

Parameter Synthesis for Probabilistic Hyperproperties

Erika *Ábrahám*, RWTH Aachen, Germany

Ezio Bartocci, TU Wien, Austria

Borzoo Bonakdarpour and Oyendrilá Dobe, Michigan State University

LPAR23 2020

13 January 2021

Trace properties and hyperproperties for reasoning about systems

[Clarkson, Schneider, 2010], [Clarkson, Finkbeiner, Koleini, Micinski, Rabe, Sánchez, 2014]

- 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.

Trace properties and hyperproperties for reasoning about systems

[Clarkson, Schneider, 2010], [Clarkson, Finkbeiner, Koleini, Micinski, Rabe, Sánchez, 2014]

- 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.

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

Trace properties and hyperproperties for reasoning about systems

[Clarkson, Schneider, 2010], [Clarkson, Finkbeiner, Koleini, Micinski, Rabe, Sánchez, 2014]

- 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.

Trace properties and hyperproperties for reasoning about systems

[Clarkson, Schneider, 2010], [Clarkson, Finkbeiner, Koleini, Micinski, Rabe, Sánchez, 2014]

- 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.

Trace properties and hyperproperties for reasoning about systems

[Clarkson, Schneider, 2010], [Clarkson, Finkbeiner, Koleini, Micinski, Rabe, Sánchez, 2014]

- 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)

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

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

HyperPCTL: PCTL extended with quantification over initial states

Example (Probabilistic noninterference)

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

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

HyperPCTL: PCTL extended with quantification over initial states

Example (Probabilistic noninterference)

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

Our contribution:

To synthesize **parametric configurations** such that the specified probabilistic **hyperproperty** is satisfied in a given **parametric deterministic probabilistic system** !

HyperPCTL is similar to PCTL **BUT**

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

$$Q_{\sigma_1} s_1 \dots Q_{\sigma_n} s_n. \psi$$

- they index atomic propositions:

$$\psi ::= a_{\sigma} \mid \psi \wedge \psi \mid \neg \psi \mid p \text{ op } p \quad \text{where op} = \{>, \geq, =, <, \leq\}$$

- they support arithmetic computations with probability expressions:

$$p ::= \mathbb{P}(\mathcal{X}\psi) \mid \mathbb{P}(\psi \mathcal{U} \psi) \mid c \mid p + p \mid p - p \mid p \cdot p$$

HyperPCTL is similar to PCTL **BUT**

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

$$Q_{\sigma_1} s_1 \dots Q_{\sigma_n} s_n. \psi$$

- they index atomic propositions:

$$\psi ::= a_{\sigma} \mid \psi \wedge \psi \mid \neg \psi \mid p \text{ op } p \quad \text{where op} = \{>, \geq, =, <, \leq\}$$

- they support arithmetic computations with probability expressions:

$$p ::= \mathbb{P}(\mathcal{X}\psi) \mid \mathbb{P}(\psi \mathcal{U} \psi) \mid c \mid p + p \mid p - p \mid p \cdot p$$

In this work, we use **ReachHyperPCTL** which is a **restrictive fragment** of HyperPCTL that does **not** allow **nested probability** operators.

Example





Figure: Parametric DTMC for randomized response protocol.


A **configuration** u for V , a finite set of parameters (here, p, q) is **valid** for the PDTMC, if,

$$\sum_{s' \in S} \mathbb{P}_u(s, s') = 1$$

Example




Example




DTMCs and their self-composition


Example



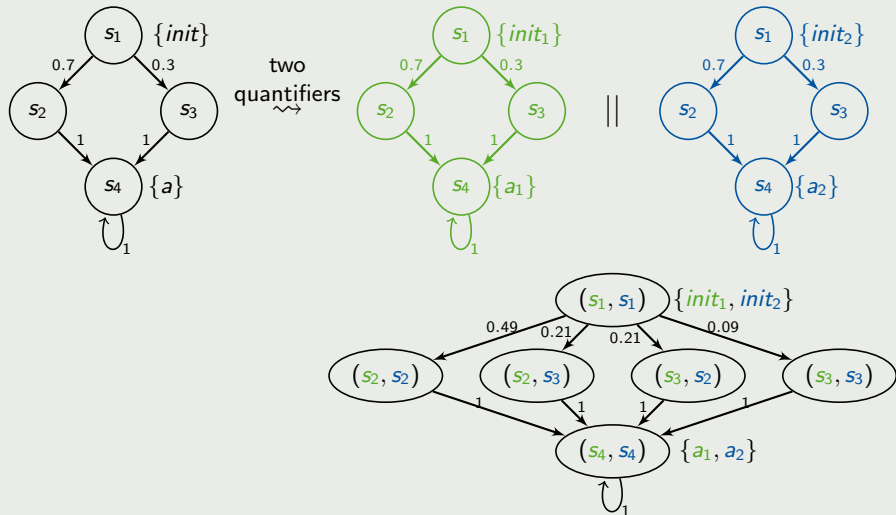
two
quantifiers
 \rightsquigarrow




||



Example




Example



$$\psi = \forall s. \forall s'. (init_s \wedge init_{s'}) \Rightarrow (\mathbb{P}(\mathcal{F}a_s) = \mathbb{P}(\mathcal{F}a_{s'}))$$


Example



$$\psi = \forall s. \forall s'. (init_s \wedge init_{s'}) \Rightarrow (\mathbb{P}(\mathcal{F}a_s) = \mathbb{P}(\mathcal{F}a_{s'}))$$

The probability of reaching a from s_0 is $0.4 + (0.2 \times 0.2) = 0.44$.

Example




$$\psi = \forall s. \forall s'. (init_s \wedge init_{s'}) \Rightarrow (\mathbb{P}(\mathcal{F}a_s) = \mathbb{P}(\mathcal{F}a_{s'}))$$

The probability of reaching a from s_0 is $0.4 + (0.2 \times 0.2) = 0.44$.

The probability of reaching a from s_1 is $0.3 + (0.7 \times 0.2) = 0.44$.

Example



$$\psi = \forall s. \forall s'. (init_s \wedge init_{s'}) \Rightarrow (\mathbb{P}(\mathcal{F}a_s) = \mathbb{P}(\mathcal{F}a_{s'}))$$

The probability of reaching a from s_0 is $0.4 + (0.2 \times 0.2) = 0.44$.

The probability of reaching a from s_1 is $0.3 + (0.7 \times 0.2) = 0.44$.

Hence, $\mathcal{M}, (s_0, s_1) \models \psi$

Given a PDTMC, D , and a ReachHyperPCTL formula, $\psi = Q_{\sigma_1} s_1 \dots Q_{\sigma_n} s_n \cdot \psi'$, we aim to provide algorithms to,

- Synthesize valid parameter configurations for D such that ψ is satisfied.

Given a PDTMC, D , and a ReachHyperPCTL formula, $\psi = Q_{\sigma_1} s_1 \dots Q_{\sigma_n} s_n \cdot \psi'$, we aim to provide algorithms to,

- Synthesize valid parameter configurations for D such that ψ is satisfied.
- For a box R of the valid parameter configurations of D , point to regions where,
 - all configurations
 - some configurations
 - none of the configurations

satisfy ψ for D .

Given a PDTMC, D , and a ReachHyperPCTL formula, $\psi = Q_{\sigma_1} s_1 \dots Q_{\sigma_n} s_n \cdot \psi'$, we aim to provide algorithms to,

- Synthesize valid parameter configurations for D such that ψ is satisfied.
- For a box R of the valid parameter configurations of D , point to regions where,
 - all configurations
 - some configurations
 - none of the configurations

satisfy ψ for D .

Theorem

The problem of parameter synthesis and checking of parameter space are both **decidable for deterministic systems**.

We propose an symbolic encoding technique for solving the model checking problem, such that

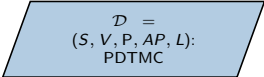
$$\mathcal{D} = (S, V, P, AP, L)$$

satisfies ψ in the region R

iff

$$\text{Symb}(\psi, s) \wedge R \text{ is satisfied}$$

Given a parameter space, we recursively divide it to find regions where the equation is **always satisfied**, **always violated**, or satisfied for some configurations.



$\mathcal{D} =$
(S, V, P, AP, L):
PDTMC

Overall Algorithm

$\mathcal{D} =$
 $(S, V, P, AP, L):$
PDTMC

$Q_{\sigma_1} s_1 \dots Q_{\sigma_n} s_n. \psi:$
ReachHyperPCTL


Overall Algorithm

R :
Parameter
Space

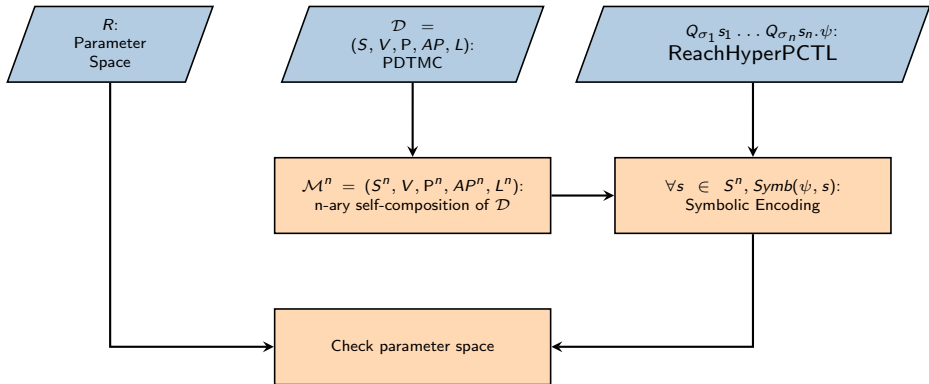
$\mathcal{D} =$
 (S, V, P, AP, L) :
PDTMC

$Q_{\sigma_1 s_1} \dots Q_{\sigma_n s_n} \psi$:
ReachHyperPCTL

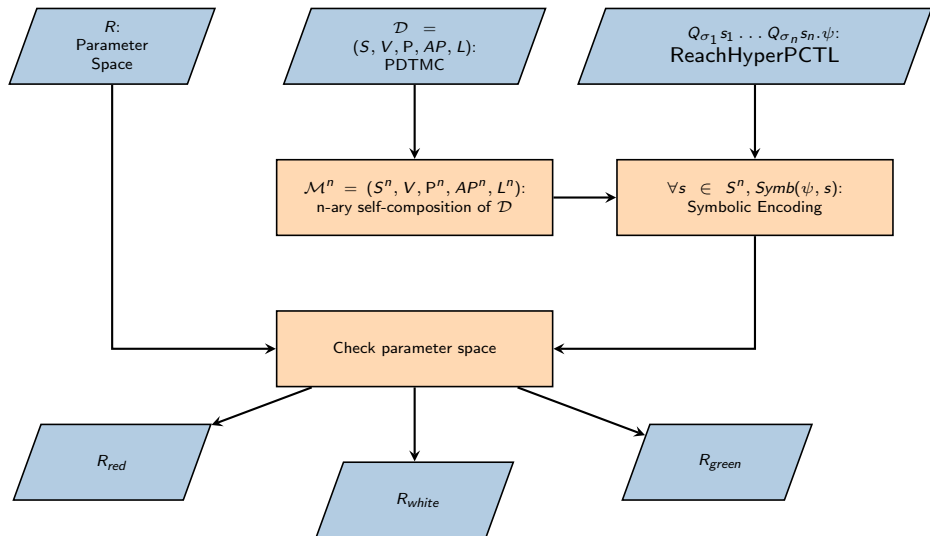
Overall Algorithm



Overall Algorithm




Overall Algorithm




We utilize the given \mathcal{D} and ReachHyperPCTL,

- to generate an encoding which is essentially a real-arithmetic formula $Symb(\psi, s)$ for each n-ary composed state, s .
- recursively parse the formula to calculate $Symb$.
- $Symb$ represents equations representing reachability probability in case of temporal operators.
- Solving these equations along with the parameter space allowed would give us parametric configurations that we aim to find.

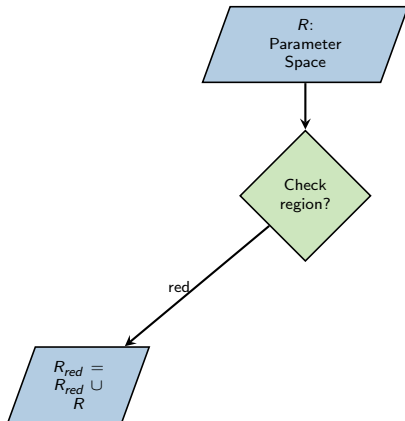
Algorithm for checking parameter space




Algorithm for checking parameter space



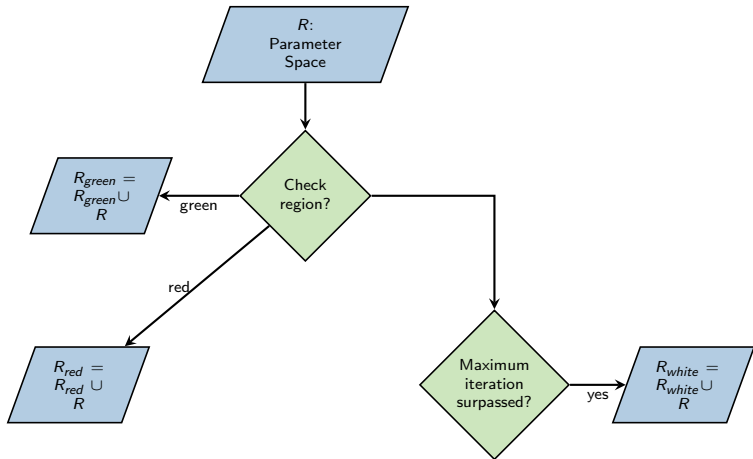
Algorithm for checking parameter space



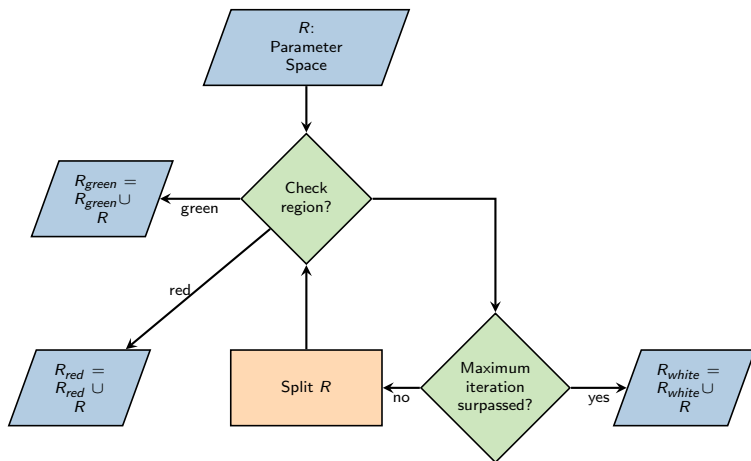
Algorithm for checking parameter space



Algorithm for checking parameter space





Algorithm for checking parameter space



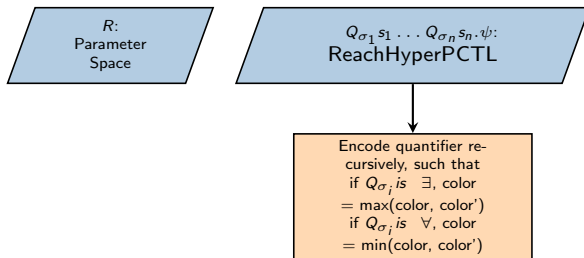
Algorithm for checking region

Here, we assume, green = 1, white = 0, red = -1.



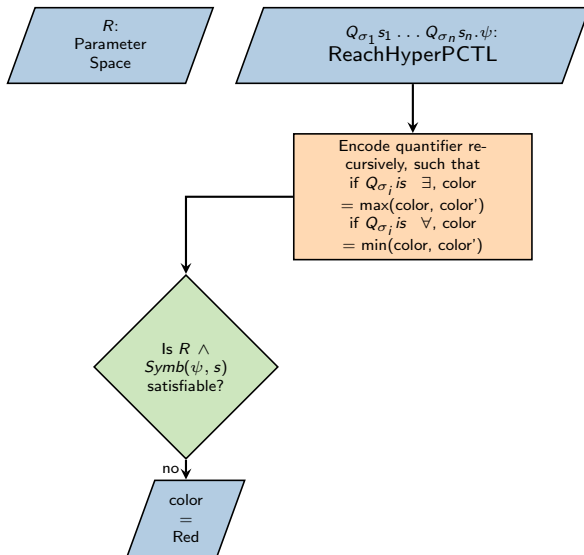
Algorithm for checking region

Here, we assume, green = 1, white = 0, red = -1.



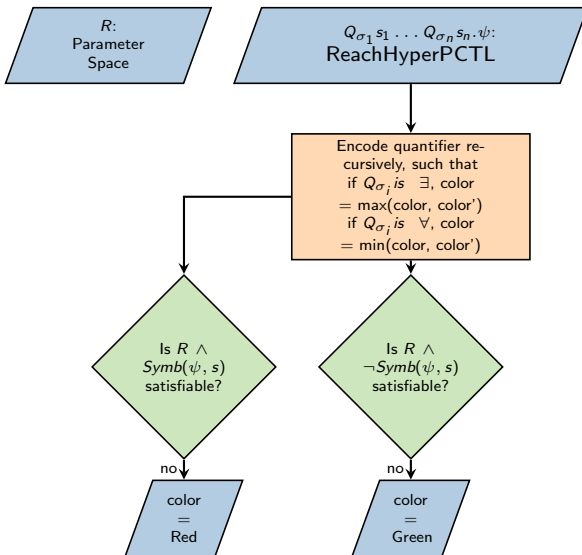
Algorithm for checking region

Here, we assume, green = 1, white = 0, red = -1.



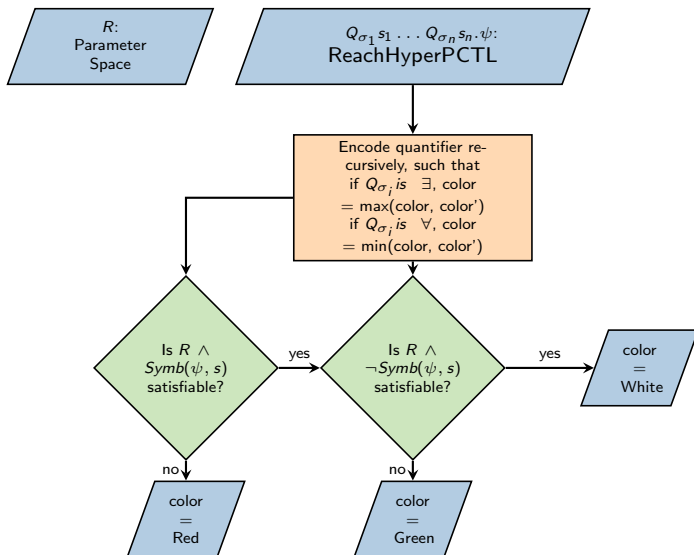
Algorithm for checking region

Here, we assume, green = 1, white = 0, red = -1.



Algorithm for checking region

Here, we assume, green = 1, white = 0, red = -1.



- Guarantees differential privacy by creating noise and providing plausible deniability.
- In this example, p is the parameter whose value we try to **calculate** depending on the ϵ we want to **achieve** in the privacy protocol.
- This algorithm should satisfy the following property:




Figure: Differential privacy

$$\varphi_{\text{dp}} = \forall \sigma. \forall \sigma'. \left[\left((t=n)_{\sigma} \wedge (t=y)_{\sigma'} \right) \Rightarrow \left(\mathbb{P}(\mathcal{F}(r=n)_{\sigma}) \leq e^{\epsilon} \cdot \mathbb{P}(\mathcal{F}(r=n)_{\sigma'}) \right) \right] \wedge$$

$$\left[\left((t=y)_{\sigma} \wedge (t=n)_{\sigma'} \right) \Rightarrow \left(\mathbb{P}(\mathcal{F}(r=y)_{\sigma}) \leq e^{\epsilon} \cdot \mathbb{P}(\mathcal{F}(r=y)_{\sigma'}) \right) \right]$$




Figure: Parameter Space for $(\ln 2)$ -differential privacy.




Figure: Parameter Space for $(\ln 3)$ -differential privacy.

We terminated the algorithm after 1500 rounds and the synthesis took 0.1 seconds.

- Check if **implementation** conforms with given **specification**.
- Synthesize a protocol such that the **reachability probability** is the **same** for the 5 final states when the **die** is rolled as well as when the die is simulated by **fair coin tosses**.




Figure: **Left:** 5-sided die
Right: parametric adaption of 5-sided die obtained by fair coin tosses

$$\varphi_{pc} = \forall \sigma. \exists \sigma'. (s_{0_\sigma} \wedge s'_{0_{\sigma'}}) \Rightarrow \bigwedge_{i=1}^5 (\mathbb{P}(\mathcal{F}d_{i_\sigma}) = \mathbb{P}(\mathcal{F}d'_{i_{\sigma'}}))$$

We synthesized the parameter $p = 0.5$ in 28 seconds. It is the only valid solution.

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 **distinguish** two computations from their **low observable outputs**, if their **high inputs** are **different**.
- It requires that the **probability** of every **low-observable** trace pattern is the **same** for every **low-equivalent initial state**.

$$\varphi_{\text{pni}} = \forall \sigma. \forall \sigma'. \left(i_{\sigma} \wedge (h=0)_{\sigma} \wedge i_{\sigma'} \wedge (h=1)_{\sigma'} \right) \Rightarrow \left(\left(\mathbb{P}(\mathcal{F}(f_{\sigma} \wedge (l=1)_{\sigma})) = \mathbb{P}(\mathcal{F}(f_{\sigma'} \wedge (l=1)_{\sigma'})) \right) \wedge \left(\mathbb{P}(\mathcal{F}(f_{\sigma} \wedge (l=2)_{\sigma})) = \mathbb{P}(\mathcal{F}(f_{\sigma'} \wedge (l=2)_{\sigma'})) \right) \right).$$

Probabilistic Noninterference in Randomized Schedulers - Results




Figure: Parameter Space for $(h = 3)_\sigma, (h = 5)_{\sigma'}$.




Figure: Parameter Space for $(h = 0)_\sigma, (h = 10)_{\sigma'}$.




Figure: Parameter Space for $(h = 14)_\sigma, (h = 8)_{\sigma'}$.

Input h	Running time (s)			#white boxes	#red boxes	red area	#samples
	Alg: Symbolic encoding	Alg: Check parameter space	Total				
(0, 1)	2.90	100.03	102.93	378	748	0.79	477
(0, 5)	15.61	143.58	159.2	374	752	0.815	421
(0, 10)	55.73	259.3	315.06	374	752	0.8164	480
(0, 15)	113.58	459.60	573.18	377	749	0.711	413
(1, 2)	8.33	114.55	122.88	368	758	0.706	425
(3, 5)	31.95	204.42	236.38	411	715	0.831	496
(4, 8)	72.23	397.91	470.14	371	755	0.6622	481
(8, 14)	213.96	2924.61	3138.07	378	748	0.825	496

Table: Experimental results for thread scheduling.

Information Leakage in Dining Cryptographer's protocol

- Three cryptographers having dinner, want to find out who paid for it, one of them or their boss, which respecting each other's privacy.
- They each flip a coin and inform the person on the right about its outcome.
- If the cryptographer didn't pay, he announces whether the two coin results he knows are the same ('a') or different, else he states the opposite.




Figure: PDTMC for dining cryptographer's protocol

$$\varphi_{dc} = \forall \sigma. \forall \sigma'. \left(\left(\bigvee_{i \in 3} \text{pay}_{\sigma}^i \right) \wedge \left(\bigvee_{i \in 3} \text{pay}_{\sigma'}^i \right) \right) \Rightarrow$$

$$\underbrace{\mathbb{P}\left(\mathcal{F}(\text{done}_{\sigma} \wedge (a_{\sigma}^1 \oplus a_{\sigma}^2 \oplus a_{\sigma}^3))\right)}_{F_1} = \underbrace{\mathbb{P}\left(\mathcal{F}(\text{done}_{\sigma'} \wedge (a_{\sigma'}^1 \oplus a_{\sigma'}^2 \oplus a_{\sigma'}^3))\right)}_{F_2}$$




Figure: Parameter space for dining cryptographer's protocol

- The whole parameter domain of $[0, 1]^3$ is green as the property is satisfied for all parameters.
- It took us approximately 40 minutes to get the result.
- This was due to the large state-space and the lack of recognition of non-reachable states in our implementation.

Provided algorithm to synthesize parameters for a given PDTMC such that it satisfies a given probabilistic hyperproperty.

Our algorithm works in two steps.

- It computes symbolic conditions for satisfying the formula, involving rational functions on the set of parameters.
- Identify regions of satisfying, unsatisfying, or grey parameter configurations by decomposing the domain of parameter configurations.

- Optimize the current prototypical implementation to lead to better scalability.
- Improve by exploiting the existence of symmetries.
- Extend the current work to more expressive logics like HyperPCTL and HyperPCTL*