Statistical Verification of Security-Related Hyperproperties

Yu Wang Miroslav Pajic

Cyber-Physical Systems Lab (CPSL)

Duke University

10/30/2020

Duke

PRATT SCHOOL of ENGINEERING

In collaboration with Siddhu Nalluri and Borzoo Bonakdarpoor

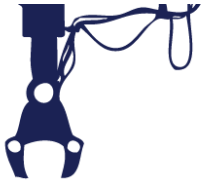

DUKE ROBOTICS

Verification by Statistical Model Checking

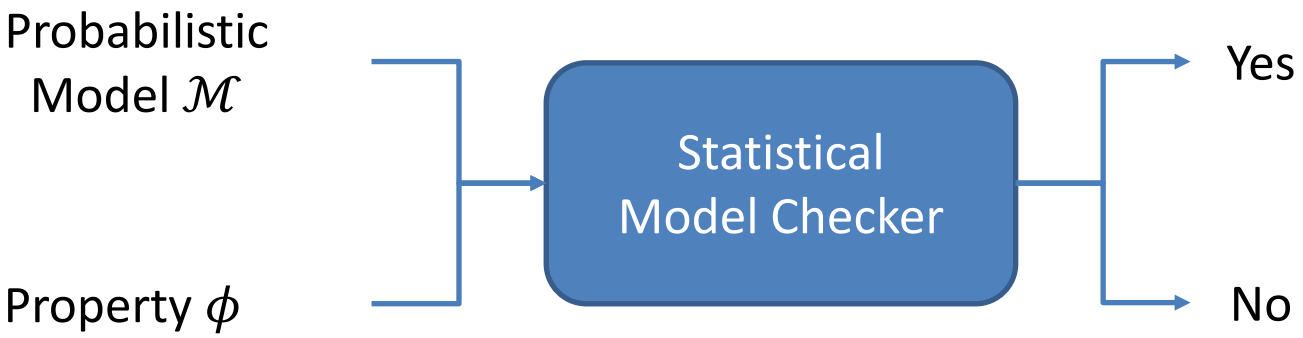

Advantages of SMC:
- Scalability
- General black-box systems

Motivation: Enable Statistical Model Checking for security-related **hyperproperties.**

Wang, Nalluri, Bonakdarpour, Pajic, Statistical model checking for hyperproperties, CSF 21

What is a Hyperproperty?

Duke
PRATT SCHOOL of ENGINEERING

Toyota Powertrain Benchmark

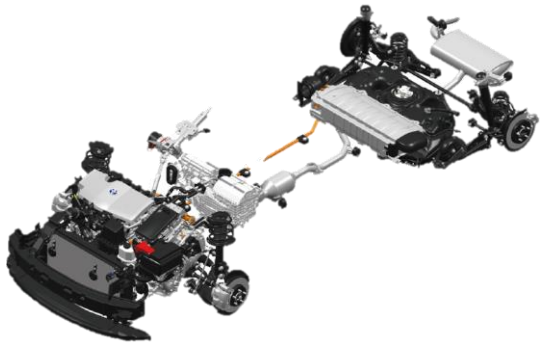

Embedded Computer

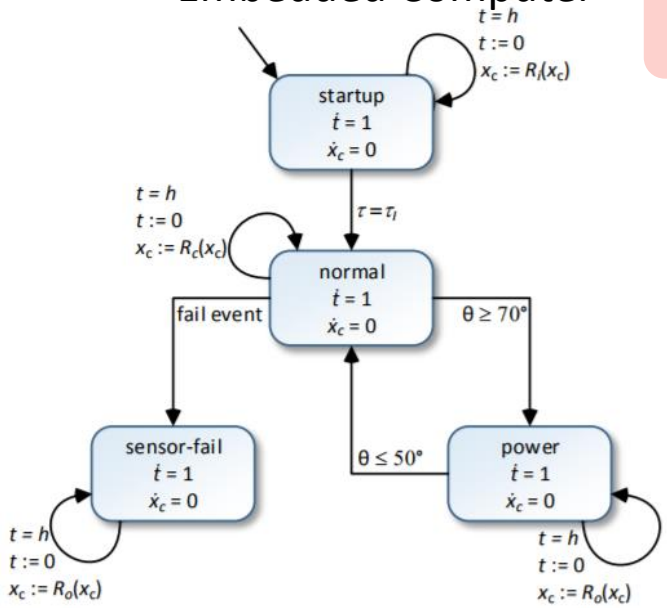

Hyperproperty: a property about the relation between multiple system paths.

AF ratio deviation:

$$e(t) = \frac{|\mu(t) - \mu_{\text{ref}}|}{\mu_{\text{ref}}}$$

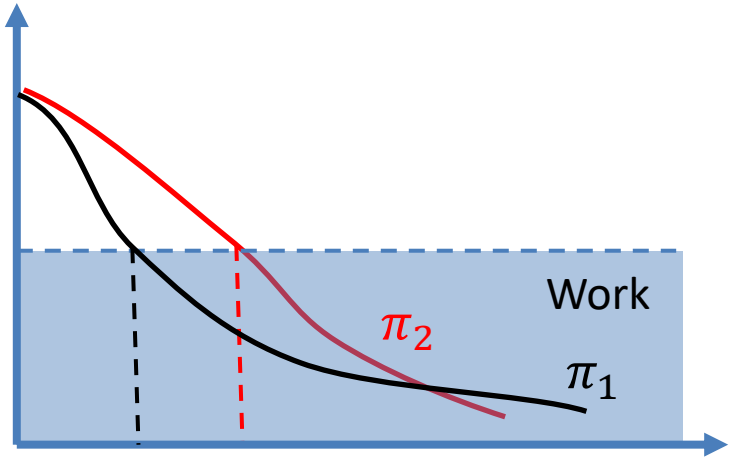

Combustion Process

State	Unit	Description
p	bar	Intake Manifold Pressure
λ_c	-	A/F Ratio in Cylinder
λ_m	-	Transfer Function Output
p_e	bar	Estimated Manifold Pressure
i	-	Integrator State, PI
$\dot{m}_{af}$	g/s	Inlet Air Mass Flow Rate
$\dot{m}_c$	g/s	Air Flow Rate to Cylinder
$\dot{m}_\phi$	g/s	Fuel Mass Aspirated into the Cylinder
$\dot{m}_\psi$	g/s	Fuel Mass Injected into Intake Manifold
θ_{in}	degrees	Throttle Angle Input
θ	degrees	Delay-Filtered Throttle Angle
$\hat{\theta}$	-	O/P of Throttle Polynomial
F_c	g/s	Command fuel
ω	rad/sec	Engine Speed
n	round/sec	Engine Speed ($\frac{\omega}{2\pi}$)

Example: Sensitivity under parameter change

$$\mathbf{Pr}_{\pi_1,\pi_2}(|\tau^{\pi_1} - \tau^{\pi_2}| \leq \delta) > 1 - \varepsilon$$

Hyperproperty in Discrete Time

Duke
PRATT SCHOOL of ENGINEERING

Probabilistic noninterference:

$$\mathbb{P}^{\pi_{h=1}}(\pi_{h=1} \text{ outputs } l = 1) \approx \mathbb{P}^{\pi_{h=2}}(\pi_{h=2} \text{ outputs } l = 1)$$

Example: Consider a parallel program $\boldsymbol{P}$ of two threads

$$\mathbf{th_1}\text{: while } h > 0 \text{ do } \{h \leftarrow h - 1; l \leftarrow 1\} \mid \mathbf{th_2}\text{: } l \leftarrow 2$$

where $h \in \{1, 2\}$ is hidden; and $l \in \{1, 2\}$ is public.

The CPU chooses to run one step of a random thread at each time.

- If $h = 1$, $\mathbf{th_1}$ has 1 steps, and $\mathbf{th_2}$ has 1 step. So when $\boldsymbol{P}$ stops, $l = 1$ w.p. $1/2$.
- If $h = 2$, $\mathbf{th_1}$ has 2 steps, and $\mathbf{th_2}$ has 1 step. So when $\boldsymbol{P}$ stops, $l = 1$ w.p. $1/3$.

So, the value of h will be (unintentionally) leaked by the value of l.

General criterion: deciding the correctness of a hyperproperty requires multiple path.

Our Work: Enabling SMC for Hyperproperties

Outline
1. New logics HyperPCTL* for discrete-time hyperproperties.
2. New SMC tool for HyperPCTL*.
3. Case Studies on challenging security problems.
4. Related work I: Continuous-time hyperproperties
5. Related work II: Conformance

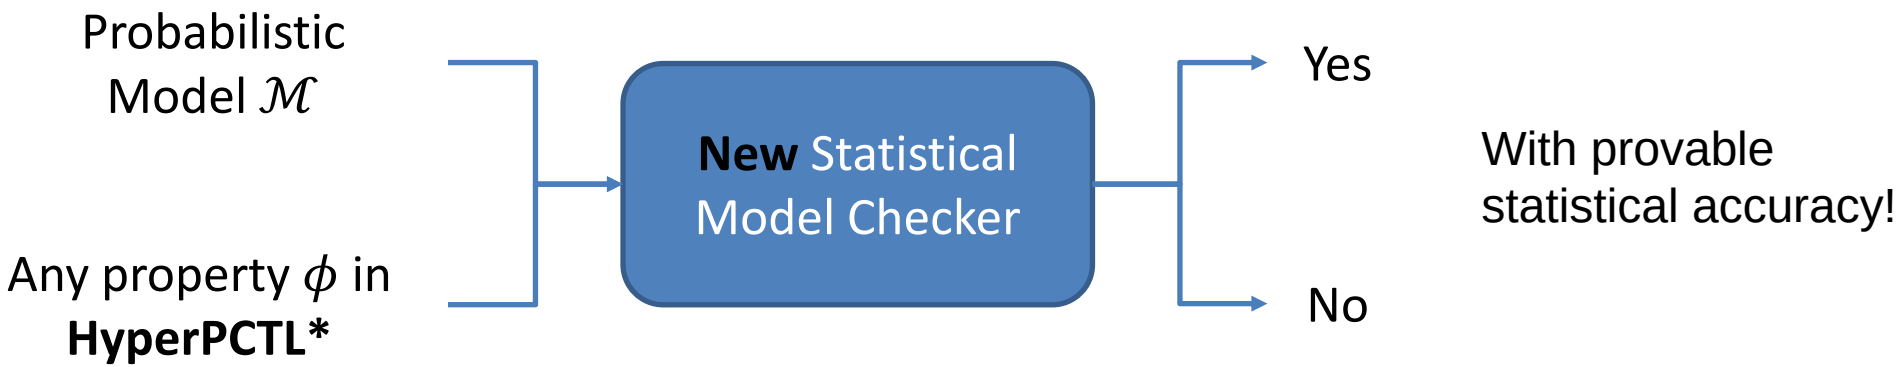

HyperPCTL*

Duke
PRATT SCHOOL of ENGINEERING

PCTL*: $\varphi ::= a \mid \neg\varphi \mid \varphi \wedge \varphi \mid \boldsymbol{X}\,\varphi \mid \varphi\ \boldsymbol{U_T}\ \varphi \mid \mathbb{P}_{\sim \boldsymbol{p}}\,\varphi$

- a is an **atomic proposition/event** (if a happens now);
- $\neg$ means "not"; $\wedge$ means "and";
- $\boldsymbol{X}\phi$ means ϕ holds next;
- $\phi_1 \mathbf{U}_T \phi_2$ means ϕ_1 holds **until** ϕ_2 becomes true;
- $\sim \in \{>, <, \geq, \leq\}$, $\mathbb{P}_{>\boldsymbol{p}}\,\varphi$ means ϕ holds with probability $> p$

Issue: all the events a and probability $\mathbb{P}_{>\boldsymbol{p}}$ are (implicitly) taken w.r.t. a single path.

HyperPCTL: $\varphi ::= a^{\pi} \mid \varphi^{\pi} \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi\ \boldsymbol{U_T}\ \varphi \mid p \sim p$

$$p ::= \mathbb{P}^{\Pi}\varphi \mid \mathbb{P}^{\Pi} p \mid f(p, \ldots, p)$$

- a replaced by a^{π}, π is a path variable,
- $\mathbb{P}$ replaced by $\mathbb{P}^{\Pi}$, Π is a set of path variables,

Example: Probabilistic Noninterference

$$\mathbb{P}^{\pi_1}\left((h = 1)_{\pi_1} \rightarrow \boldsymbol{F}(l = 1)_{\pi_1}\right) \approx \mathbb{P}^{\pi_2}\left((h = 2)_{\pi_2} \rightarrow \boldsymbol{F}(l = 1)_{\pi_2}\right)$$

- $\mathbb{P}_{\sim \boldsymbol{p}}\,\varphi$ replaced by a set of rules $p ::= \mathbb{P}^{\Pi}\varphi \mid \mathbb{P}^{\Pi} p \mid f(p, \ldots, p)$ and $p \sim p$; f is an elementary function

Example: KL-Divergence

$$KL\left(\mathbb{P}^{\pi_1}\left((h = 1)_{\pi_1} \rightarrow \boldsymbol{F}(l = 1)_{\pi_1}\right), \mathbb{P}^{\pi_2}\left((h = 2)_{\pi_2} \rightarrow \boldsymbol{F}(l = 1)_{\pi_2}\right)\right) < c$$

Theorems on HyperPCTL*

Duke
PRATT SCHOOL of ENGINEERING

Theorem 1: HyperPCTL* is well-defined.

Theorem 2: **HyperPCTL*** is strictly more expressive than **PCTL*.**

Statistical Model Checking for HyperPCTL*

Duke
PRATT SCHOOL of ENGINEERING

Model: Unknown Discrete-time Markov Chains

Challenges of SMC for HyperPCTL*

- Probabilistic quantifications of ***multiple*** parallel paths $\mathbb{P}^{(\pi_1,\pi_2)}\varphi^{(\pi_1,\pi_2)} < p$
- ***Nested*** probabilistic path quantification $\mathbb{P}^{\pi_1}\left(\mathbb{P}^{\pi_2}\varphi^{(\pi_1,\pi_2)} < p_2\right) < p_1$
- ***Joint*** probabilities $(\mathbb{P}^{\Pi_1}\varphi_1, \mathbb{P}^{\Pi_2}\varphi_2) \in D$

Statistical Model Checking for $(\mathbb{P}^{\Pi_1}\varphi_1^{\Pi_1}, \mathbb{P}^{\Pi_2}\varphi_2^{\Pi_2}) \in D$

Duke
PRATT SCHOOL of ENGINEERING

Assumption (Indifference region):
$(\mathbb{P}^{\Pi_1}\varphi_1^{\Pi_1}, \mathbb{P}^{\Pi_2}\varphi_2^{\Pi_2}) \notin D_1/D_0$
where $D_0 \subseteq D \subseteq D_1$

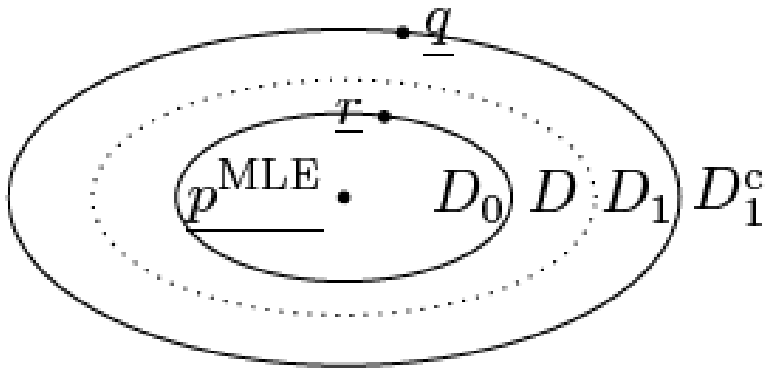

SMC as Hypothesis Testing:

$$\begin{cases} H_0 \text{ (True)}: (\mathbb{P}^{\Pi_1}\varphi_1^{\Pi_1}, \mathbb{P}^{\Pi_2}\varphi_2^{\Pi_2}) \in D_0 \\ H_1 \text{ (False)}: (\mathbb{P}^{\Pi_1}\varphi_1^{\Pi_1}, \mathbb{P}^{\Pi_2}\varphi_2^{\Pi_2}) \notin D_1 \end{cases}$$

No existing statistical method for multidimensional D_0 and D_1!

SMC for $(\mathbb{P}^{\Pi_1}\varphi_1^{\Pi_1}, \mathbb{P}^{\Pi_2}\varphi_2^{\Pi_2}) \in D$

Duke
PRATT SCHOOL of ENGINEERING

SMC as Hypothesis Testing:

$$\text{T1:} \begin{cases} H_0 \text{ (True)}: (\mathbb{P}^{\Pi_1}\varphi_1^{\Pi_1}, \mathbb{P}^{\Pi_2}\varphi_2^{\Pi_2}) \in D_0 \\ H_1 \text{ (False)}: (\mathbb{P}^{\Pi_1}\varphi_1^{\Pi_1}, \mathbb{P}^{\Pi_2}\varphi_2^{\Pi_2}) \notin D_1 \end{cases}$$

Simplified Hypothesis Testing:

$$\text{T2:} \begin{cases} H_0 \text{ (True)}: \left(\mathbb{P}^{\Pi_1}\varphi_1^{\Pi_1}, \mathbb{P}^{\Pi_2}\varphi_2^{\Pi_2}\right) = \underline{r} \\ H_1 \text{ (False)}: (\mathbb{P}^{\Pi_1}\varphi_1^{\Pi_1}, \mathbb{P}^{\Pi_2}\varphi_2^{\Pi_2}) = \underline{q} \end{cases}$$

Vector $\underline{r}, \underline{q}$ are derided from D_0, D_1, $\underline{p^{MLE}}$

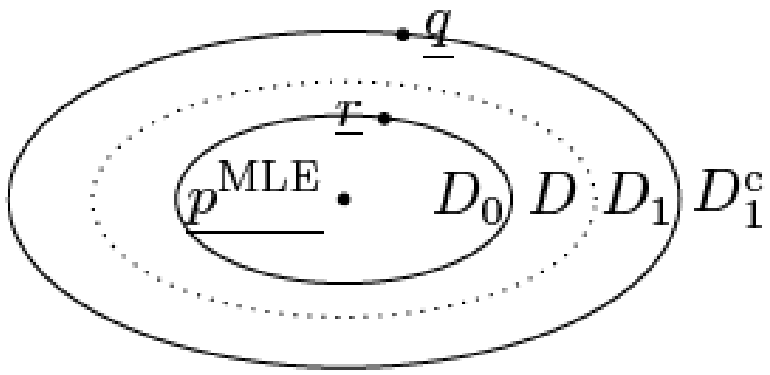

Theorem 1: T2 → T1

Remark 1: $\left(\mathbb{P}^{\Pi_1}\varphi_1^{\Pi_1}, \mathbb{P}^{\Pi_2}\varphi_2^{\Pi_2}\right)$ may be neither r or q.

Remark 2: r or q may change with samples

Statistical Model Checking for HyperPCTL*

Duke
PRATT SCHOOL of ENGINEERING

Challenges of SMC for HyperPCTL*

- Probabilistic quantifications of ***multiple*** parallel paths $\mathbb{P}^{(\pi_1,\pi_2)}\varphi^{(\pi_1,\pi_2)} < p$
- ***Nested*** probabilistic path quantification $\mathbb{P}^{\pi_1}\left(\mathbb{P}^{\pi_2}\varphi^{(\pi_1,\pi_2)} < p_2\right) < p_1$
- ***Joint*** probabilities $(\mathbb{P}^{\Pi_1}\varphi_1, \mathbb{P}^{\Pi_2}\varphi_2) \in D$

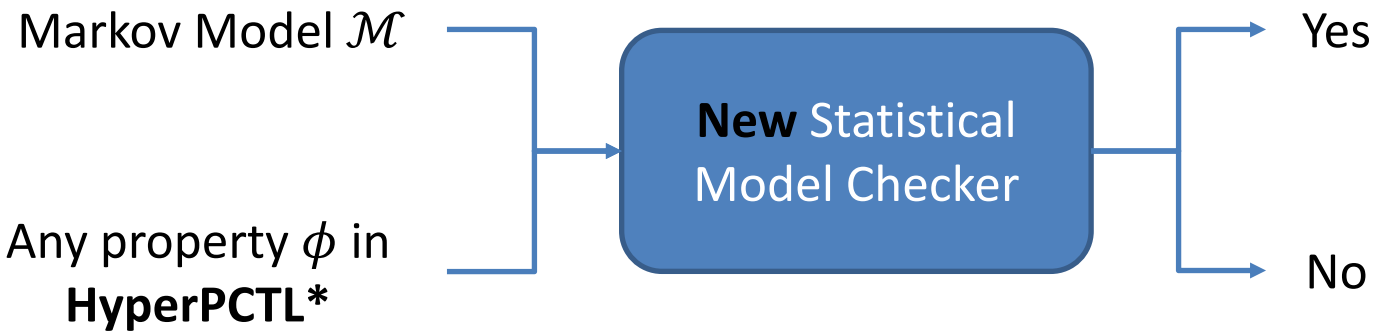

- Hyperproperties
- Provable significance level

HyperSMC Case Studies I

Duke
PRATT SCHOOL of ENGINEERING

GabFeed

- Chat server with encryption
- Encryption time may depend on private key
- Check encryption time for several baskets of keys
- We verified a time side-channel

$$\mathbb{P}^{\pi_1}\left((\bigcirc \mathrm{S}_1^{\pi_1}) \Rightarrow (\Diamond^{\leq k}\, \mathrm{F}^{\pi_1})\right) \approx_\varepsilon \mathbb{P}^{\pi_2}\left((\bigcirc \mathrm{S}_2^{\pi_2}) \Rightarrow (\Diamond^{\leq k}\, \mathrm{F}^{\pi_2})\right)$$

Horizon	ε	Significance	Acc.	No. Samples	Time (s)
60	0.05	0.01	1.00	5.5e+02	0.54
60	0.05	0.001	1.00	5.5e+03	5.76
60	0.1	0.01	1.00	6.1e+02	0.60
60	0.1	0.001	1.00	6.2e+03	7.16
90	0.05	0.01	1.00	3.7e+02	0.46
90	0.05	0.001	1.00	3.7e+03	4.94
90	0.1	0.01	1.00	4.1e+02	0.48
90	0.1	0.001	1.00	4.1e+03	5.37
120	0.05	0.01	1.00	3.8e+02	6.96
120	0.05	0.001	1.00	2.2e+03	11.24
120	0.1	0.01	1.00	3.8e+02	6.05
120	0.1	0.001	1.00	2.3e+03	9.46

HyperSMC Case Studies II

Duke
PRATT SCHOOL of ENGINEERING

Parallel Program with N threads

- Markov model of $N!$ States
- We verified probabilistic interference

$$\mathbb{P}^{\pi_1}\left((\bigcirc \mathrm{H_0}^{\pi_1}) \Rightarrow \left(\Diamond(\mathrm{F}^{\pi_1} \wedge \mathrm{L_0}^{\pi_1})\right)\right)$$
$$\approx_{\varepsilon} \mathbb{P}^{\pi_2}\left((\bigcirc \mathrm{H_1}^{\pi_2}) \Rightarrow \left(\Diamond(\mathrm{F}^{\pi_2} \wedge \mathrm{L_0}^{\pi_2})\right)\right)$$

Threads	Significance	Acc.	No. Samples	Time (s)
20	0.01	1.00	7.7e+02	0.49
20	0.001	1.00	7.6e+03	6.45
50	0.01	1.00	7.0e+02	0.48
50	0.001	1.00	6.8e+03	6.39
100	0.01	1.00	6.5e+02	0.54
100	0.001	1.00	6.6e+03	7.10

HyperSMC Case Studies III

Duke
PRATT SCHOOL of ENGINEERING

Dining Cryptographers

- [Chaum, 80s]
- Markov model of at least 2^N states
- We verified information security

$$\mathbb{P}^{\pi_1}\left(\Diamond(\neg \mathrm{S}_{ij}^{\pi_1} \wedge \Diamond \mathrm{P}^{\pi_1})\right) \approx_\varepsilon \mathbb{P}^{\pi_2}\left(\Diamond(\mathrm{S}_{ij}^{\pi_2} \wedge \Diamond \mathrm{P}^{\pi_2})\right) \approx_\varepsilon \mathbb{P}^{\pi_3}\left(\Diamond(\neg \mathrm{S}_{ij}^{\pi_3} \wedge \Diamond \mathrm{P}^{\pi_3})\right) \approx_\varepsilon \mathbb{P}^{\pi_4}\left(\Diamond(\mathrm{S}_{ij}^{\pi_4} \wedge \Diamond \mathrm{P}^{\pi_4})\right)$$

Agents	δ	Acc.	No. Samples	Time (s)
100	0.05	1.00	1.0e+03	0.91
100	0.1	1.00	5.2e+02	0.39
100	0.2	1.00	2.8e+02	0.14
1000	0.05	**0.98**	1.1e+03	3.27
1000	0.1	1.00	5.5e+02	1.52
1000	0.2	1.00	2.8e+02	0.69

[Significance level 0.01]

HyperSMC Case Studies IV

Duke
PRATT SCHOOL of ENGINEERING

Randomized Cache Replacement Policy

- [Canones, Kopf, and Reineke, '17]
- Least recently used (LRU) is not secure
- Markov model of $N!$ States
- We verified security

$$\mathbb{P}^{\pi_1}(\bigcirc^{(N)} \square^{\leq T} \mathrm{H}^{\pi_1}) > \mathbb{P}^{\pi_2}(\bigcirc^{(N)} \varphi^{\pi_2}) + \varepsilon$$

$$\varphi^{\pi_2} = \left(\mathrm{M}^{\pi_2} \wedge \bigcirc \mathrm{H}^{\pi_2} \wedge \ldots \wedge \bigcirc^{(T-1)} \mathrm{H}^{\pi_2}\right) \vee \ldots \vee \left(\mathrm{H}^{\pi_2} \wedge \ldots \wedge \bigcirc^{(T-2)} \mathrm{H}^{\pi_2} \wedge \bigcirc^{(T-1)} \mathrm{M}^{\pi_2}\right)$$

Horizon	ε	Significance	Acc.	No. Samples	Time (s)
10	0.05	0.01	1.00	1.1e+02	0.13
10	0.05	0.001	1.00	1.0e+03	2.56
10	0.01	0.01	1.00	1.2e+02	0.14
10	0.01	0.001	1.00	1.2e+03	2.79
20	0.05	0.01	1.00	6.0e+02	1.49
20	0.05	0.001	1.00	6.2e+03	16.73
20	0.01	0.01	0.99	1.2e+03	2.97
20	0.01	0.001	1.00	1.1e+04	28.99

Related Work I: HyperPSTL

$$\mathbb{P}^{(\pi_1,\pi_2)}\left((\neg Q^{\pi_1} \wedge \neg Q^{\pi_2})\boldsymbol{\mathcal{U}}\left(\left(Q^{\pi_1} \wedge \boldsymbol{\mathcal{F}}_{[0,\delta]} Q^{\pi_2}\right) \vee \boldsymbol{a}\left(Q^{\pi_2} \wedge \boldsymbol{\mathcal{F}}_{[0,\delta]} Q^{\pi_1}\right)\right)\right) > 1 - \varepsilon$$

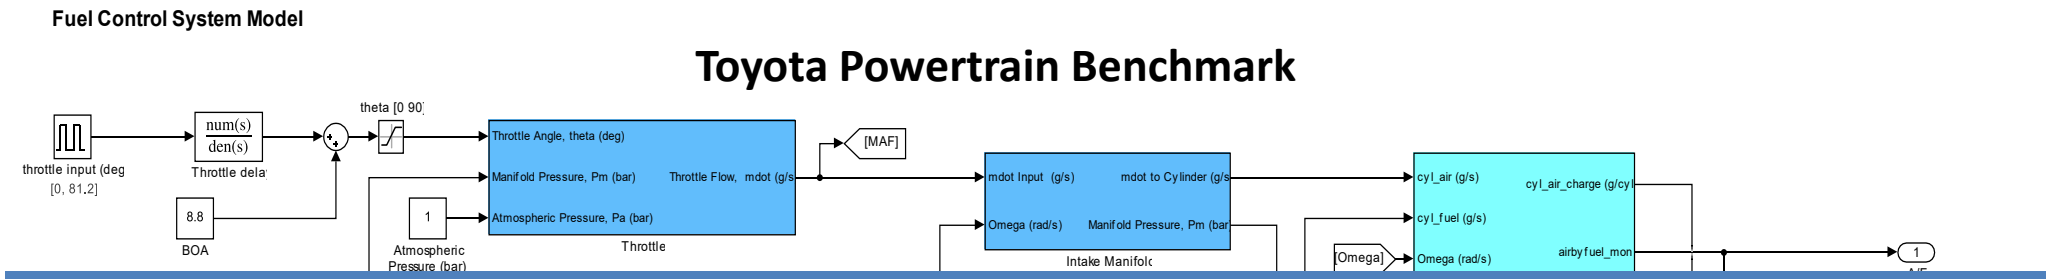

δ	ε	α	Acc.	Sam.	Time (s)	Ans.
0.15	0.95	0.05	**1.00**	5.9e+01	8.1e+00	True
0.15	0.95	0.01	**1.00**	9.0e+01	1.3e+01	True
0.15	0.99	0.05	**0.99**	6.6e+01	9.1e+00	False
0.15	0.99	0.01	**1.00**	9.7e+01	1.4e+01	False
0.20	0.95	0.05	**0.98**	5.9e+01	8.1e+00	True
0.20	0.95	0.01	**1.00**	9.0e+01	1.2e+01	True
0.20	0.99	0.05	**1.00**	3.0e+02	4.2e+01	True
0.20	0.99	0.01	**0.99**	4.6e+02	1.8e+02	True

Wang, Zarei, Bonakdarpour, Pajic, EMSOFT 19 **Best Paper Finalist**

Related Work I: HyperPSTL

$$\mathbb{P}^{\pi_1}(|\mathbb{P}^{\pi_2}((\neg \mathbb{Q}_i^{\pi_1} \wedge \neg \mathbb{Q}_j^{\pi_2}) \mathcal{U} (\mathbb{Q}_i^{\pi_1} \wedge \Diamond_{[\tau,\infty)} \mathbb{Q}_j^{\pi_2})) - \mathbb{P}^{\pi_2}((\neg \mathbb{Q}_i^{\pi_1} \wedge \neg \mathbb{Q}_j^{\pi_2}) \mathcal{U} (\mathbb{Q}_j^{\pi_2} \wedge \Diamond_{[\tau,\infty)} \mathbb{Q}_i^{\pi_1}))| \leq \delta) \geq 1 - \varepsilon.$$

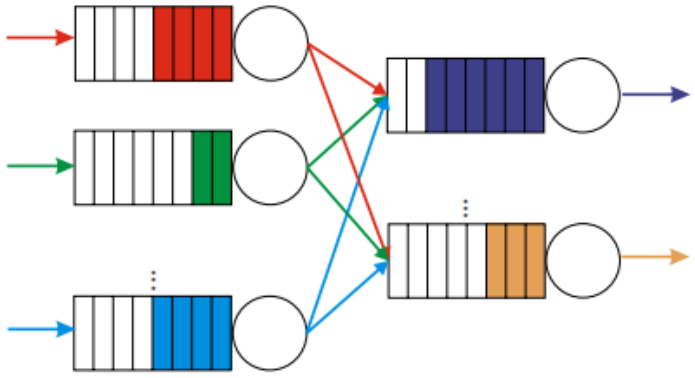

n servers *m* processors

Small: 1 server + 2 processers

t	δ	ε	α	Acc.	Sam.	Time (s)	Ans.
0.1	0.1	0.1	0.05	1.00	1.5e+02	8.0e+01	False
0.1	0.5	0.5	0.05	1.00	1.4e+02	8.2e+01	False
5.0	0.1	0.1	0.05	1.00	7.3e+02	3.9e+02	True
5.0	0.5	0.5	0.05	1.00	3.1e+01	1.9e+01	True

Queueing Network (Small)

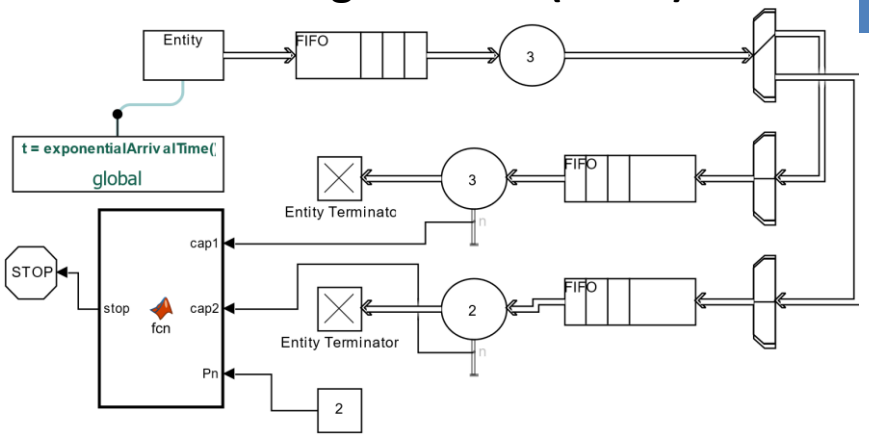

Large: 25 server + 20 processers

t	δ	ε	α	Acc.	Sam.	Time (s)	Ans.
0.1	0.1	0.1	0.05	1.00	2.1e+02	3.2e+02	False
0.1	0.5	0.5	0.05	1.00	3.3e+02	4.9e+02	False
5.0	0.1	0.1	0.05	1.00	6.8e+02	1.1e+03	True
5.0	0.5	0.5	0.05	1.00	4.2e+01	6.6e+01	True

Related Work II: Conformance as Hyperproperty

Duke
PRATT SCHOOL of ENGINEERING

Conformance: two systems simultaneously satisfy the same **set of** temporal logic specifications.

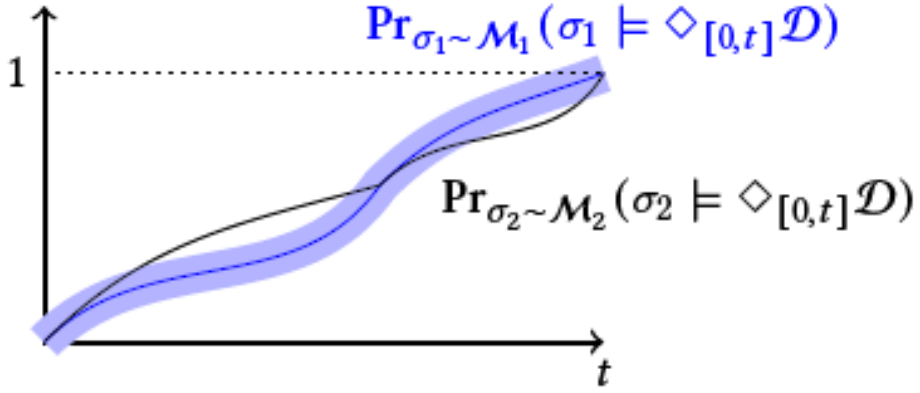

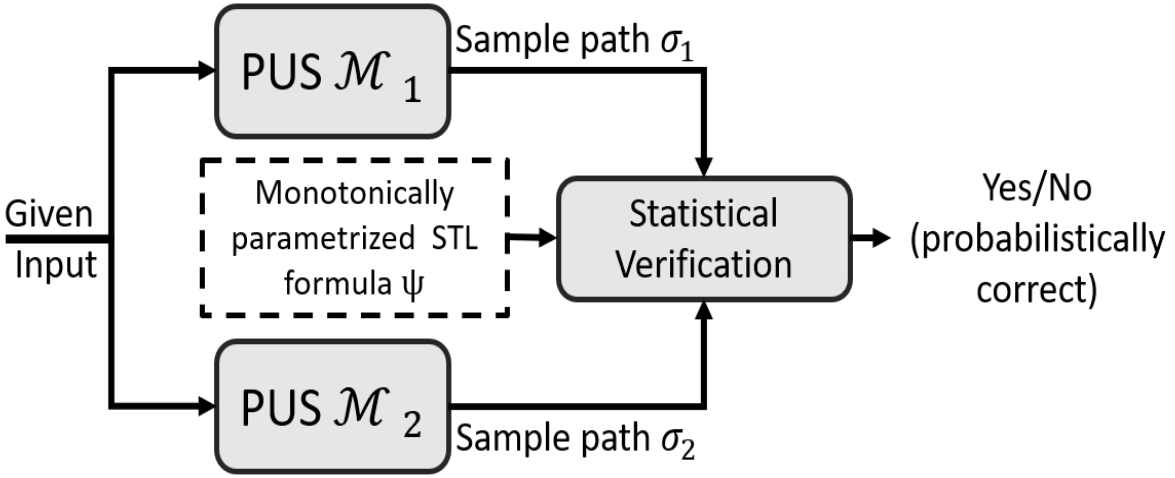

For any t, $\left|\mathbf{Pr}_{\sigma_1 \sim \mathcal{M}_1}(\sigma_1 \models \phi_t) - \mathbf{Pr}_{\sigma_2 \sim \mathcal{M}_2}(\sigma_2 \models \phi_t)\right| < c$.

Wang, Zarei, Bonakdarpour, Pajic, ICCPS 21 submitted

Application I

Duke
PRATT SCHOOL of ENGINEERING

Conformance of Simplified and Full Model of Toyota Powertrain

$$\forall \tau \geq 0. \quad \mathbf{Pr}_{\sigma_a \sim \mathcal{M}_a}(\sigma_a \vDash \Diamond_{[0.22,\tau]}(|e_{A/F}| < 0.05)) \underset{c}{\approx} \mathbf{Pr}_{\sigma_f \sim \mathcal{M}_f}(\sigma_f \vDash \Diamond_{[0.22,\tau]}(|e_{A/F}| < 0.05)).$$

c	α_d	$\delta_{n,m}$	Samples	Time (sec.)	$\mathcal{A}$
0.40	0.99	1.00	3.9e+01	1.8e-02	F
0.40	0.95	1.00	1.9e+01	4.4e-03	F
0.25	0.99	1.00	2.5e+01	4.6e-03	F
0.25	0.95	1.00	1.3e+01	2.2e-03	F
0.10	0.99	1.00	1.8e+01	3.6e-03	F
0.10	0.95	1.00	9.0e+00	1.6e-03	F
0.05	0.99	1.00	1.6e+01	2.8e-03	F
0.05	0.95	1.00	8.0e+00	1.3e-03	F

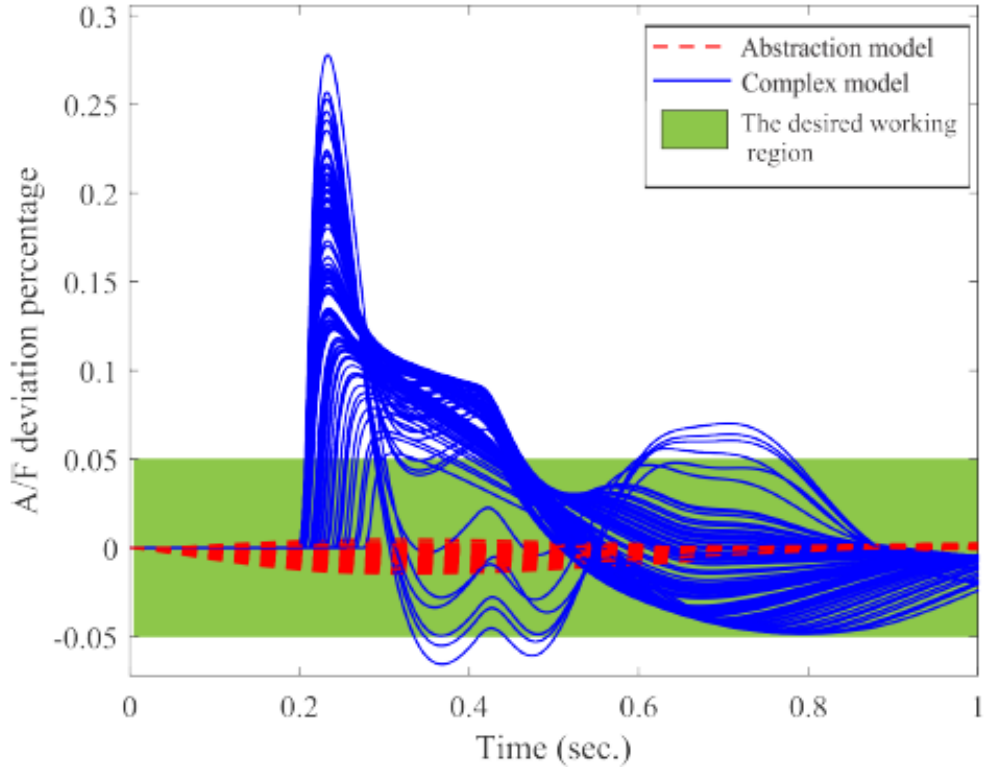

Application II

Duke
PRATT SCHOOL of ENGINEERING

Conformance of MPC and NN of Lane-Keeping System

$$\forall \tau \geq 0. \quad \mathbf{Pr}_{\sigma_1 \sim \mathcal{M}_{\mathrm{NN}}}(\sigma_1 \vDash \Diamond_{[0,\tau]}(|e_y^{\mathrm{NN}}| < \gamma)) \underset{c}{\approx} \mathbf{Pr}_{\sigma_2 \sim \mathcal{M}_{\mathrm{MPC}}}(\sigma_2 \vDash \Diamond_{[0,\tau]}(|e_y^{\mathrm{MPC}}| < \gamma)),$$

		NN (30 Neurons per Layer)				NN (45 Neurons per Layer)			
c	α_d	$\delta_{n,m}$	Samples	$T(s)$	$\mathcal{A}$	$\delta_{n,m}$	Samples	$T(s)$	$\mathcal{A}$
0.40	0.99	0.98	4.3e+01	7.4e-03	F	0.36	1.0e+04	9.6e+00	T
0.40	0.95	1.00	1.9e+01	3.1e-03	F	0.36	3.6e+03	2.0e+00	T
0.25	0.99	1.00	2.5e+01	4.1e-03	F	0.37	9.5e+02	3.2e-01	F
0.25	0.95	1.00	1.3e+01	2.1e-03	F	0.42	2.5e+02	5.9e-02	F
0.10	0.99	1.00	1.8e+01	3.0e-03	F	0.36	2.1e+02	4.2e-02	F
0.10	0.95	1.00	9.0e+00	1.6e-03	F	0.35	1.2e+02	2.2e-02	F
0.05	0.99	1.00	1.6e+01	2.7e-03	F	0.38	1.3e+02	2.5e-02	F
0.05	0.95	1.00	8.0e+00	1.2e-03	F	0.36	7.3e+01	1.4e-02	F

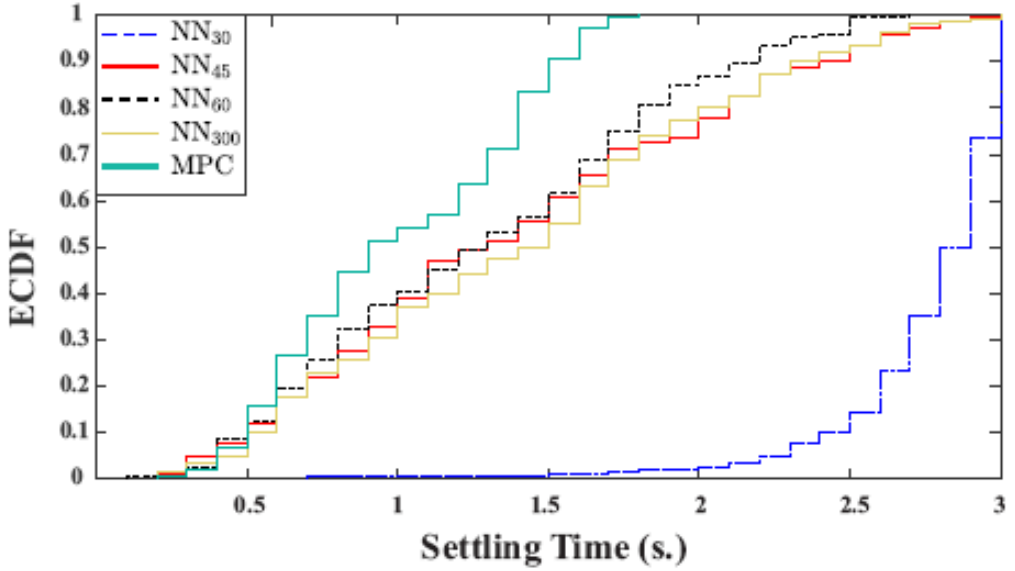

Thank you

Duke
PRATT SCHOOL of ENGINEERING

DUKE ROBOTICS