Computing Quantiles in Markov Reward Models

Michael Ummels

German Aerospace Center

michael.ummels@dlr.de

(Joint Work with Christel Baier, TU Dresden)

FOSSACS 2013

Markov Reward Models

Model: Markov decision processes with nonnegative rewards on states.

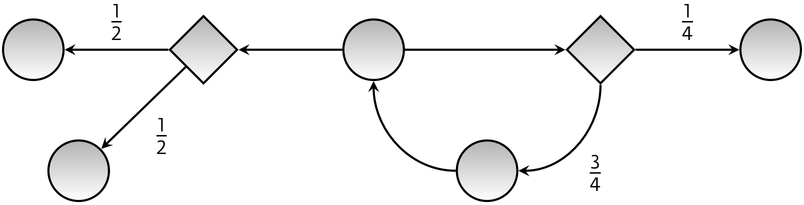

Markov Reward Models

Model: Markov decision processes with nonnegative rewards on states.

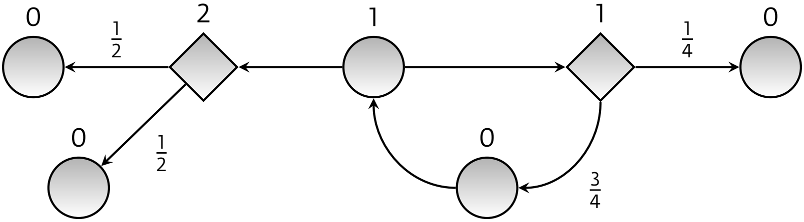

Markov Reward Models

Model: Markov decision processes with nonnegative rewards on states.

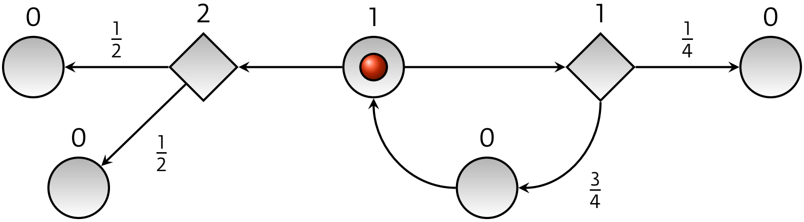

Accumulated reward: 0

Markov Reward Models

Model: Markov decision processes with nonnegative rewards on states.

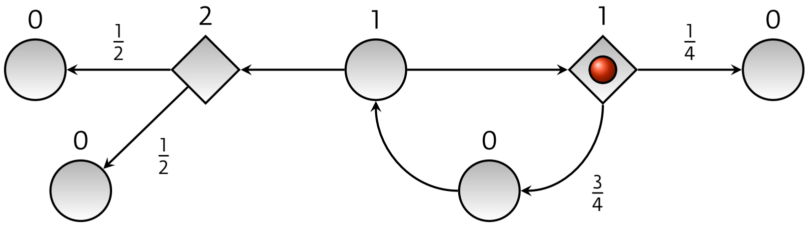

Accumulated reward: $0 + 1$

Markov Reward Models

Model: Markov decision processes with nonnegative rewards on states.

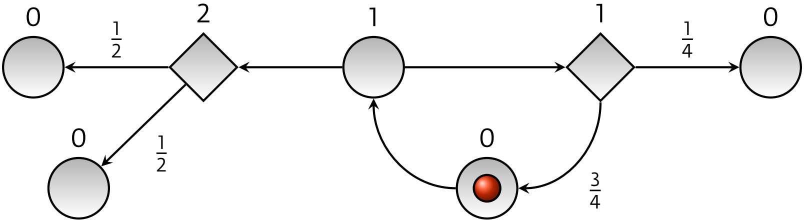

Accumulated reward: $0 + 1 + 1$

Markov Reward Models

Model: Markov decision processes with nonnegative rewards on states.

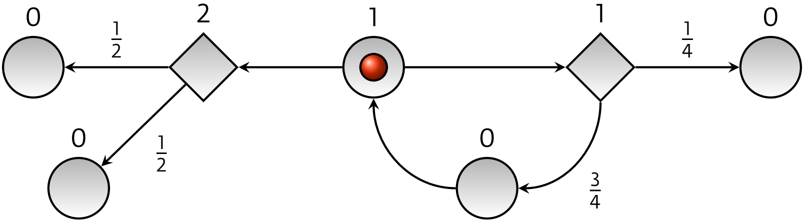

Accumulated reward: $0 + 1 + 1 + 0$

Markov Reward Models

Model: Markov decision processes with nonnegative rewards on states.

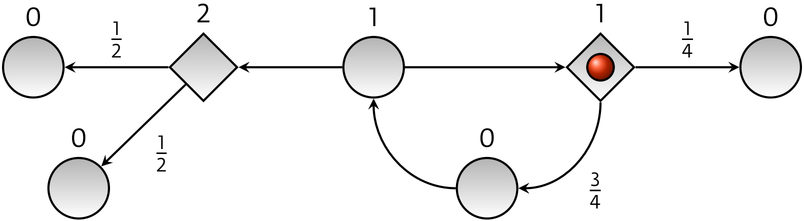

Accumulated reward: $0 + 1 + 1 + 0 + 1$

Markov Reward Models

Model: Markov decision processes with nonnegative rewards on states.

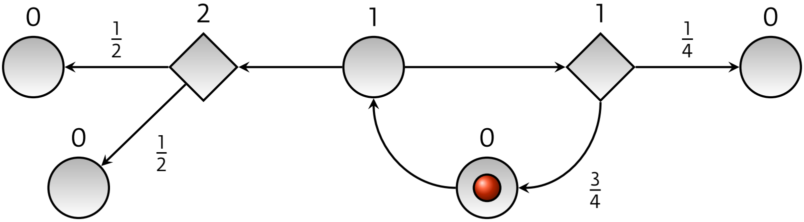

Accumulated reward: $0 + 1 + 1 + 0 + 1 + 1$

Markov Reward Models

Model: Markov decision processes with nonnegative rewards on states.

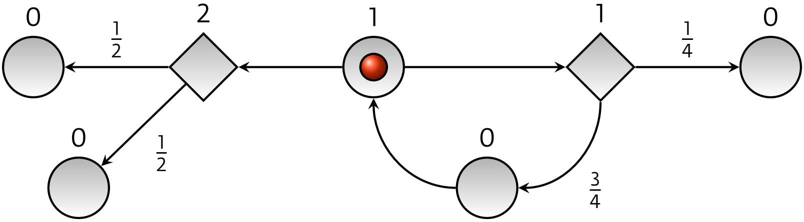

Accumulated reward: $0 + 1 + 1 + 0 + 1 + 1 + 0$

Markov Reward Models

Model: Markov decision processes with nonnegative rewards on states.

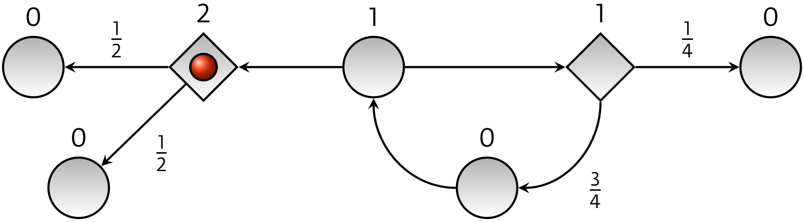

Accumulated reward: $0 + 1 + 1 + 0 + 1 + 1 + 0 + 1$

Markov Reward Models

Model: Markov decision processes with nonnegative rewards on states.

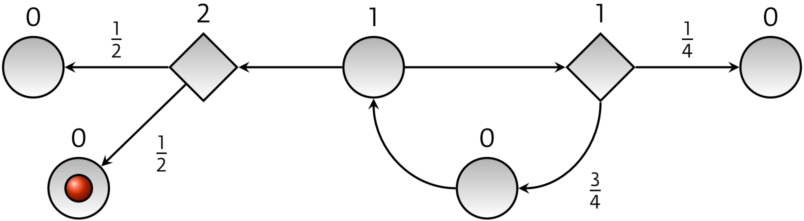

Accumulated reward: $0 + 1 + 1 + 0 + 1 + 1 + 0 + 1 + 2$

Markov Reward Models

Model: Markov decision processes with nonnegative rewards on states.

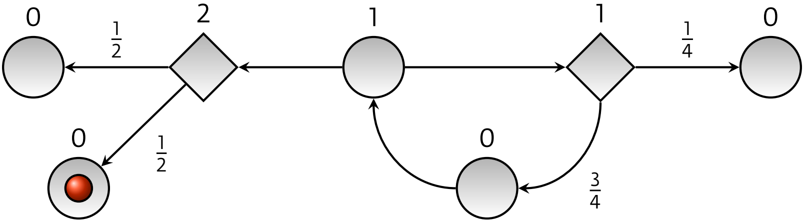

Accumulated reward: $0 + 1 + 1 + 0 + 1 + 1 + 0 + 1 + 2 = 7$

Markov Reward Models

Model: Markov decision processes with nonnegative rewards on states.

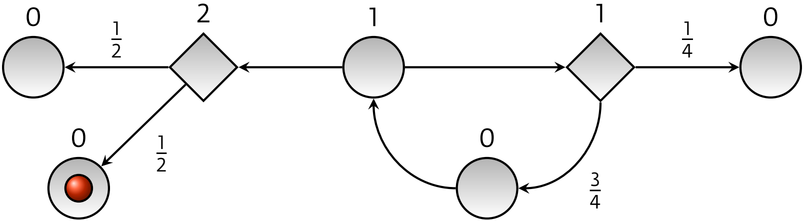

Accumulated reward: $0+1+1+0+1+1+0+1+2=7$

Note: *Scheduler* resolves nondeterminism.

PRCTL

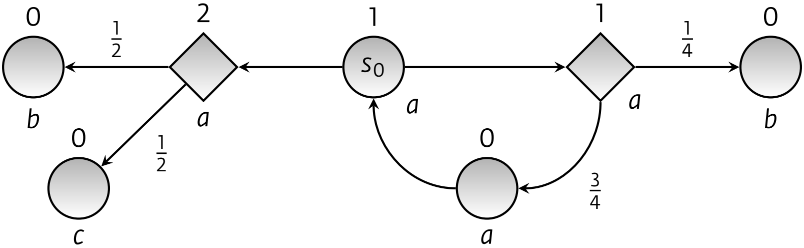

Example properties in PRCTL (Andova et al.):

PRCTL

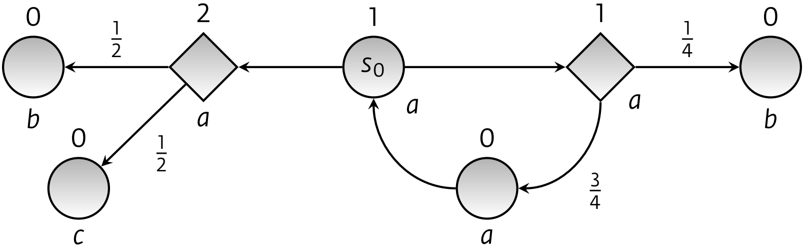

Example properties in PRCTL (Andova et al.):

- $s_0 \models P_{>0.2}(a \, U_{\leq 3} \, b)$

PRCTL

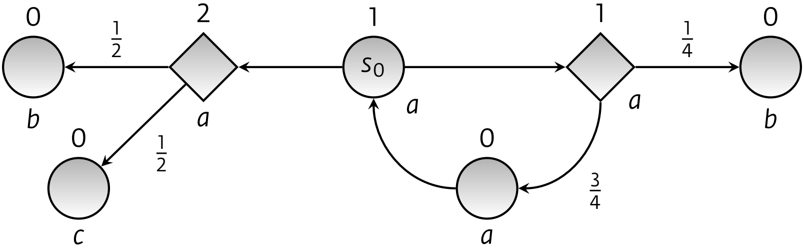

Example properties in PRCTL (Andova et al.):

- $s_0 \models P_{>0.2}(a \, U_{\leq 3} \, b)$
- $s_0 \models P_{=0}(a \, U_{\leq 1} \, b)$

PRCTL

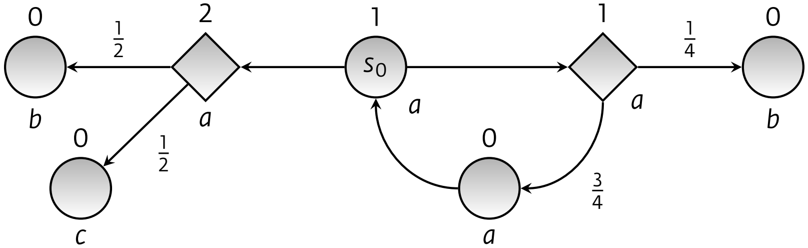

Example properties in PRCTL (Andova et al.):

- $s_0 \models \mathrm{P}_{>0.2}(a \,\mathrm{U}_{\leq 3}\, b)$
- $s_0 \not\models \mathrm{P}_{\leq 0.2}(a \,\mathrm{U}_{\leq 2}\, b)$
- $s_0 \models \mathrm{P}_{=0}(a \,\mathrm{U}_{\leq 1}\, b)$

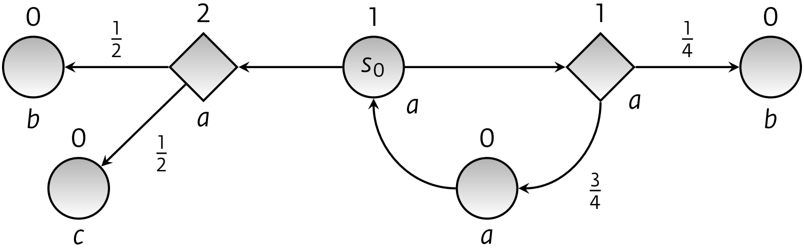

Example properties in PRCTL (Andova et al.):

- $s_0 \models \mathrm{P}_{>0.2}(a \,\mathrm{U}_{\leq 3}\, b)$
- $s_0 \models \mathrm{P}_{=0}(a \,\mathrm{U}_{\leq 1}\, b)$
- $s_0 \not\models \mathrm{P}_{\leq 0.2}(a \,\mathrm{U}_{\leq 2}\, b)$
- $s_0 \not\models \mathrm{P}_{>0}(a \,\mathrm{U}_{\leq 2}\, c)$

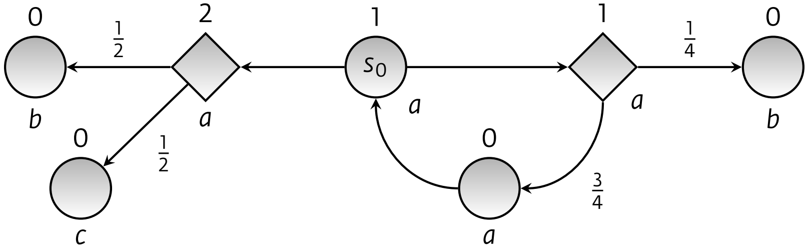

Example properties in PRCTL (Andova et al.):

- $s_0 \models \forall P_{>0.2}(a \, U_{\leq 3} \, b)$
- $s_0 \models \forall P_{=0}(a \, U_{\leq 1} \, b)$
- $s_0 \models \exists P_{>0.2}(a \, U_{\leq 2} \, b)$
- $s_0 \models \exists P_{=0}(a \, U_{\leq 2} \, b)$

Motivation

Example: Randomised Mutual exclusion.

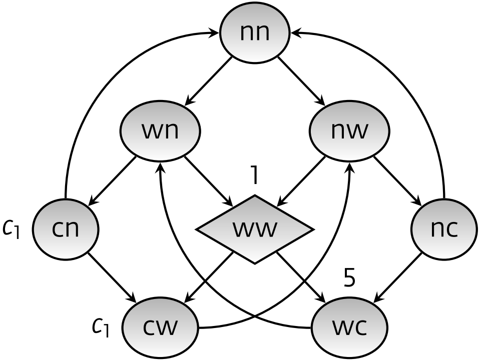

Motivation

Example: Randomised Mutual exclusion.

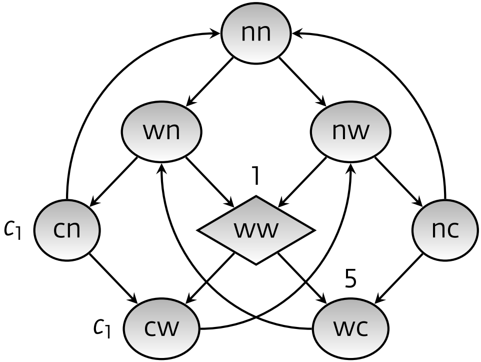

Question: How many steps may process 1 wait until with 90% chance in critical section?

Motivation

Example: Randomised Mutual exclusion.

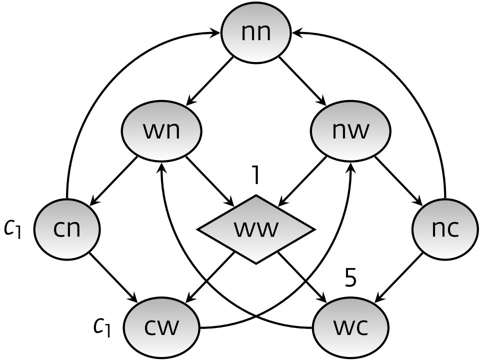

Question: How many steps may process 1 wait until with 90% chance in critical section?

Compute: least r such that wn $\models \forall \mathrm{P}_{\geq 0.9}(\mathit{true}\ \mathrm{U}_{\leq r}\ c_1)$.

More Motivation

Example: Resource Consumption.

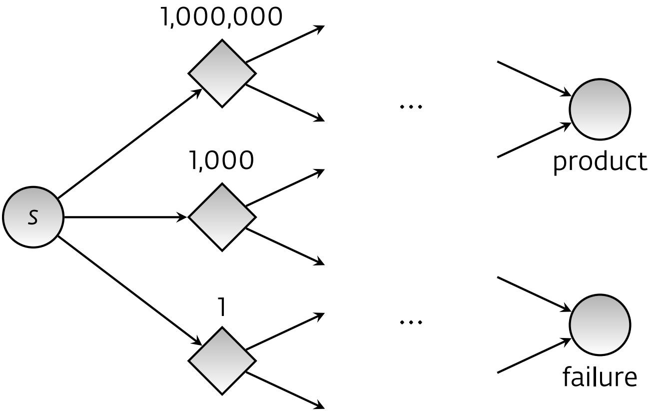

More Motivation

Example: Resource Consumption.

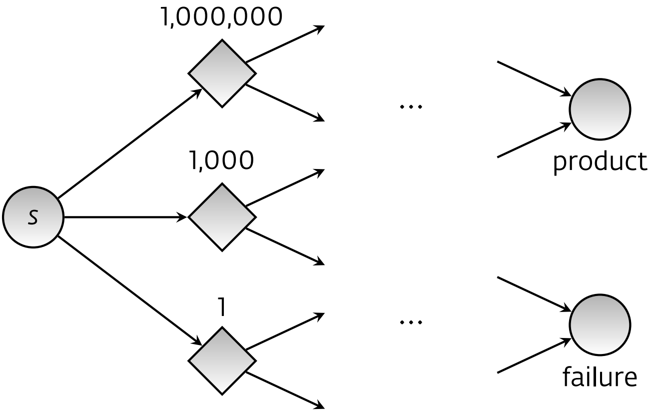

Question: How much to invest to successfully produce with 99%?

More Motivation

Example: Resource Consumption.

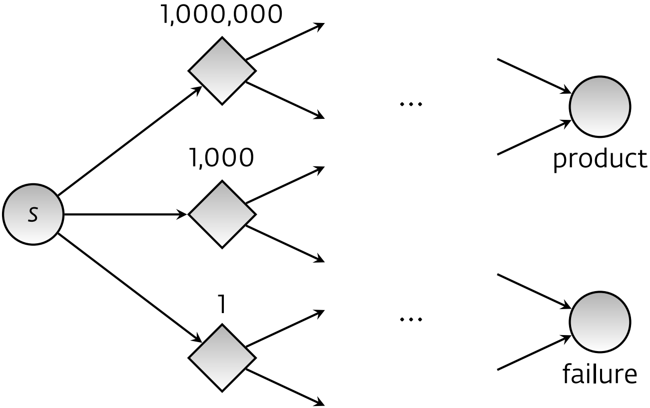

Question: How much to invest to successfully produce with 99%?

Compute: least r such that $s \models \exists \mathrm{P}_{\geq 0.99}(\mathit{true}\ \mathrm{U}_{\leq r}\ \text{product})$.

Quantile Queries

Quantile Query $\varphi = \forall \mathrm{P}_{\bowtie p}(a \,\mathrm{U}_{\leq ?}\, b)$ or $\varphi = \exists \mathrm{P}_{\bowtie p}(a \,\mathrm{U}_{\leq ?}\, b)$ where

- $a, b \in \mathrm{AP}$,
- $p \in [0, 1]$, and
- $\bowtie \in \{<, \leq, \geq, >\}$.

Quantile Queries

Quantile Query $\varphi = \forall \mathrm{P}_{\bowtie p}(a \mathbin{\mathrm{U}_{\leq ?}} b)$ or $\varphi = \exists \mathrm{P}_{\bowtie p}(a \mathbin{\mathrm{U}_{\leq ?}} b)$ where

- $a, b \in \mathrm{AP}$,
- $p \in [0, 1]$, and
- $\bowtie \in \{<, \leq, \geq, >\}$.

Write $\varphi[r]$ for the PRCTL formula that results from replacing ? by r.

Quantile Queries

Quantile Query $\varphi = \forall P_{\bowtie p}(a \, U_{\leq ?} \, b)$ or $\varphi = \exists P_{\bowtie p}(a \, U_{\leq ?} \, b)$ where

- $a, b \in AP$,
- $p \in [0,1]$, and
- $\bowtie \in \{<, \leq, \geq, >\}$.

Write $\varphi[r]$ for the PRCTL formula that results from replacing ? by r.

Define the value of s wrt. φ to be the least/largest r such that $s \models \varphi[r]$:

Quantile Queries

Quantile Query $\varphi = \forall \mathrm{P}_{\bowtie p}(a \,\mathrm{U}_{\leq ?}\, b)$ or $\varphi = \exists \mathrm{P}_{\bowtie p}(a \,\mathrm{U}_{\leq ?}\, b)$ where

- $a, b \in \mathrm{AP}$,
- $p \in [0,1]$, and
- $\bowtie \in \{<, \leq, \geq, >\}$.

Write $\varphi[r]$ for the PRCTL formula that results from replacing ? by r.

Define the value of s wrt. φ to be the least/largest r such that $s \models \varphi[r]$:

- $\mathrm{val}_\varphi(s) = \inf\{r \in \mathbb{R} : s \models \varphi[r]\}$ if $\bowtie \in \{\geq, >\}$ (minimising query).

Quantile Queries

Quantile Query $\varphi = \forall \mathrm{P}_{\bowtie p}(a \, \mathrm{U}_{\leq ?} \, b)$ or $\varphi = \exists \mathrm{P}_{\bowtie p}(a \, \mathrm{U}_{\leq ?} \, b)$ where

- $a, b \in \mathrm{AP}$,
- $p \in [0, 1]$, and
- $\bowtie \in \{<, \leq, \geq, >\}$.

Write $\varphi[r]$ for the PRCTL formula that results from replacing ? by r.

Define the value of s wrt. φ to be the least/largest r such that $s \models \varphi[r]$:

- $\mathrm{val}_\varphi(s) = \inf\{r \in \mathbb{R} : s \models \varphi[r]\}$ if $\bowtie \in \{\geq, >\}$ (minimising query).
- $\mathrm{val}_\varphi(s) = \sup\{r \in \mathbb{R} : s \models \varphi[r]\}$ if $\bowtie \in \{<, \leq\}$ (maximising query).

Quantile Queries

Quantile Query $\varphi = \forall \mathrm{P}_{\bowtie p}(a \,\mathrm{U}_{\leq ?}\, b)$ or $\varphi = \exists \mathrm{P}_{\bowtie p}(a \,\mathrm{U}_{\leq ?}\, b)$ where

- $a, b \in \mathrm{AP}$,
- $p \in [0, 1]$, and
- $\bowtie \in \{<, \leq, \geq, >\}$.

Write $\varphi[r]$ for the PRCTL formula that results from replacing ? by r.

Define the value of s wrt. φ to be the least/largest r such that $s \models \varphi[r]$:

- $\mathrm{val}_\varphi(s) = \inf\{r \in \mathbb{R} : s \models \varphi[r]\}$ if $\bowtie \in \{\geq, >\}$ (minimising query).
- $\mathrm{val}_\varphi(s) = \sup\{r \in \mathbb{R} : s \models \varphi[r]\}$ if $\bowtie \in \{<, \leq\}$ (maximising query).

Note: 1. $\mathrm{val}_\varphi(s) = -\infty$ or $\mathrm{val}_\varphi(s) \geq 0$.

2. $s \models \varphi[\mathrm{val}_\varphi(s)]$ for minimising queries with finite value.

Properties of the value

Two reasons for $\mathrm{val}_\varphi(s) = \infty$:

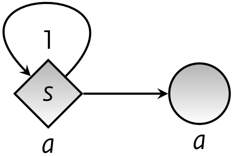

$$s \not\models \forall \mathrm{P}_{\geq 1}(a \,\mathrm{U}_{\leq \infty}\, b) \implies \mathrm{val}_\varphi(s) = \infty.$$

Properties of the value

Two reasons for $\mathrm{val}_\varphi(s) = \infty$:

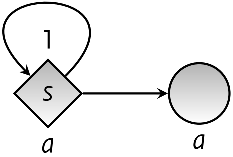

$$s \nvDash \forall P_{\geq 1}(a \; U_{\leq \infty} \; b) \implies \mathrm{val}_\varphi(s) = \infty.$$

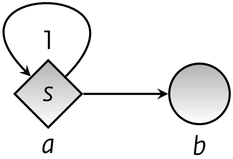

$$s \nvDash \forall P_{\geq 1}(a \; U_{\leq r} \; b) \text{ for all } r \in \mathbb{R} \implies \mathrm{val}_\varphi(s) = \infty.$$

Properties of the value

Two reasons for $\mathrm{val}_\varphi(s) = \infty$:

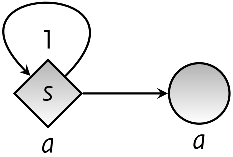

$$s \nvDash \forall \mathrm{P}_{\geq 1}(a \,\mathrm{U}_{\leq \infty}\, b) \implies \mathrm{val}_\varphi(s) = \infty.$$

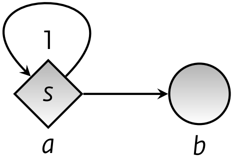

$$s \nvDash \forall \mathrm{P}_{\geq 1}(a \,\mathrm{U}_{\leq r}\, b) \text{ for all } r \in \mathbb{R} \implies \mathrm{val}_\varphi(s) = \infty.$$

Use classical PCTL model-checking algorithm to decide which is the case.

The Dual Query

Question: Do we really need all eight query types?

The Dual Query

Question: Do we really need all eight query types?

Define the dual of a query $\varphi = \forall \mathrm{P}_{\bowtie p}(a \, \mathrm{U}_{\leq ?} \, b)$ to be the query $\overline{\varphi} = \exists \mathrm{P}_{\overline{\bowtie} p}(a \, \mathrm{U}_{\leq ?} \, b)$, where e.g. $\overline{<} = \geq$, and vice versa.

The Dual Query

Question: Do we really need all eight query types?

Define the dual of a query $\varphi = \forall P_{\bowtie p}(a\,U_{\leq ?}\,b)$ to be the query $\overline{\varphi} = \exists P_{\overline{\bowtie} p}(a\,U_{\leq ?}\,b)$, where e.g. $\overline{<} = \geq$, and vice versa.

Note: $\overline{\varphi}[r] \equiv \neg\varphi[r]$ for all $r \in \mathbb{R} \cup \{\pm\infty\}$.

The Dual Query

Question: Do we really need all eight query types?

Define the dual of a query $\varphi = \forall \mathrm{P}_{\bowtie p}(a \,\mathrm{U}_{\leq ?}\, b)$ to be the query $\overline{\varphi} = \exists \mathrm{P}_{\overline{\bowtie} p}(a \,\mathrm{U}_{\leq ?}\, b)$, where e.g. $\overline{<} = \geq$, and vice versa.

Note: $\overline{\varphi}[r] \equiv \neg\varphi[r]$ for all $r \in \mathbb{R} \cup \{\pm\infty\}$.

Proposition

Let $\mathcal{M}$ be an MDP and φ a quantile query. Then $\mathrm{val}_{\varphi}(s) = \mathrm{val}_{\overline{\varphi}}(s)$ for all states s of $\mathcal{M}$.

The Dual Query

Question: Do we really need all eight query types?

Define the dual of a query $\varphi = \forall \mathrm{P}_{\bowtie p}(a \, \mathrm{U}_{\leq ?} \, b)$ to be the query $\overline{\varphi} = \exists \mathrm{P}_{\overline{\bowtie} p}(a \, \mathrm{U}_{\leq ?} \, b)$, where e.g. $\overline{<} = \geq$, and vice versa.

Note: $\overline{\varphi}[r] \equiv \neg\varphi[r]$ for all $r \in \mathbb{R} \cup \{\pm\infty\}$.

Proposition

Let $\mathcal{M}$ be an MDP and φ a quantile query. Then $\mathrm{val}_{\varphi}(s) = \mathrm{val}_{\overline{\varphi}}(s)$ for all states s of $\mathcal{M}$.

Proof:

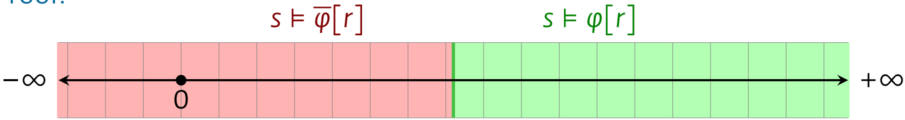

The Dual Query

Question: Do we really need all eight query types?

Define the dual of a query $\varphi = \forall \mathrm{P}_{\bowtie p}(a \, \mathrm{U}_{\leq ?} \, b)$ to be the query $\overline{\varphi} = \exists \mathrm{P}_{\overline{\bowtie} p}(a \, \mathrm{U}_{\leq ?} \, b)$, where e.g. $\overline{<} = \geq$, and vice versa.

Note: $\overline{\varphi}[r] \equiv \neg\varphi[r]$ for all $r \in \mathbb{R} \cup \{\pm\infty\}$.

Proposition

Let $\mathcal{M}$ be an MDP and φ a quantile query. Then $\mathrm{val}_{\varphi}(s) = \mathrm{val}_{\overline{\varphi}}(s)$ for all states s of $\mathcal{M}$.

Proof:

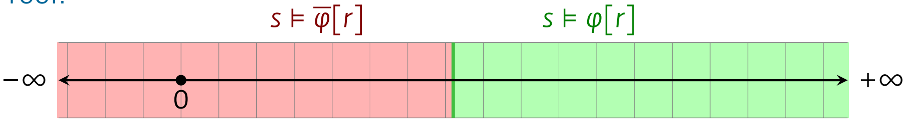

Consequence: May restrict to minimising queries.

Qualitative Queries

A quantile query $\forall \mathrm{P}_{\bowtie p}(a \, \mathrm{U}_{\leq ?} \, b)$ or $\exists \mathrm{P}_{\bowtie p}(a \, \mathrm{U}_{\leq ?} \, b)$ is qualitative if $p \in \{0, 1\}$.

Qualitative Queries

A quantile query $\forall \mathrm{P}_{\bowtie p}(a \, \mathrm{U}_{\leq ?} \, b)$ or $\exists \mathrm{P}_{\bowtie p}(a \, \mathrm{U}_{\leq ?} \, b)$ is qualitative if $p \in \{0, 1\}$.

Theorem

Qualitative queries can be evaluated in strongly polynomial time.

Qualitative Queries

A quantile query $\forall \mathrm{P}_{\bowtie p}(a \, \mathrm{U}_{\leq ?} \, b)$ or $\exists \mathrm{P}_{\bowtie p}(a \, \mathrm{U}_{\leq ?} \, b)$ is qualitative if $p \in \{0, 1\}$.

Theorem

Qualitative queries can be evaluated in strongly polynomial time.

Previous result: in P for *non-zeno* MDPs.

Qualitative Queries

A quantile query $\forall \mathrm{P}_{\bowtie p}(a \,\mathrm{U}_{\leq ?}\, b)$ or $\exists \mathrm{P}_{\bowtie p}(a \,\mathrm{U}_{\leq ?}\, b)$ is qualitative if $p \in \{0,1\}$.

Theorem

Qualitative queries can be evaluated in strongly polynomial time.

Previous result: in P for *non-zeno* MDPs.

Example: $\varphi = \forall \mathrm{P}_{>0}(a \,\mathrm{U}_{\leq ?}\, b)$

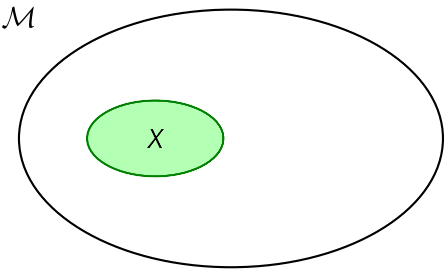

$R = \{0\}$

$X = \{s : rew(s) > 0, s \models b\}$

Qualitative Queries

A quantile query $\forall \mathrm{P}_{\bowtie p}(a \,\mathrm{U}_{\leq ?}\, b)$ or $\exists \mathrm{P}_{\bowtie p}(a \,\mathrm{U}_{\leq ?}\, b)$ is qualitative if $p \in \{0, 1\}$.

Theorem

Qualitative queries can be evaluated in strongly polynomial time.

Previous result: in P for *non-zeno* MDPs.

Example: $\varphi = \forall \mathrm{P}_{>0}(a \,\mathrm{U}_{\leq ?}\, b)$

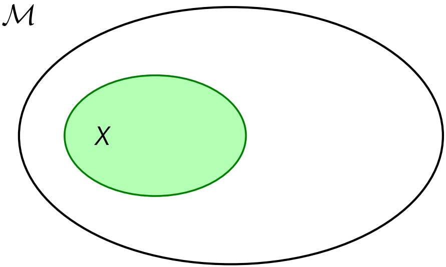

$R = \{r < r_1 < r_2 < \ldots\}$

$X = \{\text{discovered states}\}$

Qualitative Queries

A quantile query $\forall \mathrm{P}_{\bowtie p}(a \, \mathrm{U}_{\leq ?} \, b)$ or $\exists \mathrm{P}_{\bowtie p}(a \, \mathrm{U}_{\leq ?} \, b)$ is qualitative if $p \in \{0, 1\}$.

Theorem

Qualitative queries can be evaluated in strongly polynomial time.

Previous result: in P for *non-zeno* MDPs.

Example: $\varphi = \forall \mathrm{P}_{>0}(a \, \mathrm{U}_{\leq ?} \, b)$

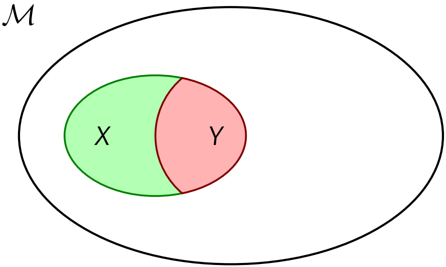

$R = \{r < r_1 < r_2 < \ldots\}$

$X = \{\text{discovered states}\}$

$Y = \{s \in X : \mathit{rew}(s) > 0, \mathrm{val}_\varphi(s) \leq r\}$

Qualitative Queries

A quantile query $\forall \mathrm{P}_{\bowtie p}(a \,\mathrm{U}_{\leq ?}\, b)$ or $\exists \mathrm{P}_{\bowtie p}(a \,\mathrm{U}_{\leq ?}\, b)$ is qualitative if $p \in \{0,1\}$.

Theorem

Qualitative queries can be evaluated in strongly polynomial time.

Previous result: in P for *non-zeno* MDPs.

Example: $\varphi = \forall \mathrm{P}_{>0}(a \,\mathrm{U}_{\leq ?}\, b)$

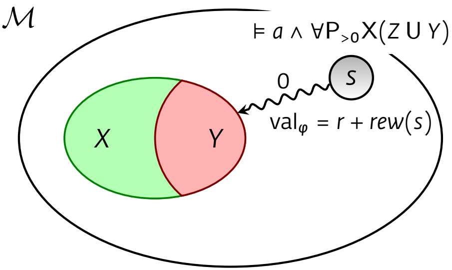

$R = \{r < r_1 < r_2 < \ldots\}$

$X = \{\text{discovered states}\}$

$Y = \{s \in X : rew(s) > 0, \mathrm{val}_\varphi(s) \leq r\}$

$Z = \{s : s \models a \wedge \neg b, rew(s) = 0\}$

Qualitative Queries

A quantile query $\forall \mathrm{P}_{\bowtie p}(a \, \mathrm{U}_{\leq ?} \, b)$ or $\exists \mathrm{P}_{\bowtie p}(a \, \mathrm{U}_{\leq ?} \, b)$ is qualitative if $p \in \{0, 1\}$.

Theorem

Qualitative queries can be evaluated in strongly polynomial time.

Previous result: in P for *non-zeno* MDPs.

Example: $\varphi = \forall \mathrm{P}_{>0}(a \, \mathrm{U}_{\leq ?} \, b)$

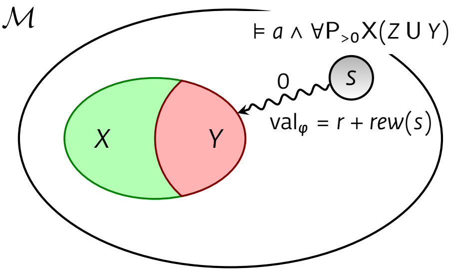

$R = \{r < r_1 < r_2 < \dots\}$

$X = \{\text{discovered states}\}$

$Y = \{s \in X : rew(s) > 0, \mathrm{val}_\varphi(s) \leq r\}$

$Z = \{s : s \models a \wedge \neg b, rew(s) = 0\}$

$X \leftarrow \{s \models a \wedge \forall \mathrm{P}_{>0} \mathrm{X}(Z \, \mathrm{U} \, Y)\}$

Qualitative Queries

A quantile query $\forall \mathrm{P}_{\bowtie p}(a \, \mathrm{U}_{\leq ?} \, b)$ or $\exists \mathrm{P}_{\bowtie p}(a \, \mathrm{U}_{\leq ?} \, b)$ is qualitative if $p \in \{0,1\}$.

Theorem

Qualitative queries can be evaluated in strongly polynomial time.

Previous result: in P for *non-zeno* MDPs.

Example: $\varphi = \forall \mathrm{P}_{>0}(a \, \mathrm{U}_{\leq ?} \, b)$

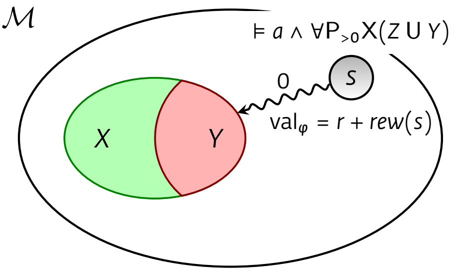

$R = \{r < r_1 < r_2 < \ldots\}$

$X = \{\text{discovered states}\}$

$Y = \{s \in X : rew(s) > 0, \mathrm{val}_\varphi(s) \leq r\}$

$Z = \{s : s \models a \wedge \neg b, rew(s) = 0\}$

$X \leftarrow \{s \models a \wedge \forall \mathrm{P}_{>0}\mathrm{X}(Z \, \mathrm{U} \, Y)\}$

$R \leftarrow R \setminus \{r\}$

Existential queries

Now let $\varphi = \exists \mathrm{P}_{>p}(a \,\mathrm{U}_{\leq ?}\, b)$ with $p \in (0,1)$.

Existential queries

Now let $\varphi = \exists \mathsf{P}_{>p}(a \,\mathsf{U}_{\leq ?}\, b)$ with $p \in (0,1)$.

Fact 1: For each $r \in \mathbb{N}$ the probabilities $\max_\sigma \Pr_s^\sigma(a \,\mathsf{U}_{\leq i}\, b)$, $0 \leq i \leq r$ can be computed in time $\mathrm{poly}(r \cdot |\mathcal{M}|)$.

Existential queries

Now let $\varphi = \exists \mathrm{P}_{>p}(a \, \mathrm{U}_{\leq ?} \, b)$ with $p \in (0,1)$.

Fact 1: For each $r \in \mathbb{N}$ the probabilities $\max_\sigma \Pr_s^\sigma(a \, \mathrm{U}_{\leq i} \, b)$, $0 \leq i \leq r$ can be computed in time $\mathrm{poly}(r \cdot |\mathcal{M}|)$.

Fact 2: The probabilities $\max_\sigma \Pr_s^\sigma(a \, \mathrm{U}_{\leq r} \, b)$ converge to $\max_\sigma \Pr_s^\sigma(a \, \mathrm{U} \, b)$.

Existential queries

Now let $\varphi = \exists \mathrm{P}_{>p}(a \,\mathrm{U}_{\leq ?}\, b)$ with $p \in (0,1)$.

Fact 1: For each $r \in \mathbb{N}$ the probabilities $\max_\sigma \Pr_s^\sigma(a \,\mathrm{U}_{\leq i}\, b)$, $0 \leq i \leq r$ can be computed in time $\mathrm{poly}(r \cdot |\mathcal{M}|)$.

Fact 2: The probabilities $\max_\sigma \Pr_s^\sigma(a \,\mathrm{U}_{\leq r}\, b)$ converge to $\max_\sigma \Pr_s^\sigma(a \,\mathrm{U}\, b)$.

Algorithm for computing $\mathrm{val}_\varphi(s)$:

- If $p \geq \max_\sigma \Pr_s^\sigma(a \,\mathrm{U}\, b)$, then return ∞.
- Otherwise compute $\max_\sigma \Pr_s^\sigma(a \,\mathrm{U}_{\leq i}\, b)$ for $i = 0, 1, 2, 3, \ldots$ until probability exceeds p; return i.

Existential queries

Now let $\varphi = \exists \mathrm{P}_{>p}(a \, \mathrm{U}_{\leq ?} \, b)$ with $p \in (0,1)$.

Fact 1: For each $r \in \mathbb{N}$ the probabilities $\max_\sigma \mathrm{Pr}_s^\sigma(a \, \mathrm{U}_{\leq i} \, b)$, $0 \leq i \leq r$ can be computed in time $\mathrm{poly}(r \cdot |\mathcal{M}|)$.

Fact 2: The probabilities $\max_\sigma \mathrm{Pr}_s^\sigma(a \, \mathrm{U}_{\leq r} \, b)$ converge to $\max_\sigma \mathrm{Pr}_s^\sigma(a \, \mathrm{U} \, b)$.

Algorithm for computing $\mathrm{val}_\varphi(s)$:

- If $p \geq \max_\sigma \mathrm{Pr}_s^\sigma(a \, \mathrm{U} \, b)$, then return ∞.
- Otherwise compute $\max_\sigma \mathrm{Pr}_s^\sigma(a \, \mathrm{U}_{\leq i} \, b)$ for $i = 0, 1, 2, 3, \ldots$ until probability exceeds p; return i.

How big can the value get???

Bounding the Value

Lemma

Let $\mathcal{M}$ be an MDP with n states where the denominator of each transition probability is at most m, $\varphi = \exists \mathrm{P}_{>p}(a \,\mathrm{U}_{\leq ?}\, b)$, and $p < q = \max_\sigma \Pr_s^\sigma(a \,\mathrm{U}\, b)$. Then $\mathrm{val}_\varphi(s) \leq -\lfloor \ln(q-p) \rfloor \cdot n \cdot \max \mathrm{reward} \cdot m^n$.

Bounding the Value

Lemma

Let $\mathcal{M}$ be an MDP with n states where the denominator of each transition probability is at most m, $\varphi = \exists \mathrm{P}_{>p}(a \,\mathrm{U}_{\leq ?}\, b)$, and $p < q = \max_\sigma \mathrm{Pr}_s^\sigma(a \,\mathrm{U}\, b)$. Then $\mathrm{val}_\varphi(s) \leq -\lfloor \ln(q - p) \rfloor \cdot n \cdot \max \mathrm{reward} \cdot m^n$.

Theorem

Queries of the form $\exists \mathrm{P}_{>p}(a \,\mathrm{U}_{\leq ?}\, b)$ can be evaluated in exponential time.

Bounding the Value

Lemma

Let $\mathcal{M}$ be an MDP with n states where the denominator of each transition probability is at most m, $\varphi = \exists\mathrm{P}_{>p}(a \,\mathrm{U}_{\leq ?}\, b)$, and $p < q = \max_\sigma \mathrm{Pr}_s^\sigma(a \,\mathrm{U}\, b)$. Then $\mathrm{val}_\varphi(s) \leq -\lfloor \ln(q-p) \rfloor \cdot n \cdot \max \mathrm{reward} \cdot m^n$.

Theorem

Queries of the form $\exists\mathrm{P}_{>p}(a \,\mathrm{U}_{\leq ?}\, b)$ can be evaluated in exponential time.

Question: What about $\exists\mathrm{P}_{\geq p}(a \,\mathrm{U}_{\leq ?}\, b)$?

Bounding the Value

Lemma

Let $\mathcal{M}$ be an MDP with n states where the denominator of each transition probability is at most m, $\varphi = \exists \mathrm{P}_{>p}(a \, \mathrm{U}_{\leq ?} \, b)$, and $p < q = \max_\sigma \mathrm{Pr}_s^\sigma(a \, \mathrm{U} \, b)$. Then $\mathrm{val}_\varphi(s) \leq -\lfloor \ln(q-p) \rfloor \cdot n \cdot \max \mathrm{reward} \cdot m^n$.

Theorem

Queries of the form $\exists \mathrm{P}_{>p}(a \, \mathrm{U}_{\leq ?} \, b)$ can be evaluated in exponential time.

Question: What about $\exists \mathrm{P}_{\geq p}(a \, \mathrm{U}_{\leq ?} \, b)$?

Lemma still applies but $p = q$ does not entail $\mathrm{val}_\varphi(s) = \infty$!

Bounding the Value

Lemma

Let $\mathcal{M}$ be an MDP with n states where the denominator of each transition probability is at most m, $\varphi = \exists \mathrm{P}_{>p}(a \, \mathrm{U}_{\leq ?} \, b)$, and $p < q = \max_\sigma \mathrm{Pr}_s^\sigma(a \, \mathrm{U} \, b)$. Then $\mathrm{val}_\varphi(s) \leq -\lfloor \ln(q-p) \rfloor \cdot n \cdot \max \mathrm{reward} \cdot m^n$.

Theorem

Queries of the form $\exists \mathrm{P}_{>p}(a \, \mathrm{U}_{\leq ?} \, b)$ can be evaluated in exponential time.

Question: What about $\exists \mathrm{P}_{\geq p}(a \, \mathrm{U}_{\leq ?} \, b)$?

Lemma still applies but $p = q$ does not entail $\mathrm{val}_\varphi(s) = \infty$!

Lemma

Let $\mathcal{M}$ be an MDP with n states, $\varphi = \exists \mathrm{P}_{\geq p}(a \, \mathrm{U}_{\leq ?} \, b)$, and $p = \max_\sigma \mathrm{Pr}_s^\sigma(a \, \mathrm{U} \, b)$. Then $\mathrm{val}_\varphi(s) = \infty$ or $\mathrm{val}_\varphi(s) \leq n \cdot \max \mathrm{reward}$.

Universal Queries

For queries of the form $\forall \mathrm{P}_{>p}(a \, \mathrm{U}_{\leq ?} \, b)$ we can get the same result.

Universal Queries

For queries of the form $\forall \mathrm{P}_{>p}(a \,\mathrm{U}_{\leq ?}\, b)$ we can get the same result.

Fact 1: For each $r \in \mathbb{N}$ the probabilities $\min_\sigma \mathrm{Pr}_s^\sigma(a \,\mathrm{U}_{\leq i}\, b)$, $0 \leq i \leq r$ can be computed in time $\mathrm{poly}(r \cdot |\mathcal{M}|)$.

Universal Queries

For queries of the form $\forall\mathrm{P}_{>p}(a \,\mathrm{U}_{\leq ?}\, b)$ we can get the same result.

Fact 1: For each $r \in \mathbb{N}$ the probabilities $\min_\sigma \Pr_s^\sigma(a \,\mathrm{U}_{\leq i}\, b)$, $0 \leq i \leq r$ can be computed in time $\mathrm{poly}(r \cdot |\mathcal{M}|)$.

Fact 2: The probabilities $\min_\sigma \Pr_s^\sigma(a \,\mathrm{U}_{\leq r}\, b)$ converge to $\min_\sigma \Pr_s^\sigma(a \,\mathrm{U}\, b)$.

Universal Queries

For queries of the form $\forall \mathrm{P}_{>p}(a \,\mathrm{U}_{\leq ?}\, b)$ we can get the same result.

Fact 1: For each $r \in \mathbb{N}$ the probabilities $\min_\sigma \Pr_s^\sigma(a \,\mathrm{U}_{\leq i}\, b)$, $0 \leq i \leq r$ can be computed in time $\mathrm{poly}(r \cdot |\mathcal{M}|)$.

Fact 2: The probabilities $\min_\sigma \Pr_s^\sigma(a \,\mathrm{U}_{\leq r}\, b)$ converge to $\min_\sigma \Pr_s^\sigma(a \,\mathrm{U}\, b)$.

Lemma

Let $\mathcal{M}$ be an MDP with n states where the denominator of each transition probability is at most m, $\varphi = \forall \mathrm{P}_{>p}(a \,\mathrm{U}_{\leq ?}\, b)$, and $p < q = \min_\sigma \Pr_s^\sigma(a \,\mathrm{U}\, b)$. Then $\mathrm{val}_\varphi(s) \leq -\lfloor \ln(q-p) \rfloor \cdot n \cdot \max \mathrm{reward} \cdot m^n$.

Universal Queries

For queries of the form $\forall \mathrm{P}_{>p}(a \,\mathrm{U}_{\leq ?}\, b)$ we can get the same result.

Fact 1: For each $r \in \mathbb{N}$ the probabilities $\min_\sigma \Pr_s^\sigma(a \,\mathrm{U}_{\leq i}\, b)$, $0 \leq i \leq r$ can be computed in time $\mathrm{poly}(r \cdot |\mathcal{M}|)$.

Fact 2: The probabilities $\min_\sigma \Pr_s^\sigma(a \,\mathrm{U}_{\leq r}\, b)$ converge to $\min_\sigma \Pr_s^\sigma(a \,\mathrm{U}\, b)$.

Lemma

Let $\mathcal{M}$ be an MDP with n states where the denominator of each transition probability is at most m, $\varphi = \forall \mathrm{P}_{>p}(a \,\mathrm{U}_{\leq ?}\, b)$, and $p < q = \min_\sigma \Pr_s^\sigma(a \,\mathrm{U}\, b)$. Then $\mathrm{val}_\varphi(s) \leq -\lfloor \ln(q-p) \rfloor \cdot n \cdot \max \mathrm{reward} \cdot m^n$.

Theorem

Queries of the form $\forall \mathrm{P}_{>p}(a \,\mathrm{U}_{\leq ?}\, b)$ can be evaluated in exponential time.

Universal Queries

For queries of the form $\forall \mathrm{P}_{>p}(a \,\mathrm{U}_{\leq ?}\, b)$ we can get the same result.

Fact 1: For each $r \in \mathbb{N}$ the probabilities $\min_\sigma \Pr_s^\sigma(a \,\mathrm{U}_{\leq i}\, b)$, $0 \leq i \leq r$ can be computed in time $\mathrm{poly}(r \cdot |\mathcal{M}|)$.

Fact 2: The probabilities $\min_\sigma \Pr_s^\sigma(a \,\mathrm{U}_{\leq r}\, b)$ converge to $\min_\sigma \Pr_s^\sigma(a \,\mathrm{U}\, b)$.

Lemma

Let $\mathcal{M}$ be an MDP with n states where the denominator of each transition probability is at most m, $\varphi = \forall \mathrm{P}_{>p}(a \,\mathrm{U}_{\leq ?}\, b)$, and $p < q = \min_\sigma \Pr_s^\sigma(a \,\mathrm{U}\, b)$. Then $\mathrm{val}_\varphi(s) \leq -\lfloor \ln(q-p) \rfloor \cdot n \cdot \max \mathrm{reward} \cdot m^n$.

Theorem

Queries of the form $\forall \mathrm{P}_{>p}(a \,\mathrm{U}_{\leq ?}\, b)$ can be evaluated in exponential time.

Open: Algorithm for evaluating $\forall \mathrm{P}_{\geq p}(a \,\mathrm{U}_{\leq ?}\, b)$.

A Counter-Example

Question: Assume $\varphi = \forall \mathrm{P}_{\geq p}(a \,\mathrm{U}_{\leq ?}\, b)$, where $p = \min_\sigma \Pr_s^\sigma(a \,\mathrm{U}\, b)$, and $\mathrm{val}_\varphi(s) < \infty$. Then $\mathrm{val}_\varphi(s) \leq |S| \cdot \max \text{reward}$?

A Counter-Example

Question: Assume $\varphi = \forall \mathrm{P}_{\geq p}(a \,\mathrm{U}_{\leq ?}\, b)$, where $p = \min_\sigma \Pr_s^\sigma(a \,\mathrm{U}\, b)$, and $\mathrm{val}_\varphi(s) < \infty$. Then $\mathrm{val}_\varphi(s) \leq |S| \cdot \max \mathrm{reward}$?

Answer: No! Let $p = \frac{1}{2}$ and $0 < q < 1$.

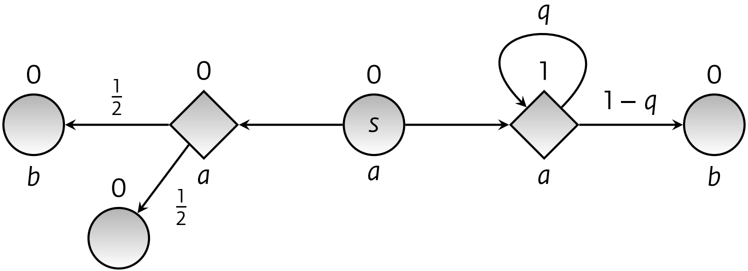

A Counter-Example

Question: Assume $\varphi = \forall \mathsf{P}_{\geq p}(a \,\mathsf{U}_{\leq ?}\, b)$, where $p = \min_\sigma \Pr_s^\sigma(a \,\mathsf{U}\, b)$, and $\mathrm{val}_\varphi(s) < \infty$. Then $\mathrm{val}_\varphi(s) \leq |S| \cdot \max \mathrm{reward}$?

Answer: No! Let $p = \frac{1}{2}$ and $0 < q < 1$.

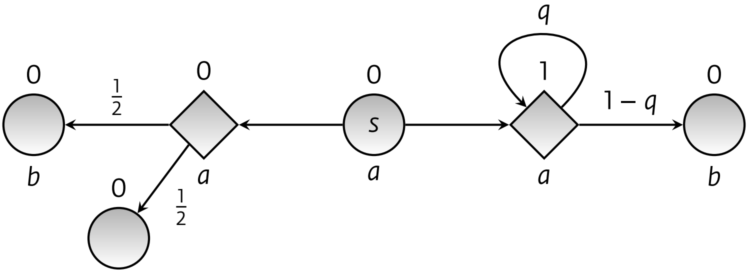

Note: $\mathrm{val}_\varphi(s) = -\lfloor 1/\log_2 q \rfloor$.

Markov Chains

Reduce MCs with non-negative rewards to MCs with rewards 0 and 1:

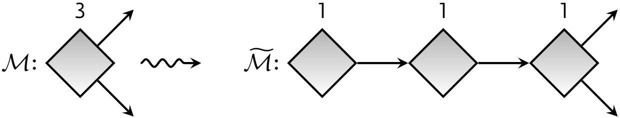

$|\widetilde{\mathcal{M}}| \le |\mathcal{M}| \cdot$ max reward

Markov Chains

Reduce MCs with non-negative rewards to MCs with rewards 0 and 1:

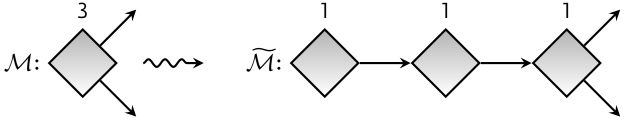

$|\widetilde{\mathcal{M}}| \leq |\mathcal{M}| \cdot \text{max reward}$

Fact: For a Markov chain with rewards 0 and 1, the probabilities $\Pr_s(a \, \mathsf{U}_{\leq r} \, b)$ can be computed in time $\text{poly}(|\mathcal{M}|) \cdot \log r$.

Markov Chains

Reduce MCs with non-negative rewards to MCs with rewards 0 and 1:

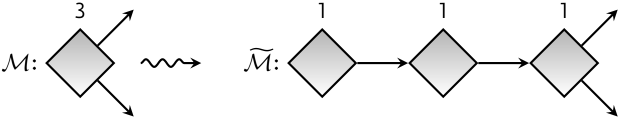

$|\widetilde{\mathcal{M}}| \le |\mathcal{M}| \cdot \text{max reward}$

Fact: For a Markov chain with rewards 0 and 1, the probabilities $\Pr_s(a \,\mathrm{U}_{\le r}\, b)$ can be computed in time $\text{poly}(|\mathcal{M}|) \cdot \log r$.

Hence: Binary search in the interval $[0, -\lfloor \ln(q-p) \rfloor \cdot n \cdot \text{max reward} \cdot m^n]$ with $q = \Pr_s(a \,\mathrm{U}\, b)$ can be used to determine $\text{val}_\varphi(s)$ for $\varphi = \mathrm{P}_{>p}(a \,\mathrm{U}_{\le ?}\, b)$.

Markov Chains

Reduce MCs with non-negative rewards to MCs with rewards 0 and 1:

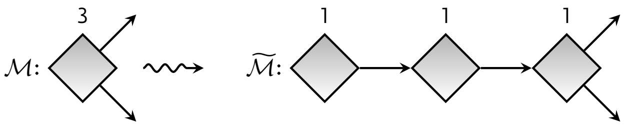

$|\widetilde{\mathcal{M}}| \leq |\mathcal{M}| \cdot \text{max reward}$

Fact: For a Markov chain with rewards 0 and 1, the probabilities $\Pr_s(a \, \mathsf{U}_{\leq r} \, b)$ can be computed in time $\text{poly}(|\mathcal{M}|) \cdot \log r$.

Hence: Binary search in the interval $[0, -\lfloor \ln(q-p) \rfloor \cdot n \cdot \text{max reward} \cdot m^n]$ with $q = \Pr_s(a \, \mathsf{U} \, b)$ can be used to determine $\text{val}_\varphi(s)$ for $\varphi = \mathsf{P}_{>p}(a \, \mathsf{U}_{\leq ?} \, b)$.

Note: $-\lfloor \ln(q-p) \rfloor \leq \text{poly}(|\mathcal{M}|) + \|p\|$.

Markov Chains

Reduce MCs with non-negative rewards to MCs with rewards 0 and 1:

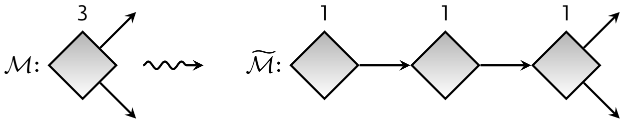

$|\widetilde{\mathcal{M}}| \leq |\mathcal{M}| \cdot \text{max reward}$

Fact: For a Markov chain with rewards 0 and 1, the probabilities $\Pr_s(a \, \mathrm{U}_{\leq r} \, b)$ can be computed in time $\mathrm{poly}(|\mathcal{M}|) \cdot \log r$.

Hence: Binary search in the interval $[0, -\lfloor \ln(q-p) \rfloor \cdot n \cdot \text{max reward} \cdot m^n]$ with $q = \Pr_s(a \, \mathrm{U} \, b)$ can be used to determine $\mathrm{val}_\varphi(s)$ for $\varphi = \mathrm{P}_{>p}(a \, \mathrm{U}_{\leq ?} \, b)$.

Note: $-\lfloor \ln(q-p) \rfloor \leq \mathrm{poly}(|\mathcal{M}|) + \|p\|$.

Theorem

Quantile Queries can be evaluated in pseudo-polynomial time on Markov chains.

Conclusion

Results:

- Polynomial algorithm for qualitative queries.
- Exponential algorithm for quantitative queries.
- Pseudo-polynomial algorithm for Markov Chains.

Conclusion

Results:

- Polynomial algorithm for qualitative queries.
- Exponential algorithm for quantitative queries.
- Pseudo-polynomial algorithm for Markov Chains.

Future work:

- Queries of the form $Q(a \, \mathsf{U}_{>r} \, b)$.
- Long-run average rewards.
- PRCTL with parameters.