Abstract Monitors for Quantitative Specifications

RV 2022

Thomas A. Henzinger   Nicolas Mazzocchi   N. Ege Saraç

Institute of Science and Technology Austria

Online Black-Box Monitoring

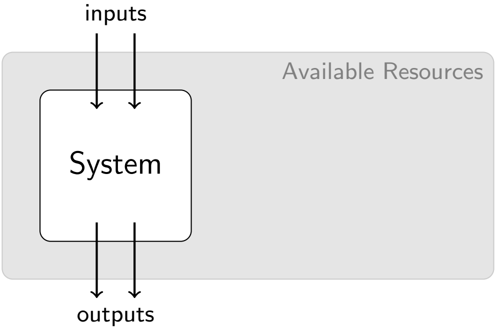

- Monitor runs in parallel and outputs a stream of verdicts
- Computation is deterministic and online (a.k.a. real-time)

Online Black-Box Monitoring

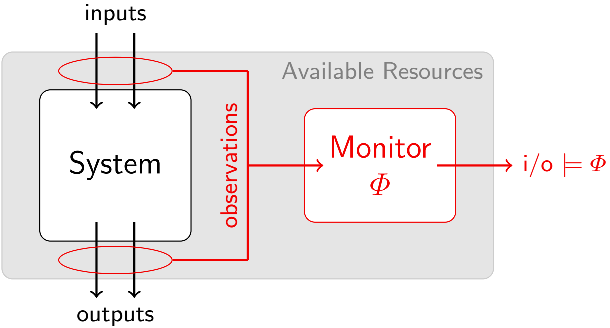

- Monitor runs in parallel and outputs a stream of verdicts
- Computation is deterministic and online (a.k.a. real-time)

Motivation

Model-Checking

- Small systems (state explosion)
- Open access (exhaustive exploration)
- Constance (verify after each update)

Model-Monitoring

- Conceptually easy (trace inclusion vs. trace membership)
- Cheap (background verification, immediate violation witness)
- System independent (black-box verification)

Goals

Quantitative verification

- Specifications map trace to a real value (instead of a Boolean)
- To capture properties on system performance (e.g. buffer length)
- Approximation (add/remove trace vs. measurable transformation)

Our framework

- Formalism that captures and abstract all monitors
- Enable to reason on approximation quality and resource availability

Example: Maximal Response $\Phi_{\max}$

$\Sigma = \{r, g, t, o\}$ $\qquad$ $f \in \Sigma^\omega$

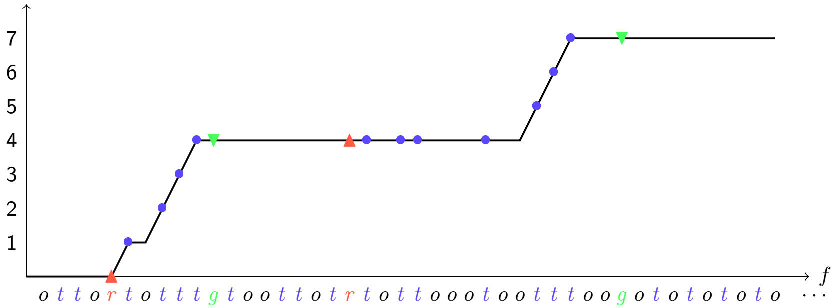

Example: Maximal Response $\Phi_{\max}$

$\Sigma = \{r, g, t, o\}$ $\qquad$ $f \in \Sigma^\omega$

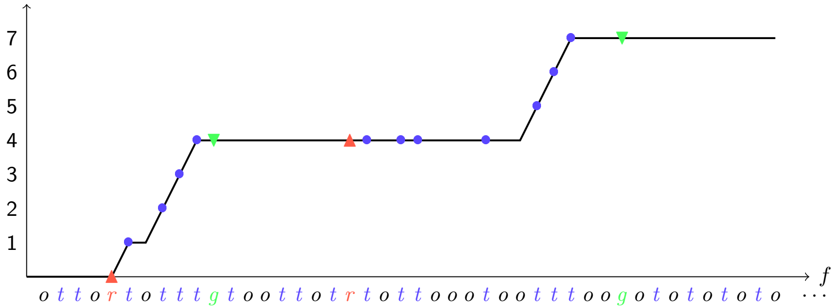

Limit behavior

- $\Phi_{\max}(f) = \infty$ if f admits some r is never followed by g, otherwise
- $\Phi_{\max}(f) = \max\{|u|_t : f \in \{\Sigma^* r u g \Sigma^\omega\}, u \in \{o, t, r\}^*\}$

Example: Discounted Response Φ_{disc}

$\Sigma = \{r, g, t, o\}$ $\qquad f \in \Sigma^\omega$

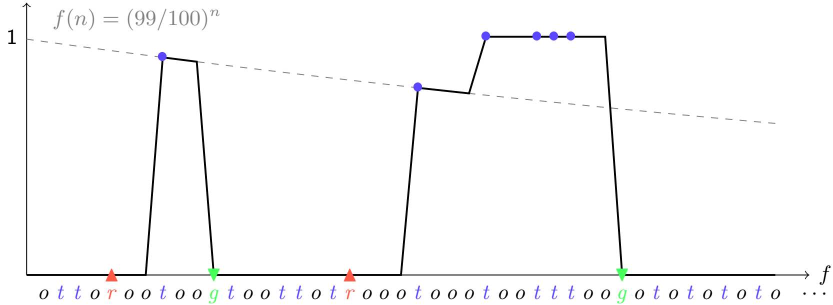

Example: Discounted Response Φ_{disc}

$\Sigma = \{r, g, t, o\}$ $\quad\quad$ $f \in \Sigma^\omega$

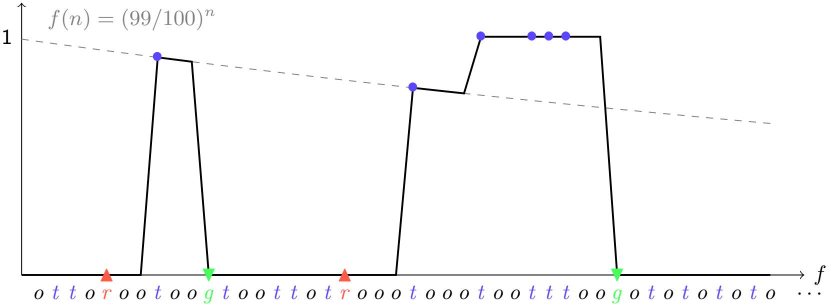

Limit behavior

- $\Phi_{\text{disc}}(f) = 1$ if f admits some r is followed by 2 t but no g, otherwise
- $\Phi_{\text{disc}}(f) = 0$

Specification

Definition

Syntax $\Phi = (\pi, \ell)$ where $\pi \colon \Sigma^* \to \mathbb{R}$ and $\ell \in \{\liminf, \limsup\}$

Semantics $[\![\Phi]\!] \colon \Sigma^* \cup \Sigma^\omega \to \mathbb{R}$ such that

finite words $[\![\Phi]\!](s) = \pi(s)$ for all $s \in \Sigma^*$

infinite words $[\![\Phi]\!](f) = \ell(\pi(f))$ for all $f \in \Sigma^\omega$

- where $\pi(f) = (\pi(s_i))_{i \in \mathbb{N}}$ and $s_i \prec f$ with $|s_i| = i$

Specification

Definition

Syntax $\Phi = (\pi, \ell)$ where $\pi \colon \Sigma^* \to \mathbb{R}$ and $\ell \in \{\liminf, \limsup\}$

Semantics $[\![\Phi]\!] \colon \Sigma^* \cup \Sigma^\omega \to \mathbb{R}$ such that

finite words $[\![\Phi]\!](s) = \pi(s)$ for all $s \in \Sigma^*$

infinite words $[\![\Phi]\!](f) = \ell(\pi(f))$ for all $f \in \Sigma^\omega$

- where $\pi(f) = (\pi(s_i))_{i \in \mathbb{N}}$ and $s_i \prec f$ with $|s_i| = i$

$$\text{resp}(s) = \begin{cases} 0 & \text{if each } r \text{ in } s \text{ has a succeeding } g \\ |s|_t - |r|_t & \text{otherwise, where } r \prec s \text{ is longest with } \text{resp}(r) = 0 \end{cases}$$

Maximal Response

- $\Phi_{\max} = (\pi_{\max}, \limsup)$ where $\pi_{\max}(s) = \max_{r \preceq s} \text{resp}(r)$

Specification

Definition

Syntax $\Phi = (\pi, \ell)$ where $\pi \colon \Sigma^* \to \mathbb{R}$ and $\ell \in \{\liminf, \limsup\}$

Semantics $[\![\Phi]\!] \colon \Sigma^* \cup \Sigma^\omega \to \mathbb{R}$ such that

finite words $[\![\Phi]\!](s) = \pi(s)$ for all $s \in \Sigma^*$

infinite words $[\![\Phi]\!](f) = \ell(\pi(f))$ for all $f \in \Sigma^\omega$

- where $\pi(f) = (\pi(s_i))_{i \in \mathbb{N}}$ and $s_i \prec f$ with $|s_i| = i$

$$\mathsf{resp}(s) = \begin{cases} 0 & \text{if each } r \text{ in } s \text{ has a succeeding } g \\ |s|_t - |r|_t & \text{otherwise, where } r \prec s \text{ is longest with } \mathsf{resp}(r) = 0 \end{cases}$$

Discounted Response

- $\Phi_{\mathsf{disc}} = (\pi_{\mathsf{disc}}, \liminf)$ where $\pi_{\mathsf{disc}}(s) = \begin{cases} 0 & \text{if } \mathsf{resp}(s) = 0 \\ (99/100)^{|s|} & \text{if } \mathsf{resp}(s) = 1 \\ 1 & \text{if } \mathsf{resp}(s) > 1 \end{cases}$

Monitors

Definition

Syntax $\mathcal{M} = (\sim, \gamma)$ where

$\sim \subseteq \Sigma^* \times \Sigma^*$ is a right-monotonic equivalence relation

$\gamma \colon (\Sigma^* / \sim) \to \mathbb{R}$ is a function

Semantics $\mathcal{M}$ is a $(\delta_{\text{prompt}}, \delta_{\text{limit}})$-monitor for $\Phi = (\pi, \ell)$ iff

prompt-error: $|\pi(s) - \gamma([s])| \leq \delta_{\text{prompt}}$ for all $s \in \Sigma^*$

limit-error: $|\ell(\pi(f)) - \ell(\gamma([f]))| \leq \delta_{\text{limit}}$ for all $f \in \Sigma^\omega$

- where $\gamma([f]) = (\gamma([s_i]))_{i \in \mathbb{N}}$ and $s_i \prec f$ with $|s_i| = i$

Monitors

Definition

Syntax $\mathcal{M} = (\sim, \gamma)$ where

$\sim \subseteq \Sigma^* \times \Sigma^*$ is a right-monotonic equivalence relation

$\gamma \colon (\Sigma^* / \sim) \to \mathbb{R}$ is a function

Semantics $\mathcal{M}$ is a $(\delta_{\text{prompt}}, \delta_{\text{limit}})$-monitor for $\Phi = (\pi, \ell)$ iff

prompt-error: $|\pi(s) - \gamma([s])| \leq \delta_{\text{prompt}}$ for all $s \in \Sigma^*$

limit-error: $|\ell(\pi(f)) - \ell(\gamma([f]))| \leq \delta_{\text{limit}}$ for all $f \in \Sigma^\omega$

- where $\gamma([f]) = (\gamma([s_i]))_{i \in \mathbb{N}}$ and $s_i \prec f$ with $|s_i| = i$

Exact-value monitor

- $\mathcal{M}_\Phi = (\sim_\Phi^*, s \mapsto \pi(s))$ for a given $\Phi = (\pi, \ell)$ where
  $$\forall s_1, s_2 \in \Sigma^* : \big(s_1 \sim_\Phi^* s_2 \iff \forall r \in \Sigma^* : \pi(s_1 r) = \pi(s_2 r)\big)$$

Example: Approximate Maximal Response $\mathcal{M}_{\max}$

$\Sigma = \{r, g, t, o\}$ $\qquad$ $f \in \Sigma^\omega$

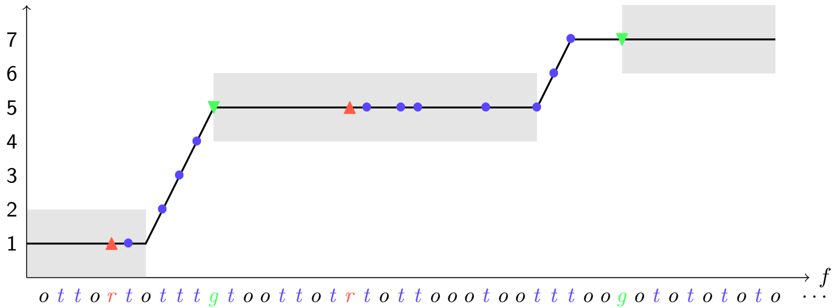

Example: Approximate Maximal Response $\mathcal{M}_{\max}$

$\Sigma = \{r, g, t, o\}$ $f \in \Sigma^\omega$

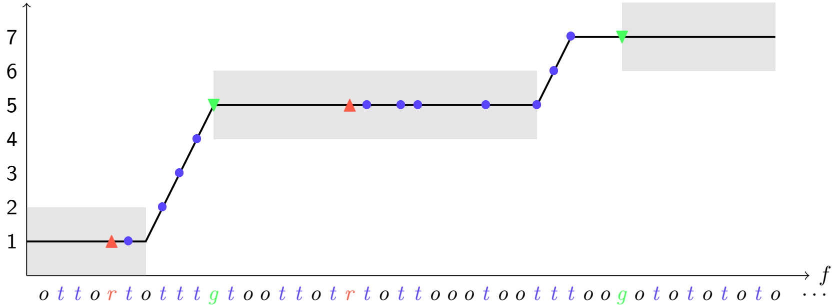

Limit behavior

- $\mathcal{M}_{\max}(f) = \infty$ if f admits some r never followed by g, otherwise
- $\mathcal{M}_{\max}(f) = \Phi_{\max}(f) + \big(\Phi_{\max}(f) + 1 \mod 2\big)$

Example: Approximate Discounted Response $\mathcal{M}_{\text{disc}}$

$\Sigma = \{r, g, t, o\}$ $\qquad$ $f \in \Sigma^\omega$

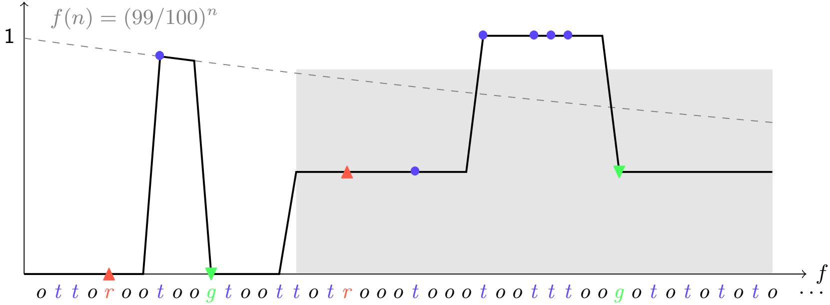

Example: Approximate Discounted Response $\mathcal{M}_{\text{disc}}$

$\Sigma = \{r, g, t, o\}$ $\qquad f \in \Sigma^\omega$

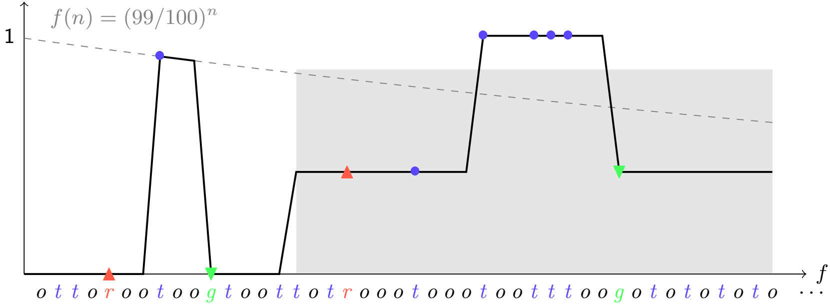

Limit behavior

- $\mathcal{M}_{\text{disc}}(f) = 1$ if f admits some r is followed by 2 t but no g, otherwise
- $\mathcal{M}_{\text{disc}}(f) = \frac{(99/100)^{15}}{2}$

Resource use

Definition

Let $\mathcal{M} = (\sim, \gamma)$ be a monitor.

- $\mathbf{r}_n(\mathcal{M}) = |\Sigma^{\leq n}/\sim| - |\Sigma^{<n}/\sim|$
- $\mathbf{R}_n(\mathcal{M}) = \sum_{i=0}^{n} \mathbf{r}_i(\mathcal{M}) = |\Sigma^{\leq n}/\sim|$

Optimality

- $\mathcal{M}$ is resource-optimal when it uses at most as many resources as any other monitor $\mathcal{M}'$ with the same error thresholds

Definition

Given a specification Φ and a $(\delta_{\text{prompt}}, \delta_{\text{limit}})$-monitor $\mathcal{M}$ for Φ, we say that $\mathcal{M}$ is *resource-optimal* for Φ when for every $(\delta_{\text{prompt}}, \delta_{\text{limit}})$-monitor $\mathcal{M}'$ for Φ we have $\mathbf{r}_n(\mathcal{M}) \leq \mathbf{r}_n(\mathcal{M}')$ for all n.

Approximate Monitoring

Prompt-error Monitoring

- Bounds the prompt-error (i.e., $\delta_{\text{prompt}} \neq \infty$)
- Prompt-error guarantees implies limit-error guarantees
- Provides a constant approximation precision

Limit-error Monitoring

- No limit-error (i.e., $\delta_{\text{limit}} = 0$)
- Targets a perfect precision on the limit
- Supports speculative monitor (i.e., non-monotonic verdict)

Prompt-error Monitoring is NOT canonical

Theorem

For all $x > 0$ and $y \leq x$ there exists a specification Φ that admits multiple resource-optimal (x, y)-monitors.

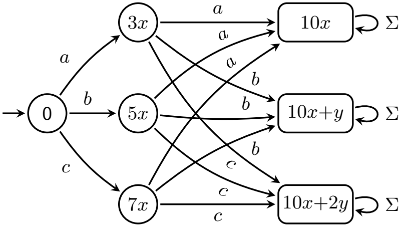

Prompt-error Monitoring is NOT canonical

Theorem

For all $x > 0$ and $y \leq x$ there exists a specification Φ that admits multiple resource-optimal (x, y)-monitors.

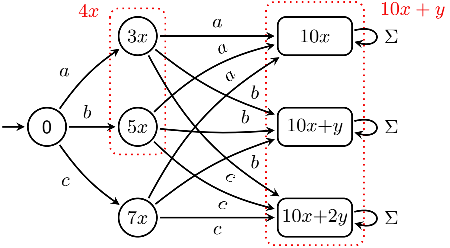

Prompt-error Monitoring is NOT canonical

Theorem

For all $x > 0$ and $y \leq x$ there exists a specification Φ that admits multiple resource-optimal (x, y)-monitors.

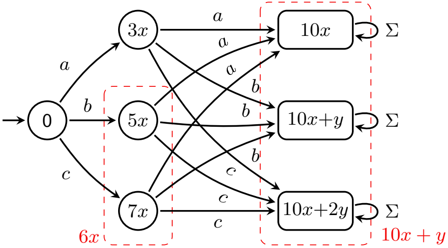

Prompt-error Monitoring is NOT canonical

Theorem

For all $x > 0$ and $y \leq x$ there exists a specification Φ that admits multiple resource-optimal (x, y)-monitors.

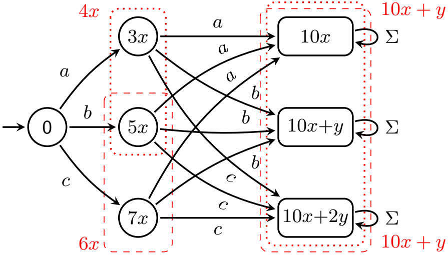

- The exact-value monitor is the unique resource-optimal $(0, 0)$-monitor.

Prompt-error Monitoring is NOT hierarchical

Theorem

There exists an optimal $(1, 0)$-monitor $\mathcal{M} = (\sim, \gamma)$ for some specification Φ such that for every other $(1, 0)$-monitor $\mathcal{M}' = (\sim', \gamma')$ we have that $\sim_\Phi \subseteq \sim'$ implies $\mathcal{M}'$ non-optimal.

$$\begin{aligned} \varepsilon &\mapsto 0 \\ c &\mapsto 3 \\ a, ca &\mapsto 6 \\ b, cb &\mapsto 9 \\ cab &\mapsto 12 \\ ab, ba &\mapsto 14 \\ cba &\mapsto 16 \\ * &\mapsto 19 \end{aligned}$$

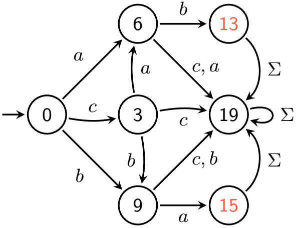

Prompt-error Monitoring is NOT greedy

Theorem

There exists a specification Φ admitting a $(1,1)$-monitor $\mathcal{M} = (\sim, \gamma)$ such that for all equivalence relations $\approx$ over Σ^* and $n \in \mathbb{N}$ we have that $|\Sigma^{\leq n}/{\sim}|$ is strictly greater than

$$\min \left\{ |\Sigma^{\leq n}/{\approx}| \;\middle|\; \forall s_1, s_2 \in \Sigma^{\leq n} : s_1 \approx s_2 \Rightarrow \bigwedge \begin{array}{l} \forall r \in \Sigma^* : s_1 r \approx s_2 r \\ |\Phi(s_1) - \Phi(s_2)| \leq 1 \end{array} \right\}$$

$$\begin{array}{lcc} \varepsilon & \mapsto & 8 \times 0 = 0 \\ a & \mapsto & 8 \times 1 - 2 = 6 \\ b & \mapsto & 8 \times 1 = 8 \end{array}$$

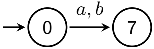

Prompt-error Monitoring is NOT greedy

Theorem

There exists a specification Φ admitting a $(1,1)$-monitor $\mathcal{M} = (\sim, \gamma)$ such that for all equivalence relations $\approx$ over Σ^* and $n \in \mathbb{N}$ we have that $|\Sigma^{\leq n}/{\sim}|$ is strictly greater than

$$\min \left\{ |\Sigma^{\leq n}/{\approx}| \;\middle|\; \forall s_1, s_2 \in \Sigma^{\leq n} : s_1 \approx s_2 \Rightarrow \bigwedge \begin{array}{l} \forall r \in \Sigma^* : s_1 r \approx s_2 r \\ |\Phi(s_1) - \Phi(s_2)| \leq 1 \end{array} \right\}$$

$$\begin{aligned}
\varepsilon &\mapsto 8 \times 0 = 0 \\
a &\mapsto 8 \times 1 - 2 = 6 \\
b &\mapsto 8 \times 1 = 8 \\
aa &\mapsto 8 \times 2 - 2 = 14 \\
ab &\mapsto 8 \times 2 - 16 \times 1 + 10 = 10 \\
ba &\mapsto 8 \times 2 - 16 \times 1 + 4 = 4 \\
bb &\mapsto 8 \times 2 = 16
\end{aligned}$$

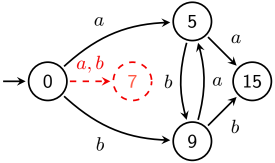

Prompt-error Monitoring is NOT greedy

Theorem

There exists a specification Φ admitting a $(1,1)$-monitor $\mathcal{M} = (\sim, \gamma)$ such that for all equivalence relations $\approx$ over Σ^* and $n \in \mathbb{N}$ we have that $|\Sigma^{\leq n}/\sim|$ is strictly greater than

$$\min \left\{ |\Sigma^{\leq n}/\approx| \;\middle|\; \forall s_1, s_2 \in \Sigma^{\leq n} : s_1 \approx s_2 \Rightarrow \bigwedge \begin{array}{l} \forall r \in \Sigma^* : s_1 r \approx s_2 r \\ |\Phi(s_1) - \Phi(s_2)| \leq 1 \end{array} \right\}$$

$$\begin{aligned}
aaa &\mapsto 8 \times 3 - 2 = 22 \\
aab &\mapsto 8 \times 3 - 16 \times 1 + 10 = 18 \\
aba &\mapsto 8 \times 3 - 16 \times 1 - 4 = 4 \\
abb &\mapsto 8 \times 3 - 16 \times 1 + 10 = 18 \\
baa &\mapsto 8 \times 3 - 16 \times 1 + 4 = 12 \\
bab &\mapsto 8 \times 3 - 16 \times 1 + 2 = 10 \\
bba &\mapsto 8 \times 3 - 16 \times 1 + 4 = 12 \\
bbb &\mapsto 8 \times 3 = 24
\end{aligned}$$

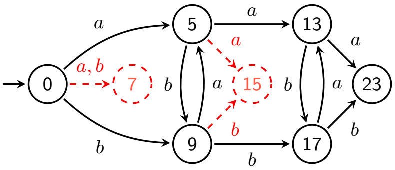

Prompt-error Monitoring is NOT greedy

Theorem

There exists a specification Φ admitting a $(1,1)$-monitor $\mathcal{M} = (\sim, \gamma)$ such that for all equivalence relations $\approx$ over Σ^* and $n \in \mathbb{N}$ we have that $|\Sigma^{\leq n}/{\sim}|$ is strictly greater than

$$\min \left\{ |\Sigma^{\leq n}/{\approx}| \;\middle|\; \forall s_1, s_2 \in \Sigma^{\leq n} : s_1 \approx s_2 \Rightarrow \bigwedge \begin{array}{l} \forall r \in \Sigma^* : s_1 r \approx s_2 r \\ |\Phi(s_1) - \Phi(s_2)| \leq 1 \end{array} \right\}$$

$$\pi(s) = \begin{cases} 8|s| & \text{if } s \in b^* \\ 8|s| - 16k + 4 & \text{if } s \in (b^+a^+)^k \text{ for some } k \geq 1 \\ 8|s| - 16k + 2 & \text{if } s \in (b^+a^+)^k b^+ \text{ for some } k \geq 1 \\ 8|s| - 2 & \text{if } s \in a^+ \\ 8|s| - 16k + 10 & \text{if } s \in (a^+b^+)^k \text{ for some } k \geq 1 \\ 8|s| - 16k - 4 & \text{if } s \in (a^+b^+)^k a^+ \text{ for some } k \geq 1 \end{cases}$$

Prompt-error Monitoring is NOT greedy

Theorem

There exists a specification Φ admitting a $(1,1)$-monitor $\mathcal{M} = (\sim, \gamma)$ such that for all equivalence relations $\approx$ over Σ^* and $n \in \mathbb{N}$ we have that $|\Sigma^{\leq n}/\!\sim|$ is strictly greater than

$$\min\left\{ |\Sigma^{\leq n}/\!\approx| \;\middle|\; \forall s_1, s_2 \in \Sigma^{\leq n} : s_1 \approx s_2 \Rightarrow \bigwedge \begin{array}{l} \forall r \in \Sigma^* : s_1 r \approx s_2 r \\ |\Phi(s_1) - \Phi(s_2)| \leq 1 \end{array} \right\}$$

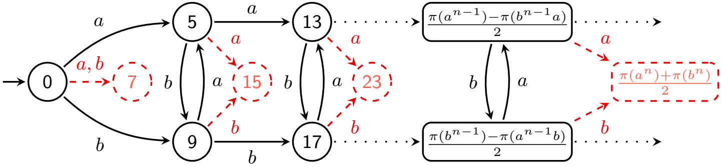

- At each step n, attempting to minimize $\mathbf{R}_n$ results in taking a^n and b^n as equivalent, leading to violate any congruence for step $n+1$.

Prompt Monitoring saves resources

Theorem

For all $\delta \in \mathbb{N}$, there exists a (δ, δ)-monitor $\mathcal{M}_\delta$ for the maximal response specification $\Phi_{\max}$. Furthermore, for all $\delta_i > \delta_j$, $\mathbf{r}_n(\mathcal{M}_{\delta_i}) \leq \mathbf{r}_n(\mathcal{M}_{\delta_j})$ for all n and $\mathbf{r}_k(\mathcal{M}_{\delta_i}) < \mathbf{r}_k(\mathcal{M}_{\delta_j})$ for some k.

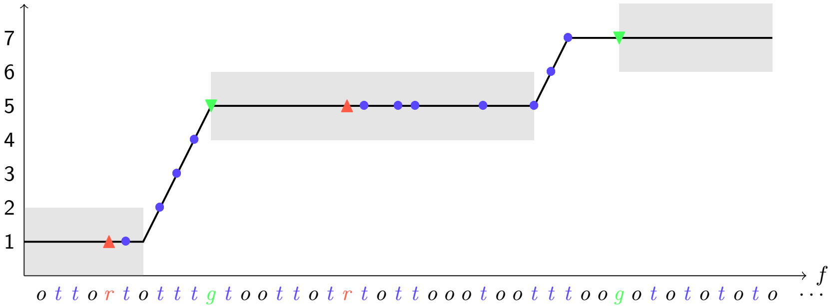

Prompt Monitoring saves resources

Theorem

For all $\delta \in \{x \in \mathbb{R} \mid 0 < x \leq 1\}$, there exists a (δ, δ)-monitor $\mathcal{M}_\delta$ for the discounted response specification Φ_{disc}. Furthermore, for all $\delta_i > \delta_j$, $\mathbf{r}_n(\mathcal{M}_{\delta_i}) \leq \mathbf{r}_n(\mathcal{M}_{\delta_j})$ for all n and and $\mathbf{r}_k(\mathcal{M}_{\delta_i}) < \mathbf{r}_k(\mathcal{M}_{\delta_j})$ for some k.

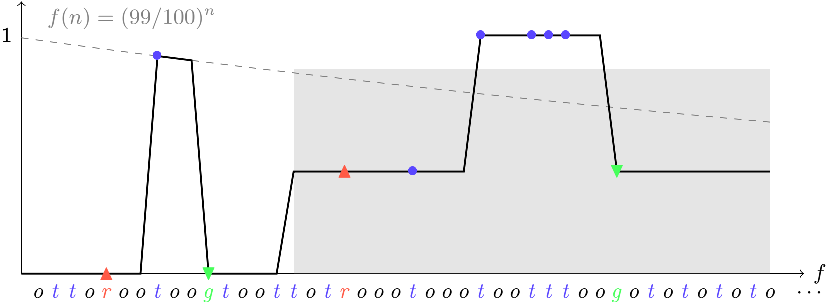

Limit Monitoring

Exact-value vs. Exact-limit

- $\mathcal{M}_\Phi = (\sim_\Phi^*, s \mapsto \pi(s))$ for a given $\Phi = (\pi, \ell)$ where
  $$\forall s_1, s_2 \in \Sigma^* : \big(s_1 \sim_\Phi^* s_2 \iff \forall r \in \Sigma^* : \pi(s_1 r) = \pi(s_2 r)\big)$$
- $\mathcal{M}_\Phi^\omega = (\sim_\Phi^\omega, s \mapsto \pi(s))$ for a given $\Phi = (\pi, \ell)$ where
  $$\forall s_1, s_2 \in \Sigma^* : \big(s_1 \sim_\Phi^\omega s_2 \iff \forall f \in \Sigma^\omega : \ell(\pi(s_1 f)) = \ell(\pi(s_2 f))\big)$$

Limit Monitoring

Exact-value vs. Exact-limit

- $\mathcal{M}_\Phi = (\sim_\Phi^*, s \mapsto \pi(s))$ for a given $\Phi = (\pi, \ell)$ where
  $$\forall s_1, s_2 \in \Sigma^* : \big(s_1 \sim_\Phi^* s_2 \iff \forall r \in \Sigma^* : \pi(s_1 r) = \pi(s_2 r)\big)$$
- $\mathcal{M}_\Phi^\omega = (\sim_\Phi^\omega, s \mapsto \pi(s))$ for a given $\Phi = (\pi, \ell)$ where
  $$\forall s_1, s_2 \in \Sigma^* : \big(s_1 \sim_\Phi^\omega s_2 \iff \forall f \in \Sigma^\omega : \ell(\pi(s_1 f)) = \ell(\pi(s_2 f))\big)$$

Theorem

Let Φ be a specification. If $\sim_\Phi^* = \sim_\Phi^\omega$ then its exact-value monitor $\mathcal{M}_\Phi$ is a resource-optimal $(\delta, 0)$-monitor for any $\delta \geq 0$.

Limit Monitoring

Exact-value vs. Exact-limit

- $\mathcal{M}_\Phi = (\sim_\Phi^*, s \mapsto \pi(s))$ for a given $\Phi = (\pi, \ell)$ where
  $$\forall s_1, s_2 \in \Sigma^* : \big(s_1 \sim_\Phi^* s_2 \iff \forall r \in \Sigma^* : \pi(s_1 r) = \pi(s_2 r)\big)$$
- $\mathcal{M}_\Phi^\omega = (\sim_\Phi^\omega, s \mapsto \pi(s))$ for a given $\Phi = (\pi, \ell)$ where
  $$\forall s_1, s_2 \in \Sigma^* : \big(s_1 \sim_\Phi^\omega s_2 \iff \forall f \in \Sigma^\omega : \ell(\pi(s_1 f)) = \ell(\pi(s_2 f))\big)$$

Theorem

Let Φ be a specification. If $\sim_\Phi^* = \sim_\Phi^\omega$ then its exact-value monitor $\mathcal{M}_\Phi$ is a resource-optimal $(\delta, 0)$-monitor for any $\delta \geq 0$.

Example: Maximal Response

- $s_1 \not\sim_{\Phi_{Max}}^* s_2 \implies \Phi_{Max}(s_1 r) \neq \Phi_{Max}(s_2 r)$ for some $r \in \Sigma^*$
- if $\Phi_{Max}(s_1 r) \neq \Phi_{Max}(s_2 r)$ then $\Phi_{Max}(s_1 r (g)^\omega) \neq \Phi_{Max}(s_2 r (g)^\omega)$
- $\Phi_{Max}(s_1 r (g)^\omega) \neq \Phi_{Max}(s_2 r (g)^\omega) \implies s_1 \not\sim_{\Phi_{Max}}^\omega s_2$

Conclusion

Our framework

- Formalism that captures and abstract all monitors
- Enable to reason on approximation quality and resource availability

Future work

- Dynamic resource allocation
- Conditions enabling finite-state approximations
- Transformations allowing to adjust the resource/precision trade-off

Conclusion

Our framework

- Formalism that captures and abstract all monitors
- Enable to reason on approximation quality and resource availability

Future work

- Dynamic resource allocation
- Conditions enabling finite-state approximations
- Transformations allowing to adjust the resource/precision trade-off

Thank you