

Gray-box Monitoring of Hyperproperties

Sandro Stucki¹ César Sánchez²
Gerardo Schneider¹ Borzoo Bonakdarpour³

¹GU | Chalmers, Sweden ²IMDEA SW, Spain ³ISU, USA

FM '19, Porto, Portugal, 11 October 2019

sandro.stucki@gu.se @stuckintheory




UNIVERSITY OF
GOTHENBURG



CHALMERS
UNIVERSITY OF TECHNOLOGY

The monitorability cube



Motivation: distributed data minimality

- Distributed data minimality (DDM)
 - privacy property (GDPR)
 - generalization of data minimality to a multi-input setting

Motivation: distributed data minimality

- Distributed data minimality (DDM)
 - privacy property (GDPR)
 - generalization of data minimality to a multi-input setting
 - $\forall\forall\exists\exists$ -hyperproperty

$$\varphi_i = \forall\pi.\forall\pi'.\exists\tau.\exists\tau'. \neg\text{same}_i(\pi, \pi') \rightarrow \left(\begin{array}{l} \text{same}_i(\pi, \tau) \wedge \text{same}_i(\pi', \tau') \wedge \\ \text{almost}_i(\tau, \tau') \wedge \neg\text{output}(\tau, \tau') \end{array} \right)$$

Motivation: distributed data minimality

- Distributed data minimality (DDM)
 - privacy property (GDPR)
 - generalization of data minimality to a multi-input setting
 - $\forall\forall\exists\exists$ -hyperproperty

$$\varphi_i = \forall\pi.\forall\pi'.\exists\tau.\exists\tau'. \neg \text{same}_i(\pi, \pi') \rightarrow \left(\begin{array}{l} \text{same}_i(\pi, \tau) \wedge \text{same}_i(\pi', \tau') \wedge \\ \text{almost}_i(\tau, \tau') \wedge \neg \text{output}(\tau, \tau') \end{array} \right)$$

- Challenges:
 - Not black-box monitorable.
 - Undecidable.
 - Defined over arbitrary domains/datatypes.

Motivation: distributed data minimality

- Distributed data minimality (DDM)
 - privacy property (GDPR)
 - generalization of data minimality to a multi-input setting
 - $\forall\forall\exists\exists$ -hyperproperty

$$\varphi_i = \forall\pi.\forall\pi'.\exists\tau.\exists\tau'. \neg \text{same}_i(\pi, \pi') \rightarrow \left(\begin{array}{l} \text{same}_i(\pi, \tau) \wedge \text{same}_i(\pi', \tau') \wedge \\ \text{almost}_i(\tau, \tau') \wedge \neg \text{output}(\tau, \tau') \end{array} \right)$$

- Challenges:
 - Not black-box monitorable.
 - Undecidable.
 - Defined over arbitrary domains/datatypes.

Yet, we have a monitor. . .

Motivation: distributed data minimality

- Distributed data minimality (DDM)
 - privacy property (GDPR)
 - generalization of data minimality to a multi-input setting
 - $\forall\forall\exists\exists$ -hyperproperty


$$\varphi_i = \forall\pi.\forall\pi'.\exists\tau.\exists\tau'. \neg \text{same}_i(\pi, \pi') \rightarrow \left(\begin{array}{l} \text{same}_i(\pi, \tau) \wedge \text{same}_i(\pi', \tau') \wedge \\ \text{almost}_i(\tau, \tau') \wedge \neg \text{output}(\tau, \tau') \end{array} \right)$$

- Challenges:
 - Not black-box monitorable.
 - Undecidable.
 - Defined over arbitrary domains/datatypes.

Yet, we have a monitor. . .

what's going on here?

Trace properties – LTL





Trace properties – LTL

$$\varphi_s = \square \text{☕}$$

$$\varphi_l = \diamond \text{☕}$$

$$\varphi_r = \square \diamond \text{☕}$$

Trace properties – LTL

$$\varphi_s = \square \text{☕}$$

$$\varphi_l = \diamond \text{☕}$$

$$\varphi_r = \square \diamond \text{☕}$$

$t_1 = \text{☕☕☕☕☕☕} \dots$

$t_1 \models \varphi_s$

$t_1 \models \varphi_l$

$t_1 \models \varphi_r$

$t_2 = \text{☕☕☕☕☕☕} \dots$

$t_2 \not\models \varphi_s$

$t_2 \models \varphi_l$

$t_2 \not\models \varphi_r$

$t_3 = \text{☕☕☕☕☕☕} \dots$

$t_3 \not\models \varphi_s$

$t_3 \models \varphi_l$

$t_3 \models \varphi_r$

Trace properties – LTL

$$\varphi_s = \square \text{☕}$$

$$\varphi_l = \diamond \text{☕}$$

$$\varphi_r = \square \diamond \text{☕}$$

$$t_1 = \text{☕☕☕☕☕☕} \dots$$

$$t_1 \models \varphi_s$$

$$t_1 \models \varphi_l$$

$$t_1 \models \varphi_r$$

$$t_2 = \text{☕☕☕☕☕☕} \dots$$

$$t_2 \not\models \varphi_s$$

$$t_2 \models \varphi_l$$

$$t_2 \not\models \varphi_r$$

$$t_3 = \text{☕☕☕☕☕☕} \dots$$

$$t_3 \not\models \varphi_s$$

$$t_3 \models \varphi_l$$

$$t_3 \models \varphi_r$$

$$\varphi ::= a \mid \neg\varphi \mid \varphi \vee \varphi \mid \bigcirc\varphi \mid \varphi \mathcal{U} \varphi$$

$$\diamond\varphi \equiv \text{true } \mathcal{U} \varphi$$

$$\square\varphi \equiv \neg\diamond\neg\varphi$$

$$t \models p$$

$$\text{iff } p \in t[0]$$

$$t \models \neg\varphi$$

$$\text{iff } t \not\models \varphi$$

$$t \models \varphi_1 \vee \varphi_2$$

$$\text{iff } t \models \varphi_1 \text{ or } t \models \varphi_2$$

$$t \models \bigcirc\varphi$$

$$\text{iff } t[1, \dots] \models \varphi$$

$$t \models \varphi_1 \mathcal{U} \varphi_2$$

$$\text{iff for some } i, t[i, \dots] \models \varphi_2 \text{ and for all } j < i, t[j, \dots] \models \varphi_1$$

Monitoring LTL

$$\varphi_s = \square \text{☕} \quad \varphi_l = \diamond \text{☕} \quad \varphi_r = \square \diamond \text{☕}$$

Monitoring LTL

$$\varphi_s = \square \text{☕} \quad \varphi_l = \diamond \text{☕} \quad \varphi_r = \square \diamond \text{☕}$$

- **Observation:** the world today at 10am

$$u_{10} = \text{☕☕☕☕}$$

Monitoring LTL

$$\varphi_s = \square \text{☕} \quad \varphi_l = \diamond \text{☕} \quad \varphi_r = \square \diamond \text{☕}$$

- **Observation:** the world today at 10am

$$u_{10} = \text{☕☕☕☕}$$

- **Update:** the world at 11am

$$u_{11} = \text{☕☕☕☕☕☕☕☕}$$

Monitoring LTL

$$\varphi_s = \square \text{☕} \quad \varphi_l = \diamond \text{☕} \quad \varphi_r = \square \diamond \text{☕}$$

- **Observation:** the world today at 10am

$$u_{10} = \text{☕☕☕☕}$$

- **Update:** the world at 11am

$$u_{11} = \text{☕☕☕☕☕☕☕☕}$$

φ_s Is there always coffee?

Monitoring LTL

$$\varphi_s = \square \text{☕} \quad \varphi_l = \diamond \text{☕} \quad \varphi_r = \square \diamond \text{☕}$$

- **Observation:** the world today at 10am

$$u_{10} = \text{☕☕☕☕}$$

- **Update:** the world at 11am

$$u_{11} = \text{☕☕☕☕☕☕☕☕}$$

φ_s Is there always coffee?

$u_{10} \rightarrow ?$

Monitoring LTL

$$\varphi_s = \square \text{☕} \quad \varphi_l = \diamond \text{☕} \quad \varphi_r = \square \diamond \text{☕}$$

- **Observation:** the world today at 10am

$$u_{10} = \text{☕☕☕☕}$$

- **Update:** the world at 11am

$$u_{11} = \text{☕☕☕☕☕☕☕☕}$$

φ_s Is there always coffee?

$u_{10} \rightarrow ?$, $u_{11} \rightarrow \text{✗}$

Monitoring LTL

$$\varphi_s = \square \text{☕} \quad \varphi_l = \diamond \text{☕} \quad \varphi_r = \square \diamond \text{☕}$$

- **Observation:** the world today at 10am

$$u_{10} = \text{☕☕☕☕}$$

- **Update:** the world at 11am

$$u_{11} = \text{☕☕☕☕☕☕☕☕}$$

φ_s Is there always coffee?

$u_{10} \rightarrow ?$, $u_{11} \rightarrow \times$

φ_l Is there eventually coffee?

$u_{10} \rightarrow \checkmark$, $u_{11} \rightarrow \checkmark$

Monitoring LTL

$$\varphi_s = \square \text{☕} \quad \varphi_l = \diamond \text{☕} \quad \varphi_r = \square \diamond \text{☕}$$

- **Observation:** the world today at 10am

$$u_{10} = \text{☕☕☕☕}$$

- **Update:** the world at 11am

$$u_{11} = \text{☕☕☕☕☕☕☕☕}$$

φ_s Is there always coffee?

$u_{10} \rightarrow ?$, $u_{11} \rightarrow \times$

φ_l Is there eventually coffee?

$u_{10} \rightarrow \checkmark$, $u_{11} \rightarrow \checkmark$

φ_r Is there always eventually coffee?

$u_{10} \rightarrow ?$, $u_{11} \rightarrow ?$

Monitoring LTL

Monitoring: decide whether a given property φ is **permanently** satisfied (✓), violated (✗), or neither (?), at runtime.

Monitoring LTL

Monitoring: decide whether a given property φ is permanently satisfied (✓), violated (✗), or neither (?), given a finite observation u .

Monitoring LTL

Monitoring: decide whether a given property φ is **permanently satisfied** (✓), violated (✗), or neither (?), given a **finite observation** u .

Definition

A finite observation u **permanently satisfies** (resp. **violates**) φ , if every infinite extension of u satisfies (resp. violates) φ :

$$u \models^s \varphi \quad \text{iff} \quad \text{for all } t \in \Sigma^\omega \text{ such that } u \preceq t, t \models \varphi$$

$$u \models^v \varphi \quad \text{iff} \quad \text{for all } t \in \Sigma^\omega \text{ such that } u \preceq t, t \not\models \varphi$$

Monitoring LTL

Monitoring: decide whether a given property φ is **permanently** satisfied (✓), violated (✗), or neither (?), given a **finite observation** u .

Definition

A finite observation u **permanently satisfies** (resp. **violates**) φ , if every infinite extension of u satisfies (resp. violates) φ :

$$u \models^s \varphi \quad \text{iff} \quad \text{for all } t \in \Sigma^\omega \text{ such that } u \preceq t, t \models \varphi$$

$$u \models^v \varphi \quad \text{iff} \quad \text{for all } t \in \Sigma^\omega \text{ such that } u \preceq t, t \not\models \varphi$$

$$u_{11} = \text{☕☕☕☕☕☕☕☕☕☕}$$

$$u_{11} \not\models^s \square \text{☕}$$

$$u_{11} \models^v \square \text{☕}$$

$$u_{11} \models^s \diamond \text{☕}$$

$$u_{11} \not\models^v \diamond \text{☕}$$

$$u_{11} \not\models^s \square \diamond \text{☕}$$

$$u_{11} \not\models^v \square \diamond \text{☕}$$

Monitors for LTL

Monitoring: decide whether a given property φ is **permanently** satisfied (✓), violated (✗), or neither (?), given a **finite observation** u .

Monitors for LTL

Monitoring: decide whether a given property φ is **permanently** satisfied (✓), violated (✗), or neither (?), given a **finite observation** u .

A **monitor** for a property φ is a **computable** function $M_\varphi: \Sigma^* \rightarrow \{\checkmark, \times, ?\}$ that decides a **verdict** for φ given a finite u .

Monitors for LTL

Monitoring: decide whether a given property φ is **permanently** satisfied (\checkmark), violated (\times), or neither ($?$), given a **finite observation** u .

A **monitor** for a property φ is a **computable** function $M_\varphi: \Sigma^* \rightarrow \{\checkmark, \times, ?\}$ that decides a **verdict** for φ given a finite u .

The monitor M_φ is **sound** if

$$u \models^s \varphi \quad \text{if} \quad M_\varphi(u) = \checkmark, \quad u \models^v \varphi \quad \text{if} \quad M_\varphi(u) = \times$$

Monitors for LTL

Monitoring: decide whether a given property φ is **permanently** satisfied (\checkmark), violated (\times), or neither ($?$), given a **finite observation** u .

A **monitor** for a property φ is a **computable** function $M_\varphi: \Sigma^* \rightarrow \{\checkmark, \times, ?\}$ that decides a **verdict** for φ given a finite u .

The monitor M_φ is **sound** if

$$u \models^s \varphi \quad \text{if} \quad M_\varphi(u) = \checkmark, \quad u \models^v \varphi \quad \text{if} \quad M_\varphi(u) = \times$$

The monitor M_φ is **complete** if

$$M_\varphi(u) = \checkmark \text{ if } u \models^s \varphi, \quad M_\varphi(u) = \times \text{ if } u \models^v \varphi, \quad M_\varphi(u) = ? \text{ o/w.}$$

Monitors for LTL

Monitoring: decide whether a given property φ is **permanently** satisfied (\checkmark), violated (\times), or neither ($?$), given a **finite observation** u .

A **monitor** for a property φ is a **computable** function $M_\varphi: \Sigma^* \rightarrow \{\checkmark, \times, ?\}$ that decides a **verdict** for φ given a finite u .

The monitor M_φ is **sound** if

$$u \models^s \varphi \quad \text{if} \quad M_\varphi(u) = \checkmark, \quad u \models^v \varphi \quad \text{if} \quad M_\varphi(u) = \times$$

The monitor M_φ is **complete** if

$$M_\varphi(u) = \checkmark \text{ if } u \models^s \varphi, \quad M_\varphi(u) = \times \text{ if } u \models^v \varphi, \quad M_\varphi(u) = ? \text{ o/w.}$$

Fact: every LTL formula has a sound and complete monitor.

Monitorability of LTL formulas

Monitoring: decide whether a given property φ is **permanently** satisfied (✓), violated (✗), or neither (?), given a **finite observation** u .

$$\varphi_r = \square \diamond \text{☕}$$

$$u_{11} = \text{☕☕☕☕☕☕☕☕☕☕}$$

$$u_{11} \not\models^s \varphi_r$$

$$u_{11} \not\models^v \varphi_r$$

Monitorability of LTL formulas

Monitoring: decide whether a given property φ is **permanently** satisfied (✓), violated (✗), or neither (?), given a **finite observation** u .

$$\varphi_r = \square \diamond \text{☕}$$

$$u_{11} = \text{☕☕☕☕☕☕☕☕☕☕}$$

$$u_{11} \not\models^s \varphi_r$$

$$u_{11} \not\models^v \varphi_r$$

Observation: $u \not\models^s \varphi_r$ and $u \not\models^v \varphi_r$ for any u .

Monitorability of LTL formulas

Monitoring: decide whether a given property φ is **permanently** satisfied (✓), violated (✗), or neither (?), given a **finite observation** u .

$$\varphi_r = \square \diamond \text{☕}$$

$$u_{11} = \text{☕☕☕☕☕☕☕☕☕☕}$$

$$u_{11} \not\models^s \varphi_r$$

$$u_{11} \not\models^v \varphi_r$$

Observation: $u \not\models^s \varphi_r$ and $u \not\models^v \varphi_r$ for any u .

There's no point in monitoring φ_r !

Monitorability of LTL formulas

Monitoring: decide whether a given property φ is **permanently** satisfied (✓), violated (✗), or neither (?), given a **finite observation** u .

$$\begin{aligned}\varphi_r &= \square \diamond \text{☕} & u_{11} &= \text{☕☕☕☕☕☕☕☕} \\ u_{11} &\not\models^s \varphi_r & u_{11} &\not\models^v \varphi_r\end{aligned}$$

Observation: $u \not\models^s \varphi_r$ and $u \not\models^v \varphi_r$ for any u .


There's no point in monitoring φ_r !

Definition (Pnueli & Zaks 2006)

A formula φ is (*semantically*) *monitorable* if every observation u has an extension $v \succeq u$, such that either $v \models^s \varphi$ or $v \models^v \varphi$.


LTL – Summary

- Properties defined over individual traces.
⇒ Properties describe sets of traces.
- Sound and complete monitors can be constructed for any formula.
- Not every formula is monitorable. For example,
 - safety and liveness properties are monitorable,
 - recurrence properties ($\square\diamond$) are not.



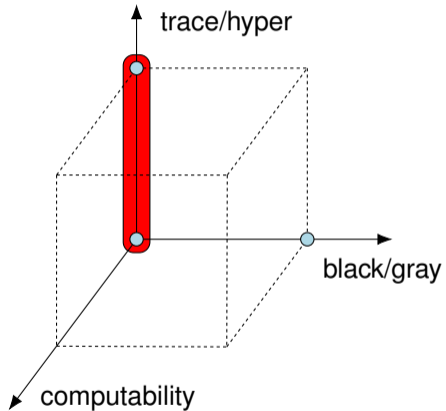
LTL – Summary

- Properties defined over individual traces.
⇒ Properties describe sets of traces.
- Sound and complete monitors can be constructed for any formula.
- Not every formula is monitorable. For example,
 - safety and liveness properties are monitorable,
 - recurrence properties ($\square\diamond$) are not.



- [9] A. Pnueli and A. Zaks. *PSL Model Checking and Run-time Verification via Testers.*, FM'06, Springer, 2006.
- [5] Y. Falcone, J-C. Fernandez, and L. Mounier. *What can you verify and enforce at runtime?*, STTT 14(3), 2012.
- [8] K. Havelund and D. Peled. *Runtime Verification: From Propositional to First-Order Temporal Logic.* RV'18, Springer, 2018.
- ... and many more!

Hyperproperties – HyperLTL



Hyperproperties – HyperLTL

$$\varphi_u = \forall \pi. \forall \tau. \square (\text{☕}_{\pi} \rightarrow \text{☕}_{\tau})$$

$$\varphi_a = \forall \pi. \exists \tau. \square (\text{☕}_{\pi} \rightarrow \text{☕}_{\tau})$$

Hyperproperties – HyperLTL

$$\varphi_u = \forall \pi. \forall \tau. \square (\text{☕}_\pi \rightarrow \text{☕}_\tau)$$

$$\varphi_a = \forall \pi. \exists \tau. \square (\text{☕}_\pi \rightarrow \text{☕}_\tau)$$

$$T_1 = \{\text{☕} \text{☕} \text{☕} \text{☕} \text{☕} \text{☕} \dots\}$$

$$T_1 \models \varphi_u \quad T_1 \models \varphi_a$$

$$T_2 = \{\text{☕} \text{☕} \text{☕} \text{☕} \text{☕} \text{☕} \dots, \text{☕} \text{☕} \text{☕} \text{☕} \text{☕} \text{☕} \dots\}$$

$$T_2 \not\models \varphi_u \quad T_2 \models \varphi_a$$

$$T_3 = \{\text{☕} \text{☕} \text{☕} \text{☕} \text{☕} \text{☕} \dots, \text{☕} \text{☕} \text{☕} \text{☕} \text{☕} \text{☕} \dots, \\ \text{☕} \text{☕} \text{☕} \text{☕} \text{☕} \text{☕} \dots, \text{☕} \text{☕} \text{☕} \text{☕} \text{☕} \text{☕} \dots, \dots\}$$

$$T_3 \not\models \varphi_u \quad T_3 \not\models \varphi_a$$

Hyperproperties – HyperLTL

$$\varphi_u = \forall \pi. \forall \tau. \square (\text{☕}_\pi \rightarrow \text{☕}_\tau)$$

$$\varphi_a = \forall \pi. \exists \tau. \square (\text{☕}_\pi \rightarrow \text{☕}_\tau)$$

$$T_1 = \{\text{☕}_1 \text{☕}_1 \text{☕}_1 \text{☕}_1 \text{☕}_1 \text{☕}_1 \dots\}$$

$$T_1 \models \varphi_u \quad T_1 \models \varphi_a$$

$$T_2 = \{\text{☕}_1 \text{☕}_1 \text{☕}_1 \text{☕}_1 \text{☕}_1 \text{☕}_1 \dots, \text{☕}_2 \text{☕}_2 \text{☕}_2 \text{☕}_2 \text{☕}_2 \text{☕}_2 \dots\}$$

$$T_2 \not\models \varphi_u \quad T_2 \models \varphi_a$$

$$T_3 = \{\text{☕}_1 \text{☕}_2 \text{☕}_1 \text{☕}_1 \text{☕}_1 \text{☕}_1 \dots, \text{☕}_2 \text{☕}_1 \text{☕}_1 \text{☕}_1 \text{☕}_1 \text{☕}_1 \dots, \\ \text{☕}_1 \text{☕}_1 \text{☕}_1 \text{☕}_1 \text{☕}_1 \text{☕}_1 \dots, \text{☕}_2 \text{☕}_2 \text{☕}_2 \text{☕}_2 \text{☕}_2 \text{☕}_2 \dots, \dots\}$$

$$T_3 \not\models \varphi_u \quad T_3 \not\models \varphi_a$$

$$\varphi ::= \forall \pi. \varphi \mid \exists \pi. \varphi \mid \psi$$

$$\psi ::= a_\pi \mid \neg \psi \mid \psi \vee \psi \mid \bigcirc \psi \mid \psi \mathcal{U} \psi$$

$\Pi \models a_\pi$	iff	$a \in \Pi(\pi)[0]$
$\Pi \models \psi_1 \vee \psi_2$	iff	$\Pi \models \psi_1$ or $\Pi \models \psi_2$
$\Pi \models \neg \psi$	iff	$\Pi \not\models \psi$
$\Pi \models \bigcirc \psi$	iff	$\Pi[1..] \models \psi$
$\Pi \models \psi_1 \mathcal{U} \psi_2$	iff	for some i , $\Pi[i, ..] \models \psi_2$, and for all $j < i$ $\Pi[j, ..] \models \psi_1$

$T, \Pi \models \forall \pi. \varphi$	iff	$T, \Pi[\pi \rightarrow t] \models \varphi$ for all $t \in T$
$T, \Pi \models \exists \pi. \varphi$	iff	$T, \Pi[\pi \rightarrow t] \models \varphi$ for some $t \in T$
$T, \Pi \models \psi$	iff	$\Pi \models \psi$

Monitoring HyperLTL

$$\varphi_u = \forall \pi. \forall \tau. \square (\text{☕}_\pi \rightarrow \text{☕}_\tau) \quad \varphi_a = \forall \pi. \exists \tau. \square (\text{☕}_\pi \rightarrow \text{☕}_\tau)$$

Monitoring HyperLTL

$$\varphi_u = \forall \pi. \forall \tau. \square (\text{☕}_\pi \rightarrow \text{☕}_\tau) \quad \varphi_a = \forall \pi. \exists \tau. \square (\text{☕}_\pi \rightarrow \text{☕}_\tau)$$

- **Observation:** the world today at 10am

$$U_{10} = \{\text{☕☕☕☕}\}$$

Monitoring HyperLTL

$$\varphi_u = \forall \pi. \forall \tau. \square (\text{☕}_\pi \rightarrow \text{☕}_\tau) \quad \varphi_a = \forall \pi. \exists \tau. \square (\text{☕}_\pi \rightarrow \text{☕}_\tau)$$

- **Observation:** the world today at 10am

$$U_{10} = \{\text{☕☕☕☕}\}$$

- **Update:** the world at 11am

$$U_{11} = \{\text{☕☕☕☕☕☕☕☕}, \text{☕☕☕☕☕☕☕☕}\}$$

Monitoring HyperLTL

$$\varphi_u = \forall \pi. \forall \tau. \square (\text{☕}_\pi \rightarrow \text{☕}_\tau) \quad \varphi_a = \forall \pi. \exists \tau. \square (\text{☕}_\pi \rightarrow \text{☕}_\tau)$$

- **Observation:** the world today at 10am

$$U_{10} = \{\text{☕☕☕☕}\}$$

- **Update:** the world at 11am

$$U_{11} = \{\text{☕☕☕☕☕☕☕☕}, \text{☕☕☕☕☕☕☕☕}, \text{☕☕☕}\}$$

Monitoring HyperLTL

$$\varphi_u = \forall \pi. \forall \tau. \square (\text{☕}_\pi \rightarrow \text{☕}_\tau) \quad \varphi_a = \forall \pi. \exists \tau. \square (\text{☕}_\pi \rightarrow \text{☕}_\tau)$$

- **Observation:** the world today at 10am

$$U_{10} = \{\text{☕☕☕☕}\}$$

- **Update:** the world at 11am

$$U_{11} = \{\text{☕☕☕☕☕☕☕☕}, \text{☕☕☕☕☕☕☕☕}, \text{☕☕☕}\}$$

φ_u Is there always coffee everywhere at the same time?

Monitoring HyperLTL

$$\varphi_u = \forall \pi. \forall \tau. \square (\text{☕}_\pi \rightarrow \text{☕}_\tau) \quad \varphi_a = \forall \pi. \exists \tau. \square (\text{☕}_\pi \rightarrow \text{☕}_\tau)$$

- **Observation:** the world today at 10am

$$U_{10} = \{\text{☕☕☕☕}\}$$

- **Update:** the world at 11am

$$U_{11} = \{\text{☕☕☕☕☕☕☕☕}, \text{☕☕☕☕☕☕☕☕}, \text{☕☕☕}\}$$

φ_u Is there always coffee everywhere at the same time? $U_{10} \rightarrow ?$,

Monitoring HyperLTL

$$\varphi_u = \forall \pi. \forall \tau. \Box (\text{☕}_\pi \rightarrow \text{☕}_\tau) \quad \varphi_a = \forall \pi. \exists \tau. \Box (\text{☕}_\pi \rightarrow \text{☕}_\tau)$$

- **Observation:** the world today at 10am

$$U_{10} = \{\text{☕☕☕☕}\}$$

- **Update:** the world at 11am

$$U_{11} = \{\text{☕☕☕☕☕☕☕☕}, \text{☕☕☕☕☕☕☕☕☕☕}, \text{☕☕☕}\}$$

φ_u Is there always coffee everywhere at the same time? $U_{10} \rightarrow ?$, $U_{11} \rightarrow \times$

Monitoring HyperLTL

$$\varphi_u = \forall \pi. \forall \tau. \square (\text{☕}_\pi \rightarrow \text{☕}_\tau) \quad \varphi_a = \forall \pi. \exists \tau. \square (\text{☕}_\pi \rightarrow \text{☕}_\tau)$$

- **Observation:** the world today at 10am

$$U_{10} = \{\text{☕}\text{☕}\text{☕}\text{☕}\}$$

- **Update:** the world at 11am

$$U_{11} = \{\text{☕}\text{☕}\text{☕}\text{☕}\text{☕}\text{☕}\text{☕}, \text{☕}\text{☕}\text{☕}\text{☕}\text{☕}\text{☕}\text{☕}, \text{☕}\text{☕}\text{☕}\}$$

φ_u Is there always coffee everywhere at the same time? $U_{10} \rightarrow ?$, $U_{11} \rightarrow \times$

φ_a Is there always coffee somewhere? $U_{10} \rightarrow ?$, $U_{11} \rightarrow ?$

Monitoring HyperLTL

Monitoring: decide whether a given property φ is permanently satisfied (✓), violated (✗), or neither (?), given a finite observation U .

Monitoring HyperLTL

Monitoring: decide whether a given property φ is **permanently satisfied** (✓), violated (✗), or neither (?), given a **finite observation** U .

Definition

A finite observation $U \in \mathcal{P}_{fin}(\Sigma^*)$ **permanently satisfies** (resp. **violates**) φ , if every infinite extension of U satisfies (resp. violates) φ :

$U \models^s \varphi$ iff for all $T \in \mathcal{P}(\Sigma^\omega)$ such that $U \preceq T$, $T \models \varphi$

$U \models^v \varphi$ iff for all $T \in \mathcal{P}(\Sigma^\omega)$ such that $U \preceq T$, $T \not\models \varphi$

Monitoring HyperLTL

Monitoring: decide whether a given property φ is **permanently** satisfied (✓), violated (✗), or neither (?), given a **finite observation** U .

Definition

A finite observation $U \in \mathcal{P}_{fin}(\Sigma^*)$ **permanently satisfies** (resp. **violates**) φ , if every infinite extension of U satisfies (resp. violates) φ :

$$U \models^s \varphi \quad \text{iff} \quad \text{for all } T \in \mathcal{P}(\Sigma^\omega) \text{ such that } U \preceq T, T \models \varphi$$

$$U \models^v \varphi \quad \text{iff} \quad \text{for all } T \in \mathcal{P}(\Sigma^\omega) \text{ such that } U \preceq T, T \not\models \varphi$$

$$U_{11} = \{ \text{☕☕☕☕☕☕☕☕}, \text{☕☕☕☕☕☕☕☕☕☕}, \text{☕☕☕☕☕☕☕☕} \}$$

$$U_{11} \not\models^s \forall \pi. \forall \tau. \square(\text{☕}_\pi \rightarrow \text{☕}_\tau)$$

$$U_{11} \not\models^s \forall \pi. \exists \tau. \square(\text{☕}_\pi \rightarrow \text{☕}_\tau)$$

$$U_{11} \models^v \forall \pi. \forall \tau. \square(\text{☕}_\pi \rightarrow \text{☕}_\tau)$$

$$U_{11} \not\models^v \forall \pi. \exists \tau. \square(\text{☕}_\pi \rightarrow \text{☕}_\tau)$$

Monitorability of HyperLTL formulas

Monitoring: decide whether a given property φ is **permanently** satisfied (✓), violated (✗), or neither (?), given a **finite observation** U .

$$U_{11} = \{ \text{☕☕☕☕☕☕☕☕}, \text{☕☕☕☕☕☕☕☕☕☕☕☕☕☕}, \text{☕☕☕☕☕} \}$$

$$\varphi_a = \forall \pi. \exists \tau. \square (\text{☕}_\pi \rightarrow \text{☕}_\tau) \quad U_{11} \not\models^s \varphi_a \quad U_{11} \not\models^v \varphi_a$$

Monitorability of HyperLTL formulas

Monitoring: decide whether a given property φ is **permanently** satisfied (✓), violated (✗), or neither (?), given a **finite observation** U .

$$U_{11} = \{ \text{☕☕☕☕☕☕☕☕}, \text{☕☕☕☕☕☕☕☕☕☕☕☕☕☕}, \text{☕☕☕☕☕} \}$$

$$\varphi_a = \forall \pi. \exists \tau. \square (\text{☕}_\pi \rightarrow \text{☕}_\tau) \quad U_{11} \not\models^s \varphi_a \quad U_{11} \not\models^v \varphi_a$$

Observation: $U \not\models^s \varphi_a$ and $U \not\models^v \varphi_a$ for any U .

Monitorability of HyperLTL formulas

Monitoring: decide whether a given property φ is **permanently** satisfied (✓), violated (✗), or neither (?), given a **finite observation** U .

$$U_{11} = \{ \text{☕☕☕☕☕☕☕☕}, \text{☕☕☕☕☕☕☕☕☕☕☕☕☕☕}, \text{☕☕☕☕☕} \}$$

$$\varphi_a = \forall \pi. \exists \tau. \square (\text{☕}_\pi \rightarrow \text{☕}_\tau) \quad U_{11} \not\models^s \varphi_a \quad U_{11} \not\models^v \varphi_a$$

Observation: $U \not\models^s \varphi_a$ and $U \not\models^v \varphi_a$ for any U .

There's no point in monitoring φ_a !

Monitorability of HyperLTL formulas

Monitoring: decide whether a given property φ is **permanently** satisfied (✓), violated (✗), or neither (?), given a **finite observation** U .

$$U_{11} = \{ \text{☕☕☕☕☕☕☕☕}, \text{☕☕☕☕☕☕☕☕☕☕☕☕☕☕}, \text{☕☕☕☕} \}$$

$$\varphi_a = \forall \pi. \exists \tau. \square (\text{☕}_\pi \rightarrow \text{☕}_\tau) \quad U_{11} \not\models^s \varphi_a \quad U_{11} \not\models^v \varphi_a$$

Observation: $U \not\models^s \varphi_a$ and $U \not\models^v \varphi_a$ for any U .


There's no point in monitoring φ_a !

Definition (Agrawal & Bonakdarpour 2016)

A formula φ is (*semantically*) *monitorable* if every observation U has an extension $V \succeq U$, such that $V \models^s \varphi$ or $V \models^v \varphi$.


HyperLTL – Summary

- Properties defined over sets of traces.
⇒ Properties describe sets of sets of traces.
- Sound and complete monitors can be constructed for some formulas.
 - For example, for formulas without **quantifier alternations**.
 - But what about formulas with alternations?
- Most formulas are not monitorable.
 - For example, $\forall^+\exists^+$ -properties are not!



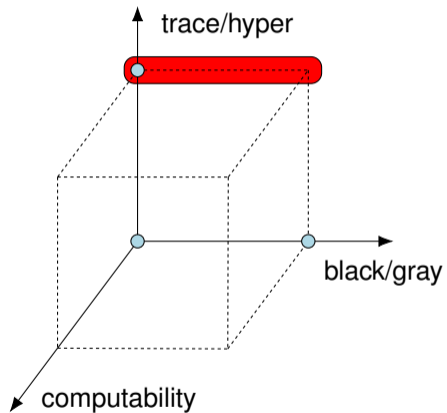
HyperLTL – Summary

- Properties defined over sets of traces.
⇒ Properties describe sets of sets of traces.
- Sound and complete monitors can be constructed for some formulas.
 - For example, for formulas without **quantifier alternations**.
 - But what about formulas with alternations?
- Most formulas are not monitorable.
 - For example, $\forall^+\exists^+$ -properties are not!



- [1] S. Agrawal and B. Bonakdarpour. *Runtime Verification of k -Safety Hyperproperties in HyperLTL*. CSF'16, IEEE CS Press, 2016.
- [8] K. Havelund and D. Peled. *Runtime Verification: From Propositional to First-Order Temporal Logic*. RV'18, Springer, 2018.
- [7] C. Hahn. *Algorithms for Monitoring Hyperproperties*. RV'19, Springer, 2019.

Gray-box monitoring (of hyperproperties)



Why is φ_a not monitorable?

Theorem

Let $\varphi_a = \forall \pi. \exists \tau. \square (\text{☕}_\pi \rightarrow \text{☕}_\tau)$. Then $U \not\models^s \varphi_a$ and $U \not\models^v \varphi_a$ for all $U \in \mathcal{P}_{fin}(\Sigma^*)$.

Why is φ_a not monitorable?

Theorem

Let $\varphi_a = \forall \pi. \exists \tau. \square (\text{☕}_\pi \rightarrow \text{☕}_\tau)$. Then $U \not\models^s \varphi_a$ and $U \not\models^v \varphi_a$ for all $U \in \mathcal{P}_{fin}(\Sigma^*)$.

Proof.

$U \not\models^v \varphi_a$ $U \preceq \Sigma^\omega$, and $\Sigma^\omega \models \varphi_a$ because $\text{☕}\text{☕}\text{☕}\dots \in \Sigma^\omega$;

Why is φ_a not monitorable?

Theorem

Let $\varphi_a = \forall \pi. \exists \tau. \square(\text{☕}_\pi \rightarrow \text{☕}_\tau)$. Then $U \not\models^s \varphi_a$ and $U \not\models^v \varphi_a$ for all $U \in \mathcal{P}_{fin}(\Sigma^*)$.

Proof.

$U \not\models^v \varphi_a$ $U \preceq \Sigma^\omega$, and $\Sigma^\omega \models \varphi_a$ because $\text{☕}\text{☕}\text{☕}\dots \in \Sigma^\omega$;

$U \not\models^s \varphi_a$ define T as $T = \{t \in \Sigma^\omega \mid w = u\text{☕}\text{☕}\text{☕}\dots \text{ for } u \in U\}$;
then $U \preceq T$ and $T \not\models \varphi_a$. □

Why is φ_a not monitorable?

Theorem

Let $\varphi_a = \forall \pi. \exists \tau. \Box(\text{☕}_\pi \rightarrow \text{☕}_\tau)$. Then $U \not\models^s \varphi_a$ and $U \not\models^v \varphi_a$ for all $U \in \mathcal{P}_{fin}(\Sigma^*)$.

Proof.

$U \not\models^v \varphi_a$ $U \preceq \Sigma^\omega$, and $\Sigma^\omega \models \varphi_a$ because $\text{☕☕☕} \dots \in \Sigma^\omega$;

$U \not\models^s \varphi_a$ define T as $T = \{t \in \Sigma^\omega \mid w = u\text{☕☕☕} \dots \text{ for } u \in U\}$;
then $U \preceq T$ and $T \not\models \varphi_a$. □

This theorem can be generalized to all formulas $\varphi = \forall \pi. \exists \tau. \Box P(\pi, \tau)$ where P is

- a binary (non-temporal) predicate,
- serial,
- non-reflexive.

Why is φ_a not monitorable?

Theorem

Let $\varphi_a = \forall \pi. \exists \tau. \Box(\text{☕}_\pi \rightarrow \text{☕}_\tau)$. Then $U \not\models^s \varphi_a$ and $U \not\models^v \varphi_a$ for all $U \in \mathcal{P}_{fin}(\Sigma^*)$.

Proof.

$U \not\models^v \varphi_a$ $U \preceq \Sigma^\omega$, and $\Sigma^\omega \models \varphi_a$ because $\text{☕☕☕} \dots \in \Sigma^\omega$;

$U \not\models^s \varphi_a$ define T as $T = \{t \in \Sigma^\omega \mid w = u\text{☕☕☕} \dots \text{ for } u \in U\}$;
then $U \preceq T$ and $T \not\models \varphi_a$. □

This theorem can be generalized to all formulas $\varphi = \forall \pi. \exists \tau. \Box P(\pi, \tau)$ where P is

- a binary (non-temporal) predicate,
- serial,
- non-reflexive.

OK, but let's have a closer look at this proof. . .

Why is φ_a not monitorable?

Theorem

Let $\varphi_a = \forall \pi. \exists \tau. \square(\text{☕}_\pi \rightarrow \text{☕}_\tau)$. Then $U \not\models^s \varphi_a$ and $U \not\models^v \varphi_a$ for all $U \in \mathcal{P}_{fin}(\Sigma^*)$.

Proof.

$U \not\models^v \varphi_a$ $U \preceq \Sigma^\omega$, and $\Sigma^\omega \models \varphi_a$ because ☕☕☕... $\in \Sigma^\omega$;

...

This step is somewhat dubious.

Why is φ_a not monitorable?

Theorem

Let $\varphi_a = \forall \pi. \exists \tau. \square(\text{☕}_\pi \rightarrow \text{☕}_\tau)$. Then $U \not\models^s \varphi_a$ and $U \not\models^v \varphi_a$ for all $U \in \mathcal{P}_{fin}(\Sigma^*)$.

Proof.

$U \not\models^v \varphi_a$ $U \preceq \Sigma^\omega$, and $\Sigma^\omega \models \varphi_a$ because ☕☕☕... $\in \Sigma^\omega$;

...

This step is somewhat dubious.

- Realistic systems don't realize every possible trace.

Why is φ_a not monitorable?

Theorem

Let $\varphi_a = \forall \pi. \exists \tau. \square(\text{☕}_\pi \rightarrow \text{☕}_\tau)$. Then $U \not\models^s \varphi_a$ and $U \not\models^v \varphi_a$ for all $U \in \mathcal{P}_{fin}(\Sigma^*)$.

Proof.

$U \not\models^v \varphi_a$ $U \preceq \Sigma^\omega$, and $\Sigma^\omega \models \varphi_a$ because ☕☕☕... $\in \Sigma^\omega$;

...

This step is somewhat dubious.

- Realistic systems don't realize every possible trace.
- There is only a finite number of coffee dispensers in the world (sadly).

Why is φ_a not monitorable?

Theorem

Let $\varphi_a = \forall \pi. \exists \tau. \square(\text{☕}_\pi \rightarrow \text{☕}_\tau)$. Then $U \not\models^s \varphi_a$ and $U \not\models^v \varphi_a$ for all $U \in \mathcal{P}_{fin}(\Sigma^*)$.

Proof.

$U \not\models^v \varphi_a$ $U \preceq \Sigma^\omega$, and $\Sigma^\omega \models \varphi_a$ because ☕☕☕... $\in \Sigma^\omega$;

...

This step is somewhat dubious.

- Realistic systems don't realize every possible trace.
- There is only a finite number of coffee dispensers in the world (sadly).

*When monitoring hyperproperties, we'd like to take into account
some information about the system
(gray-box monitoring).*

Gray-box monitoring of HyperLTL properties

Monitoring: decide whether a given property φ is **permanently satisfied** (✓), violated (✗), or neither (?), given a **finite observation** U .

Definition

A finite observation $U \in \mathcal{P}_{fin}(\Sigma^*)$ **permanently satisfies** (resp. **violates**) φ , if every infinite extension of U satisfies (resp. violates) φ :

$$U \models^s \varphi \quad \text{iff} \quad \text{for all } T \in \mathcal{P}(\Sigma^\omega) \text{ such that } U \preceq T, T \models \varphi$$

$$U \models^v \varphi \quad \text{iff} \quad \text{for all } T \in \mathcal{P}(\Sigma^\omega) \text{ such that } U \preceq T, T \not\models \varphi$$

Gray-box monitoring of HyperLTL properties

Monitoring: decide whether a given property φ is **permanently satisfied** (✓), violated (✗), or neither (?), given a **finite observation** U **of a system** \mathcal{S} .

Definition

A finite observation $U \in \mathcal{P}_{fin}(\Sigma^*)$ **permanently satisfies** (resp. **violates**) φ , if every infinite extension of U satisfies (resp. violates) φ :

$U \models^s \varphi$ iff for all $T \in \mathcal{P}(\Sigma^\omega)$ such that $U \preceq T$, $T \models \varphi$

$U \models^v \varphi$ iff for all $T \in \mathcal{P}(\Sigma^\omega)$ such that $U \preceq T$, $T \not\models \varphi$

Gray-box monitoring of HyperLTL properties

Monitoring: decide whether a given property φ is **permanently satisfied** (✓), **violated** (✗), or **neither** (?), given a **finite observation** U **of a system** \mathcal{S} .

Definition

Given a set of system behaviors $\mathcal{S} \subseteq \mathcal{P}(\Sigma^\omega)$,
a finite observation $U \in \mathcal{P}_{fin}(\Sigma^*)$ **permanently satisfies** (resp. **violates**) φ ,
if every infinite extension of U **in** \mathcal{S} satisfies (resp. violates) φ :

$$U \models_{\mathcal{S}}^s \varphi \quad \text{iff} \quad \text{for all } T \in \mathcal{S} \text{ such that } U \preceq T, T \models \varphi$$

$$U \models_{\mathcal{S}}^v \varphi \quad \text{iff} \quad \text{for all } T \in \mathcal{S} \text{ such that } U \preceq T, T \not\models \varphi$$

Gray-box monitoring of HyperLTL properties

Monitoring: decide whether a given property φ is **permanently satisfied** (✓), violated (✗), or neither (?), given a **finite observation** U **of a system** S .

Definition

Given a set of system behaviors $S \subseteq \mathcal{P}(\Sigma^\omega)$,
a finite observation $U \in \mathcal{P}_{fin}(\Sigma^*)$ **permanently satisfies** (resp. **violates**) φ ,
if every infinite extension of U **in** S satisfies (resp. violates) φ :

$$U \models_S^s \varphi \quad \text{iff} \quad \text{for all } T \in S \text{ such that } U \preceq T, T \models \varphi$$
$$U \models_S^v \varphi \quad \text{iff} \quad \text{for all } T \in S \text{ such that } U \preceq T, T \not\models \varphi$$

$$S = \{T \in \mathcal{P}(\Sigma^\omega) \mid |T| = 3\} \quad U = \{\text{☕☕☕☕☕}, \text{☕☕☕☕☕}, \text{☕☕☕☕☕}\}$$

$$U \not\models^s \forall \pi. \exists \tau. \square(\text{☕}_\pi \rightarrow \text{☕}_\tau)$$

$$U \models^v \forall \pi. \exists \tau. \square(\text{☕}_\pi \rightarrow \text{☕}_\tau)$$

Gray-box monitoring in general

Monitoring: decide whether a given property φ is **permanently** satisfied (✓), violated (✗), or neither (?), given a **finite observation** O of a system S .

Definition

Given a set of **system behaviors** $S \subseteq \mathcal{B}$,
a finite observation $O \in \mathcal{O}$ **permanently satisfies** (resp. **violates**) φ ,
if every infinite extension of O in S satisfies (resp. violates) φ :

$$O \models_S^s \varphi \quad \text{iff} \quad \text{for all } B \in S \text{ such that } O \preceq B, B \models \varphi$$

$$O \models_S^v \varphi \quad \text{iff} \quad \text{for all } B \in S \text{ such that } O \preceq B, B \not\models \varphi$$

Gray-box monitoring in general

Monitoring: decide whether a given property φ is **permanently satisfied** (✓), violated (✗), or neither (?), given a **finite observation** O of a system \mathcal{S} .

Definition

Given a set of **system behaviors** $\mathcal{S} \subseteq \mathcal{B}$,
a finite observation $O \in \mathcal{O}$ **permanently satisfies** (resp. **violates**) φ ,
if every infinite extension of O in \mathcal{S} satisfies (resp. violates) φ :

$$O \models_{\mathcal{S}}^s \varphi \quad \text{iff} \quad \text{for all } B \in \mathcal{S} \text{ such that } O \preceq B, B \models \varphi$$

$$O \models_{\mathcal{S}}^v \varphi \quad \text{iff} \quad \text{for all } B \in \mathcal{S} \text{ such that } O \preceq B, B \not\models \varphi$$

A formula φ is **semantically gray-box monitorable for a system \mathcal{S}** if every observation O has an extension $P \succeq O$ **in \mathcal{S}** , such that $P \models_{\mathcal{S}}^s \varphi$ or $P \models_{\mathcal{S}}^v \varphi$.

Gray-box monitors for $\forall^+\exists^+$ -properties

Monitoring: decide whether a given property φ is **permanently** satisfied (✓), violated (✗), or neither (?), given a **finite observation** O of a system \mathcal{S} .

Gray-box monitors for $\forall^+\exists^+$ -properties

Monitoring: decide whether a given property φ is **permanently satisfied** (✓), violated (✗), or neither (?), given a **finite observation** O of a system \mathcal{S} .

A **monitor** for a property φ **and a system** \mathcal{S} is a **computable** function $M_{\varphi, \mathcal{S}}: \mathcal{O} \rightarrow \{\checkmark, \times, ?\}$ that decides a **verdict** for φ given a finite O **in** \mathcal{S} .

Gray-box monitors for $\forall^+\exists^+$ -properties

Monitoring: decide whether a given property φ is **permanently satisfied** (✓), violated (✗), or neither (?), given a **finite observation** O of a system \mathcal{S} .

A **monitor** for a property φ **and a system** \mathcal{S} is a **computable** function $M_{\varphi, \mathcal{S}}: \mathcal{O} \rightarrow \{\checkmark, \times, ?\}$ that decides a **verdict** for φ given a finite O **in** \mathcal{S} .

Assuming $\varphi = \forall\pi.\exists\tau.\psi(\pi, \tau)$, and a sufficiently restrictive \mathcal{S} , we may be able to **statically prove** that all extensions $T \succeq U$ of a given U permanently violate φ .

Gray-box monitors for $\forall^+\exists^+$ -properties

Monitoring: decide whether a given property φ is **permanently satisfied** (✓), violated (✗), or neither (?), given a **finite observation** O of a system S .

A **monitor** for a property φ **and a system** S is a **computable** function $M_{\varphi,S}: \mathcal{O} \rightarrow \{\checkmark, \times, ?\}$ that decides a **verdict** for φ given a finite O **in** S .

Assuming $\varphi = \forall\pi.\exists\tau.\psi(\pi, \tau)$, and a sufficiently restrictive S , we may be able to **statically prove** that all extensions $T \succeq U$ of a given U permanently violate φ .

Example: $\varphi_a = \forall\pi.\exists\tau.\square(\text{☕}_\pi \rightarrow \text{☕}_\tau)$ $S = \{T \in \mathcal{P}(\Sigma^\omega) \mid |T| = 3\}$

Gray-box monitors for $\forall^+\exists^+$ -properties

Monitoring: decide whether a given property φ is **permanently satisfied** (✓), violated (✗), or neither (?), given a **finite observation** O of a system S .

A **monitor** for a property φ **and a system** S is a **computable** function $M_{\varphi,S}: \mathcal{O} \rightarrow \{\checkmark, \times, ?\}$ that decides a **verdict** for φ given a finite O **in** S .

Assuming $\varphi = \forall\pi.\exists\tau.\psi(\pi, \tau)$, and a sufficiently restrictive S , we may be able to **statically prove** that all extensions $T \succeq U$ of a given U permanently violate φ .

Example: $\varphi_a = \forall\pi.\exists\tau.\square(\text{☕}_\pi \rightarrow \text{☕}_\tau)$ $S = \{T \in \mathcal{P}(\Sigma^\omega) \mid |T| = 3\}$

Negate φ_a : $\neg\varphi_a = \exists\pi.\neg\exists\tau.\square(\text{☕}_\pi \rightarrow \text{☕}_\tau)$

Gray-box monitors for $\forall^+\exists^+$ -properties

Monitoring: decide whether a given property φ is **permanently satisfied** (✓), violated (✗), or neither (?), given a **finite observation** O of a system S .

A **monitor** for a property φ **and a system** S is a **computable** function $M_{\varphi,S}: \mathcal{O} \rightarrow \{\checkmark, \times, ?\}$ that decides a **verdict** for φ given a finite O **in** S .

Assuming $\varphi = \forall\pi.\exists\tau.\psi(\pi, \tau)$, and a sufficiently restrictive S , we may be able to **statically prove** that all extensions $T \succeq U$ of a given U permanently violate φ .

Example: $\varphi_a = \forall\pi.\exists\tau.\square(\text{☕}_\pi \rightarrow \text{☕}_\tau)$ $S = \{T \in \mathcal{P}(\Sigma^\omega) \mid |T| = 3\}$

Negate φ_a : $\neg\varphi_a = \exists\pi.\neg\exists\tau.\square(\text{☕}_\pi \rightarrow \text{☕}_\tau)$ **instantiate**

Gray-box monitors for $\forall^+\exists^+$ -properties

Monitoring: decide whether a given property φ is **permanently** satisfied (\checkmark), violated (\times), or neither ($?$), given a **finite observation** O of a system S .

A **monitor** for a property φ **and a system** S is a **computable** function $M_{\varphi,S}: \mathcal{O} \rightarrow \{\checkmark, \times, ?\}$ that decides a **verdict** for φ given a finite O **in** S .

Assuming $\varphi = \forall\pi.\exists\tau.\psi(\pi, \tau)$, and a sufficiently restrictive S , we may be able to **statically prove** that all extensions $T \succeq U$ of a given U permanently violate φ .

Example: $\varphi_a = \forall\pi.\exists\tau.\square(\text{☕}_\pi \rightarrow \text{☕}_\tau)$ $S = \{T \in \mathcal{P}(\Sigma^\omega) \mid |T| = 3\}$

Negate φ_a : $\neg\varphi_a = \exists\pi.\neg\exists\tau.\square(\text{☕}_\pi \rightarrow \text{☕}_\tau)$ **instantiate** **solve**

Gray-box monitors for $\forall^+\exists^+$ -properties

Monitoring: decide whether a given property φ is **permanently satisfied** (✓), violated (✗), or neither (?), given a **finite observation** O of a system S .

A **monitor** for a property φ and a system S is a **computable** function $M_{\varphi,S}: \mathcal{O} \rightarrow \{\checkmark, \times, ?\}$ that decides a **verdict** for φ given a finite O in S .

Assuming $\varphi = \forall\pi.\exists\tau.\psi(\pi, \tau)$, and a sufficiently restrictive S , we may be able to **statically prove** that all extensions $T \succeq U$ of a given U permanently violate φ .

Example: $\varphi_a = \forall\pi.\exists\tau.\square(\text{☕}_\pi \rightarrow \text{☕}_\tau)$ $S = \{T \in \mathcal{P}(\Sigma^\omega) \mid |T| = 3\}$


Negate φ_a : $\neg\varphi_a = \exists\pi.\neg\exists\tau.\square(\text{☕}_\pi \rightarrow \text{☕}_\tau)$ **instantiate** **solve**

$\{\text{☕☕}, \text{☕☕☕}\} \mapsto \{\text{☕☕}\dots, \text{☕☕☕}\dots, \underline{\text{☕☕☕}}\dots\} \quad ?$


$\{\text{☕☕☕}, \text{☕☕☕☕}, \text{☕☕☕☕☕}\} \mapsto \emptyset \quad \times$

Gray-box monitoring – Summary

- Properties defined over observations (e.g. traces or sets of traces).
⇒ Properties describe sets of observations.
- Sound and complete monitors can be constructed for some formulas.
 - For example, for formulas without quantifier alternations (as for black-box).
 - But also for $\forall^+\exists^+$ -formulas when \mathcal{S} imposes enough constraints.
- Monitorability of formulas depends on set of valid system behaviors \mathcal{S} .
 - For example, $\forall^+\exists^+$ -properties are monitorable for some choices of \mathcal{S} .
 - We will see a more interesting example later...



Undecidable hyperproperties



Monitorability is not existence of monitors

A formula φ is **semantically gray-box monitorable** for a system \mathcal{S} if every observation O has an extension $P \succeq O$ in \mathcal{S} , such that $P \models_{\mathcal{S}}^s \varphi$ or $P \models_{\mathcal{S}}^v \varphi$.

Monitorability is not existence of monitors

A formula φ is **semantically gray-box monitorable** for a system \mathcal{S} if every observation O has an extension $P \succeq O$ in \mathcal{S} , such that $P \models_{\mathcal{S}}^s \varphi$ or $P \models_{\mathcal{S}}^v \varphi$.

A **monitor** for a property φ and a system \mathcal{S} is a **computable** function $M_{\varphi, \mathcal{S}}: \mathcal{O}\{\checkmark, \times, ?\}$ that decides a **verdict** for φ given a finite u .

Monitorability is not existence of monitors

A formula φ is **semantically gray-box monitorable** for a system \mathcal{S} if every observation O has an extension $P \succeq O$ in \mathcal{S} , such that $P \models_{\mathcal{S}}^s \varphi$ or $P \models_{\mathcal{S}}^v \varphi$.

A **monitor** for a property φ and a system \mathcal{S} is a **computable** function $M_{\varphi, \mathcal{S}}: \mathcal{O}\{\checkmark, \times, ?\}$ that decides a **verdict** for φ given a finite u .

Observation: Monitorability of φ in \mathcal{S} **does not guarantee** the existence of a sound and complete monitor $M_{\varphi, \mathcal{S}}$.

Monitorability is not existence of monitors

A formula φ is **semantically gray-box monitorable** for a system \mathcal{S} if every observation O has an extension $P \succeq O$ in \mathcal{S} , such that $P \models_{\mathcal{S}}^s \varphi$ or $P \models_{\mathcal{S}}^v \varphi$.

A **monitor** for a property φ and a system \mathcal{S} is a **computable** function $M_{\varphi, \mathcal{S}}: \mathcal{O}\{\checkmark, \times, ?\}$ that decides a **verdict** for φ given a finite u .

Observation: Monitorability of φ in \mathcal{S} **does not guarantee** the existence of a sound and complete monitor $M_{\varphi, \mathcal{S}}$.

Example: Let T be some Turing machine.

$$\mathcal{S} = \{t \in \Sigma^\omega \mid t_i = \text{the state of } T \text{ after } i \text{ steps}\}, \quad \varphi = \diamond \text{halt.}$$

Because T is deterministic, either $u \models_{\mathcal{S}}^s \varphi$ or $u \models_{\mathcal{S}}^v \varphi$, for any u in \mathcal{S} .

Monitorability is not existence of monitors

A formula φ is **semantically gray-box monitorable** for a system \mathcal{S} if every observation O has an extension $P \succeq O$ in \mathcal{S} , such that $P \models_{\mathcal{S}}^s \varphi$ or $P \models_{\mathcal{S}}^v \varphi$.

A **monitor** for a property φ and a system \mathcal{S} is a **computable** function $M_{\varphi, \mathcal{S}}: \mathcal{O}\{\checkmark, \times, ?\}$ that decides a **verdict** for φ given a finite u .

Observation: Monitorability of φ in \mathcal{S} **does not guarantee** the existence of a sound and complete monitor $M_{\varphi, \mathcal{S}}$.

Example: Let T be some Turing machine.

$$\mathcal{S} = \{t \in \Sigma^\omega \mid t_i = \text{the state of } T \text{ after } i \text{ steps}\}, \quad \varphi = \diamond \text{halt.}$$

Because T is deterministic, either $u \models_{\mathcal{S}}^s \varphi$ or $u \models_{\mathcal{S}}^v \varphi$, for any u in \mathcal{S} .

$\Rightarrow \varphi$ is monitorable in \mathcal{S} ;

Monitorability is not existence of monitors

A formula φ is **semantically gray-box monitorable** for a system \mathcal{S} if every observation O has an extension $P \succeq O$ in \mathcal{S} , such that $P \models_{\mathcal{S}}^s \varphi$ or $P \models_{\mathcal{S}}^v \varphi$.

A **monitor** for a property φ and a system \mathcal{S} is a **computable** function $M_{\varphi, \mathcal{S}}: \mathcal{O}\{\checkmark, \times, ?\}$ that decides a **verdict** for φ given a finite u .

Observation: Monitorability of φ in \mathcal{S} **does not guarantee** the existence of a sound and complete monitor $M_{\varphi, \mathcal{S}}$.

Example: Let T be some Turing machine.

$$\mathcal{S} = \{t \in \Sigma^\omega \mid t_i = \text{the state of } T \text{ after } i \text{ steps}\}, \quad \varphi = \diamond \text{halt.}$$

Because T is deterministic, either $u \models_{\mathcal{S}}^s \varphi$ or $u \models_{\mathcal{S}}^v \varphi$, for any u in \mathcal{S} .

$\Rightarrow \varphi$ is monitorable in \mathcal{S} ;

\Rightarrow but there is no sound **and complete** monitor $M_{\varphi, \mathcal{S}}$.

Monitorability is not existence of monitors

A formula φ is **semantically gray-box monitorable** for a system \mathcal{S} if every observation O has an extension $P \succeq O$ in \mathcal{S} , such that $P \models_{\mathcal{S}}^s \varphi$ or $P \models_{\mathcal{S}}^v \varphi$.

A **monitor** for a property φ and a system \mathcal{S} is a **computable** function $M_{\varphi, \mathcal{S}}: \mathcal{O}\{\checkmark, \times, ?\}$ that decides a **verdict** for φ given a finite u .

Observation: Monitorability of φ in \mathcal{S} **does not guarantee** the existence of a sound and complete monitor $M_{\varphi, \mathcal{S}}$.

Example: Let T be some Turing machine.


$$\mathcal{S} = \{t \in \Sigma^\omega \mid t_i = \text{the state of } T \text{ after } i \text{ steps}\}, \quad \varphi = \diamond \text{halt.}$$

Because T is deterministic, either $u \models_{\mathcal{S}}^s \varphi$ or $u \models_{\mathcal{S}}^v \varphi$, for any u in \mathcal{S} .

$\Rightarrow \varphi$ is monitorable in \mathcal{S} ;

\Rightarrow *there is a **sound** monitor $M_{\varphi, \mathcal{S}}$ that only answers \checkmark or $?!$*

Case study: distributed data minimality



Non-monitorable examples

- Storage limitation (Article 5): Personal data shall be [...] adequate relevant, and limited to what is necessary in relation to the purposes for which they are processed (data minimization) [...]
- Data minimization (attempt at formalization)
`collect (data,dataid,dsid) IMPLIES EVENTUALLY use(data, dataid, dsid)`
- But MFOTL semantics requires collected data used in **EVERY** run of the system.
 - Not finitely falsifiable (liveness) and interpretation is also too strong.
 - **Example:** when booking a long-haul flight, customers provide emergency contact for an account. In majority of cases, data is collected, not used, and deleted.
- Better would be a CTL formulation (although not monitorable on a trace)
`collect (data, dataids, dsid) IMPLIES EXISTS EVENTUALLY use(data, dataid, dsid)`

34

Slide by David Basin, *Can we Verify GDPR Compliance?*, RV'19 keynote.

Case study: distributed data minimality

- Distributed data minimality (DDM)
 - privacy property (GDPR)
 - generalization of data minimality to a multi-input setting
 - $\forall\exists\exists$ -hyperproperty

$$\varphi_i = \forall\pi.\forall\pi'.\exists\tau.\exists\tau'. \neg \text{same}_i(\pi, \pi') \rightarrow \left(\begin{array}{l} \text{same}_i(\pi, \tau) \wedge \text{same}_i(\pi', \tau') \wedge \\ \text{almost}_i(\tau, \tau') \wedge \neg \text{output}(\tau, \tau') \end{array} \right)$$

Case study: distributed data minimality

- Distributed data minimality (DDM)
 - privacy property (GDPR)
 - generalization of data minimality to a multi-input setting
 - $\forall\forall\exists\exists$ -hyperproperty

$$\varphi_i = \forall\pi.\forall\pi'.\exists\tau.\exists\tau'. \neg \text{same}_i(\pi, \pi') \rightarrow \left(\begin{array}{l} \text{same}_i(\pi, \tau) \wedge \text{same}_i(\pi', \tau') \wedge \\ \text{almost}_i(\tau, \tau') \wedge \neg \text{output}(\tau, \tau') \end{array} \right)$$

- Challenges:
 - Not black-box monitorable.
 - Undecidable.
 - Defined over arbitrary domains/datatypes.

Case study: distributed data minimality

- Distributed data minimality (DDM)
 - privacy property (GDPR)
 - generalization of data minimality to a multi-input setting
 - $\forall\forall\exists\exists$ -hyperproperty

$$\varphi_i = \forall\pi.\forall\pi'.\exists\tau.\exists\tau'. \neg \text{same}_i(\pi, \pi') \rightarrow \left(\begin{array}{l} \text{same}_i(\pi, \tau) \wedge \text{same}_i(\pi', \tau') \wedge \\ \text{almost}_i(\tau, \tau') \wedge \neg \text{output}(\tau, \tau') \end{array} \right)$$

- Challenges:
 - Not black-box monitorable.
 - Undecidable.
 - Defined over arbitrary domains/datatypes.

Yet, we have a monitor. . .

Case study: distributed data minimality

- Distributed data minimality (DDM)
 - privacy property (GDPR)
 - generalization of data minimality to a multi-input setting
 - $\forall\forall\exists\exists$ -hyperproperty

$$\varphi_i = \forall\pi.\forall\pi'.\exists\tau.\exists\tau'. \neg \text{same}_i(\pi, \pi') \rightarrow \left(\begin{array}{l} \text{same}_i(\pi, \tau) \wedge \text{same}_i(\pi', \tau') \wedge \\ \text{almost}_i(\tau, \tau') \wedge \neg \text{output}(\tau, \tau') \end{array} \right)$$

- Challenges:
 - Not black-box monitorable.
 - Undecidable.
 - Defined over arbitrary domains/datatypes.

Yet, we have a monitor. . .

here's how...

Distributed data minimality

Definition (Antignac, Sands & Schneider, 2017)

A function f is **distributed data-minimal (DDM)** if, for all input positions k and all $x, y \in I_k$ such that $x \neq y$, there is some $z \in I$, such that $f(z[k \mapsto x]) \neq f(z[k \mapsto y])$.

Distributed data minimality

$$\varphi_i = \forall \pi. \forall \pi'. \exists \tau. \exists \tau'. \neg \text{same}_i(\pi, \pi') \rightarrow \left(\begin{array}{l} \text{same}_i(\pi, \tau) \wedge \text{same}_i(\pi', \tau') \wedge \\ \text{almost}_i(\tau, \tau') \wedge \neg \text{output}(\tau, \tau') \end{array} \right)$$

$$\varphi_{\text{dm}} = \bigwedge_{i=1}^n \varphi_i, \quad \Sigma_f^\# = \{(x, y) \mid f(x) = y\}, \quad \mathcal{S}_f = \mathcal{P}(\Sigma_f^\#)$$

Distributed data minimality

$$\varphi_i = \forall \pi. \forall \pi'. \exists \tau. \exists \tau'. \neg \text{same}_i(\pi, \pi') \rightarrow \left(\begin{array}{l} \text{same}_i(\pi, \tau) \wedge \text{same}_i(\pi', \tau') \wedge \\ \text{almost}_i(\tau, \tau') \wedge \neg \text{output}(\tau, \tau') \end{array} \right)$$
$$\varphi_{\text{dm}} = \bigwedge_{i=1}^n \varphi_i, \quad \Sigma_f^\# = \{(x, y) \mid f(x) = y\}, \quad \mathcal{S}_f = \mathcal{P}(\Sigma_f^\#)$$

Using the generalized framework

- Set of observable behaviors $\mathcal{O} = \Sigma_f^\#$ are valid function applications.

Distributed data minimality

$$\varphi_i = \forall \pi. \forall \pi'. \exists \tau. \exists \tau'. \neg \text{same}_i(\pi, \pi') \rightarrow \left(\begin{array}{l} \text{same}_i(\pi, \tau) \wedge \text{same}_i(\pi', \tau') \wedge \\ \text{almost}_i(\tau, \tau') \wedge \neg \text{output}(\tau, \tau') \end{array} \right)$$
$$\varphi_{\text{dm}} = \bigwedge_{i=1}^n \varphi_i, \quad \Sigma_f^\# = \{(x, y) \mid f(x) = y\}, \quad \mathcal{S}_f = \mathcal{P}(\Sigma_f^\#)$$

Using the generalized framework

- Set of observable behaviors $\mathcal{O} = \Sigma_f^\#$ are valid function applications.
- Not black-box monitorable.

Distributed data minimality

$$\varphi_i = \forall \pi. \forall \pi'. \exists \tau. \exists \tau'. \neg \text{same}_i(\pi, \pi') \rightarrow \left(\begin{array}{l} \text{same}_i(\pi, \tau) \wedge \text{same}_i(\pi', \tau') \wedge \\ \text{almost}_i(\tau, \tau') \wedge \neg \text{output}(\tau, \tau') \end{array} \right)$$
$$\varphi_{\text{dm}} = \bigwedge_{i=1}^n \varphi_i, \quad \Sigma_f^\# = \{(x, y) \mid f(x) = y\}, \quad \mathcal{S}_f = \mathcal{P}(\Sigma_f^\#)$$

Using the generalized framework

- Set of observable behaviors $\mathcal{O} = \Sigma_f^\#$ are valid function applications.
- Not black-box monitorable, but **gray-box monitorable** (thanks to \mathcal{S}).

A sound monitor for distributed data minimality

$$\varphi_i = \forall \pi. \forall \pi'. \exists \tau. \exists \tau'. \neg \text{same}_i(\pi, \pi') \rightarrow \left(\begin{array}{l} \text{same}_i(\pi, \tau) \wedge \text{same}_i(\pi', \tau') \wedge \\ \text{almost}_i(\tau, \tau') \wedge \neg \text{output}(\tau, \tau') \end{array} \right)$$

$$\varphi_{\text{dm}} = \bigwedge_{i=1}^n \varphi_i, \quad \Sigma_f^\# = \{(x, y) \mid f(x) = y\}, \quad \mathcal{S}_f = \mathcal{P}(\Sigma_f^\#)$$

A sound monitor for distributed data minimality

$$\varphi_i = \forall \pi. \forall \pi'. \exists \tau. \exists \tau'. \neg \text{same}_i(\pi, \pi') \rightarrow \left(\begin{array}{l} \text{same}_i(\pi, \tau) \wedge \text{same}_i(\pi', \tau') \wedge \\ \text{almost}_i(\tau, \tau') \wedge \neg \text{output}(\tau, \tau') \end{array} \right)$$

$$\varphi_{\text{dm}} = \bigwedge_{i=1}^n \varphi_i, \quad \Sigma_f^\# = \{(x, y) \mid f(x) = y\}, \quad \mathcal{S}_f = \mathcal{P}(\Sigma_f^\#)$$

We build a monitor

$$M_{\text{dm}}(U) = \begin{cases} ? & \text{if } f(u_{\text{in}}) \neq u_{\text{out}} \text{ for some } u \in U, \\ ? & \text{if } \bigwedge_{i=1}^n \bigwedge_{u, u' \in U} N_{f,i}(\text{proj}_i(u_{\text{in}}), \text{proj}_i(u'_{\text{in}})) \neq \mathbf{x}, \\ \mathbf{x} & \text{otherwise.} \end{cases}$$

A sound monitor for distributed data minimality

$$\varphi_i = \forall \pi. \forall \pi'. \exists \tau. \exists \tau'. \neg \text{same}_i(\pi, \pi') \rightarrow \left(\begin{array}{l} \text{same}_i(\pi, \tau) \wedge \text{same}_i(\pi', \tau') \wedge \\ \text{almost}_i(\tau, \tau') \wedge \neg \text{output}(\tau, \tau') \end{array} \right)$$
$$\varphi_{\text{dm}} = \bigwedge_{i=1}^n \varphi_i, \quad \Sigma_f^\# = \{(x, y) \mid f(x) = y\}, \quad \mathcal{S}_f = \mathcal{P}(\Sigma_f^\#)$$

We build a monitor

$$M_{\text{dm}}(U) = \begin{cases} ? & \text{if } f(u_{\text{in}}) \neq u_{\text{out}} \text{ for some } u \in U, \\ ? & \text{if } \bigwedge_{i=1}^n \bigwedge_{u, u' \in U} N_{f,i}(\text{proj}_i(u_{\text{in}}), \text{proj}_i(u'_{\text{in}})) \neq \mathbf{x}, \\ \mathbf{x} & \text{otherwise.} \end{cases}$$

using an **oracle** $N_{f,i}(x, y)$ (implemented as symbolic execution + SMT solver):

$$N_{f,i}(x, y) = \begin{cases} \checkmark \text{ or } ? & \text{if } \exists z \in I. f(z[i \mapsto x]) \neq f(z[i \mapsto y]), \\ \mathbf{x} \text{ or } ? & \text{otherwise.} \end{cases}$$

A sound monitor for distributed data minimality

We build a monitor

$$M_{\text{dm}}(U) = \begin{cases} ? & \text{if } f(u_{\text{in}}) \neq u_{\text{out}} \text{ for some } u \in U, \\ ? & \text{if } \bigwedge_{i=1}^n \bigwedge_{u, u' \in U} N_{f,i}(\text{proj}_i(u_{\text{in}}), \text{proj}_i(u'_{\text{in}})) \neq \mathbf{x}, \\ \mathbf{x} & \text{otherwise.} \end{cases}$$

using an **oracle** $N_{f,i}(x, y)$ (implemented as symbolic execution + SMT solver):

$$N_{f,i}(x, y) = \begin{cases} \checkmark \text{ or } ? & \text{if } \exists z \in I. f(z[i \mapsto x]) \neq f(z[i \mapsto y]), \\ \mathbf{x} \text{ or } ? & \text{otherwise.} \end{cases}$$

The monitor is sound **but not complete**.

Try it out!



<https://github.com/sstucki/minion/>

Thank you!

Coauthors

- César Sánchez, IMDEA SW
- Borzoo Bonakdarpour, ISU
- Gerardo Schneider, GU/Chalmers



UNIVERSITY OF
GOTHENBURG



CHALMERS
UNIVERSITY OF TECHNOLOGY

institute
imdea
software

**IOWA STATE
UNIVERSITY**

Checkout the minion monitor for data minimality



<https://github.com/sstucki/minion/>

-  Shreya Agrawal and Borzoo Bonakdarpour.
Runtime verification of k -safety hyperproperties in HyperLTL.
In Proc. of the IEEE 29th Computer Security Foundations (CSF'16), pages 239–252. IEEE CS Press, 2016.
-  Andreas Bauer, Martin Leucker, and Chrisitan Schallhart.
Runtime verification for LTL and TLTL.
ACM T. Softw. Eng. Meth., 20(4):14, 2011.
-  Andreas Bauer, Martin Leucker, and Christian Schallhart.
The good, the bad, and the ugly—but how ugly is ugly?
In Proc. of the 7th Int'l Workshop on Runtime Verification (RV'07), volume 4839 of *LNCS*, pages 126–138. Springer, 2007.
-  Borzoo Bonakdarpour, César Sánchez, and Gerardo Schneider.
Monitoring hyperproperties by combining static analysis and runtime verification.

In *Proc. of the 8th Int'l Symp. on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA'2018). Verification. Part II*, volume 11245 of *LNCS*, pages 8–27. Springer, 2018.



Yliès Falcone, Jean-Claude Fernandez, and Laurent Mounier.

What can you verify and enforce at runtime?

International Journal on Software Tools for Technology Transfer (STTT), 14(3):349–382, 2012.



Bernd Finkbeiner, Christopher Hahn, Marvin Stenger, and Leander Tentrup.
Monitoring hyperproperties.




In *Proc. of 17th Int'l Conf. on Runtime Verification (RV'17)*, volume 10548 of *LNCS*, pages 190–207. Springer, 2017.




Christopher Hahn.

Algorithms for monitoring hyperproperties.

In Bernd Finkbeiner and Leonardo Mariani, editors, *Runtime Verification*, pages 70–90, Cham, 2019. Springer International Publishing.

-  Klaus Havelund and Doron Peled.
Runtime verification: From propositional to first-order temporal logic.
In *Proc. of the 18th Int'l Conf. on Runtime Verification (RV'18)*, volume 11237 of *LNCS*, pages 90–112. Springer, 2018.
-  Amir Pnueli and Aleksandr Zaks.
PSL model checking and run-time verification via testers.
In *Proc. of the 14th Int'l Symp on Formal Methods (FM'06)*, volume 4085 of *LNCS*, pages 573–586. Springer, 2006.
-  Xian Zhang, Martin Leucker, and Wei Dong.
Runtime verification with predictive semantics.
In *Proc. of 4th NASA Int'l Symp on Formal Methods (NFM'12)*, volume 7226 of *LNCS*, pages 418–432. Springer, 2012.



Except where otherwise noted, this work is licensed under

<http://creativecommons.org/licenses/by/3.0/>