg holds_{S_5, S_5, init_{S_5}}) \\
 \text{Encoding}(init_{S_5} \wedge init_{S_1'}) &\text{ for state}(s_0, s_0) = (holds_{S_0, S_0, (init_{S_5} \wedge init_{S_1'})} \wedge holds_{S_0, S_0, init_{S_5}} \wedge holds_{S_0, S_0, init_{S_1'}}) \vee \\
 &(\neg holds_{S_0, S_0, (init_{S_5} \wedge init_{S_1'})} \wedge (\neg holds_{S_0, S_0, init_{S_5}} \vee \neg holds_{S_0, S_0, init_{S_1'}}))
 \end{aligned}$$

Elaborating on each sub-method used: Semantics encoding



$$(init_{\xi} \wedge init_{\xi'}) \Rightarrow (\mathbb{P}(\mathcal{F}a_{\xi}) = \mathbb{P}(\mathcal{F}a_{\xi'}))$$

Encode semantics

$$\begin{aligned}
 \text{Encoding}(init_{\xi}) &= (holds_{s_0, s_0, init_{\xi}} \wedge \dots \wedge holds_{s_0, s_5, init_{\xi}}) \wedge (\neg holds_{s_1, s_0, init_{\xi}} \wedge \neg holds_{s_1, s_1, init_{\xi}} \wedge \dots \wedge \neg holds_{s_5, s_5, init_{\xi}}) \\
 \text{Encoding}(init_{\xi} \wedge init_{\xi'}) \text{ for state}(s_0, s_0) &= (holds_{s_0, s_0, (init_{\xi} \wedge init_{\xi'})} \wedge holds_{s_0, s_0, init_{\xi}} \wedge holds_{s_0, s_0, init_{\xi'}}) \vee \\
 &\quad (\neg holds_{s_0, s_0, (init_{\xi} \wedge init_{\xi'})} \wedge (\neg holds_{s_0, s_0, init_{\xi}} \vee \neg holds_{s_0, s_0, init_{\xi'}})) \\
 \text{Encoding}(\mathbb{P}(\mathcal{F}a_{\xi})) \text{ for state}(s_0, s_0) \text{ action}(\alpha_{\xi} \cdot \alpha_{\xi'}) &= (holds_{s_0, s_0, a_{\xi}} \Rightarrow prob_{s_0, s_0, \mathbb{P}(\mathcal{F}a_{\xi})} = 1) \wedge (prob_{s_0, s_0, \mathbb{P}(\mathcal{F}a_{\xi})} \geq 0) \\
 &\quad \wedge (\neg holds_{s_0, s_0, a_{\xi}} \wedge (\sigma_0 = \alpha \wedge \sigma_0 = \alpha)) \Rightarrow (prob_{s_0, s_0, \mathbb{P}(\mathcal{F}a_{\xi})} = (1 * 1 * prob_{s_1, s_1, \mathbb{P}(\mathcal{F}a_{\xi})}) \\
 &\quad \wedge (prob_{s_0, s_0, \mathbb{P}(\mathcal{F}a_{\xi})} > 0 \Rightarrow (holds_{s_1, s_1, a_{\xi}} \vee d_{s_0, s_0, a_{\xi}} > d_{s_1, s_1, a_{\xi}})))
 \end{aligned}$$

- It allows an attacker to infer the value of a secret by observing execution time of a function.
- In this example, a is an integer representing the plaintext and b is the integer encryption key.
- This algorithm should satisfy the following property:

```
1 void mexp() {
2   c = 0; d = 1; i = k;
3   while (i >= 0) {
4     i = i - 1; c = c * 2;
5     d = (d * d) % n;
6     if (b(i) = 1)
7       c = c + 1;
8     d = (d * a) % n;
9   }
10 }
11 /*****/
12 t = new Thread(mexp());
13 j = 0; m = 2 * k;
14 while (j < m & !t.stop) j++;
15 /*****/
```

Figure: Modular Exponentiation

$$\forall \hat{\sigma}_1. \forall \hat{\sigma}_2. \forall \hat{s}(\hat{\sigma}_1). \forall \hat{s}'(\hat{\sigma}_2). \left(\text{init}_{\hat{s}} \wedge \text{init}_{\hat{s}'} \right) \Rightarrow \bigwedge_{l=0}^m \left(\mathbb{P}(\mathcal{F}(j = l)_{\hat{s}}) = \mathbb{P}(\mathcal{F}(j = l)_{\hat{s}'}) \right)$$

- It allows an attacker to infer the value of a secret by observing execution time of a function.
- We want to ensure an attacker thread cannot be used to infer the number of correct bits we have in the user input.

```
1 int str_cmp(char * r){
2   char * s = 'Bg\0';
3   i = 0;
4   while (s[i] != '\0'){
5     i++;
6     if (s[i]!=r[i]) return 0;
7   }
8   return 1;
9 }
```

Figure: String comparison

Assume two threads:

$$th_1 : \text{ while } h > 0 \text{ do } \{h := h - 1\}; l := 2$$

$$th_2 : l := 1$$

- Attacker should not be able to choose a specific scheduler to control set of traces generated.
- Observational determinism should be followed across all schedulers, according to the formula,

$$\forall \hat{\sigma}. \forall \hat{s}(\hat{\sigma}). \forall \hat{s}'(\hat{\sigma}). (h_{\hat{s}} \oplus h_{\hat{s}'}) \Rightarrow \mathbb{P}\mathcal{G}\left(\bigwedge_{l \in L} ((\mathbb{P}\mathcal{X}l_{\hat{s}}) = (\mathbb{P}\mathcal{X}l_{\hat{s}'}))\right) = 1.$$

- Check if implementation conforms with given specification.
- Synthesize a protocol that simulates the 6-sided die behavior only by repeatedly tossing a fair coin.
- Given all the possible coin-implementations, our goal is to check if there exists a scheduler that gives us the DTMC from the given MDP using,




Figure: Knuth-Yao protocol for simulating fair dice using only fair coins

$$\exists \hat{\sigma}. \forall \hat{s}(\hat{\sigma}). \exists \hat{s}'(\hat{\sigma}). \left(\text{init}_{\hat{s}} \wedge \text{init}_{\hat{s}'} \right) \Rightarrow \bigwedge_{l=1}^6 \left(\mathbb{P}(\mathcal{F}(\text{die} = l)_{\hat{s}}) = \mathbb{P}(\mathcal{F}(\text{die} = l)_{\hat{s}'}) \right)$$

Case study	Running time (s)			#SMT variables	#subformulas	#states	#transitions	
	SMT encoding	SMT solving	Total					
TA	$m = 2$	5.43	0.31	5.74	8088	50654	24	46
	$m = 4$	114	20	134	50460	368062	60	136
	$m = 6$	1721	865	2586	175728	1381118	112	274
PW	$m = 2$	5.14	0.3	8.14	8088	43432	24	46
	$m = 4$	207	40	247	68670	397852	70	146
	$m = 6$	3980	1099	5079	274540	1641200	140	302
TS	$h = (0, 1)$	0.83	0.07	0.9	1379	7913	7	13
	$h = (0, 15)$	60	1607	1667	34335	251737	35	83
	$h = (4, 8)$	11.86	17.02	28.88	12369	87097	21	48
	$h = (8, 15)$	60	1606	1666	34335	251737	35	83
PC	$s=(0)$	277	1996	2273	21220	1859004	20	158
	$s=(0,1)$	822	5808	6630	21220	5349205	20	280
	$s=(0,1,2)$	1690	58095	59785	21220	11006581	20	404

Table: Experimental results. **TA:** Timing attack. **PW:** Password leakage. **TS:** Thread scheduling. **PC:** Probabilistic conformance.

Here, for **TA**, m refers to $2 \cdot$ number of bits in the encryption key, for **PW**, m refers to $2 \cdot$ length of user password, for **PC**, we have included all possible transitions from the mentioned states.

Extended the temporal logic **HyperPCTL** for DTMCs to the context of MDPs by allowing formulas to quantify over schedulers.

Extended the temporal logic **HyperPCTL** for DTMCs to the context of MDPs by allowing formulas to quantify over schedulers.

This additional expressive power leads to undecidability of the **HyperPCTL** model checking on MDPs.

Extended the temporal logic **HyperPCTL** for DTMCs to the context of MDPs by allowing formulas to quantify over schedulers.

This additional expressive power leads to undecidability of the **HyperPCTL** model checking on MDPs.

Presented a SMT-based model checking algorithm which is NP-complete (coNP-complete for universal quantifier) in the state set size of the input MDP.

Optimize the current algorithm for faster results.

Optimize the current algorithm for faster results.

Design less accurate model checking algorithms that scale better alongside providing probabilistic guarantees about the correctness of verification.

Optimize the current algorithm for faster results.

Design less accurate model checking algorithms that scale better alongside providing probabilistic guarantees about the correctness of verification.

Counter-example guided techniques to manage large state space.