

Modeling and Verification of Probabilistic Systems

Joost-Pieter Katoen

Lehrstuhl für Informatik 2
Software Modeling and Verification Group

<http://moves.rwth-aachen.de/teaching/ws-1819/movep18/>

November 26, 2018

Overview

- 1 PCTL Semantics
- 2 PCTL Model Checking
- 3 Complexity
- 4 Example: Dining Cryptographers Problem
- 5 Fairness
- 6 Summary

Probabilistic Computation Tree Logic

- ▶ PCTL is a language for formally specifying properties over DTMCs.
- ▶ It can also be used to specify properties over MDPs.
- ▶ It is a branching-time temporal logic based on CTL.
- ▶ Formula interpretation is Boolean, i.e., a state satisfies a formula or not.
- ▶ The main operator is $\mathbb{P}_J(\varphi)$
 - ▶ where φ constrains the set of paths and J is a threshold on the probability.
 - ▶ it is the probabilistic counterpart of \exists and \forall path-quantifiers in CTL.
 - ▶ ranges over all possible resolutions of nondeterminism.

$$s \models \mathbb{P}_J(\varphi) \text{ iff } \forall \text{ schedulers } \sigma. \Pr^\sigma(s \models \varphi) \in J$$

Probabilistic Computation Tree Logic: Syntax

PCTL consists of state- and path-formulas.

- ▶ PCTL *state formulas* over the set AP obey the grammar:

$$\Phi ::= \text{true} \mid a \mid \Phi_1 \wedge \Phi_2 \mid \neg \Phi \mid \mathbb{P}_J(\varphi)$$

where $a \in AP$, φ is a path formula and $J \subseteq [0, 1]$, $J \neq \emptyset$ is a non-empty interval.

- ▶ PCTL *path formulae* are formed according to the following grammar:

$$\varphi ::= \bigcirc \Phi \mid \Phi_1 \mathsf{U} \Phi_2 \mid \Phi_1 \mathsf{U}^{\leq n} \Phi_2$$

where Φ , Φ_1 , and Φ_2 are state formulae and $n \in \mathbb{N}$.

Abbreviate $\mathbb{P}_{[0,0.5]}(\varphi)$ by $\mathbb{P}_{\leq 0.5}(\varphi)$ and $\mathbb{P}_{[0,1]}(\varphi)$ by $\mathbb{P}_{>0}(\varphi)$.

Probabilistic Computation Tree Logic

- ▶ PCTL *state formulas* over the set AP obey the grammar:

$$\Phi ::= \text{true} \mid a \mid \Phi_1 \wedge \Phi_2 \mid \neg \Phi \mid \mathbb{P}_J(\varphi)$$

where $a \in AP$, φ is a path formula and $J \subseteq [0, 1]$, $J \neq \emptyset$.

- ▶ PCTL *path formulae* are formed according to the following grammar:

$$\varphi ::= \bigcirc \Phi \mid \Phi_1 \cup \Phi_2 \mid \Phi_1 \cup^{\leq n} \Phi_2 \quad \text{where } n \in \mathbb{N}.$$

Probabilistic Computation Tree Logic

- ▶ PCTL *state formulas* over the set AP obey the grammar:

$$\Phi ::= \text{true} \mid a \mid \Phi_1 \wedge \Phi_2 \mid \neg \Phi \mid \mathbb{P}_J(\varphi)$$

where $a \in AP$, φ is a path formula and $J \subseteq [0, 1]$, $J \neq \emptyset$.

- ▶ PCTL *path formulae* are formed according to the following grammar:

$$\varphi ::= \bigcirc \Phi \mid \Phi_1 \cup \Phi_2 \mid \Phi_1 \cup^{\leq n} \Phi_2 \quad \text{where } n \in \mathbb{N}.$$

Intuitive semantics

- ▶ $s_0 \alpha_0 s_1 \alpha_1 s_2 \alpha_2 \dots \models \Phi \cup^{\leq n} \Psi$ if Φ holds until Ψ holds within n steps (where $s_i \alpha_{i+1}$ is a single step).

$$\underbrace{s_0 \alpha_0 s_1 \alpha_1 s_2 \alpha_2 \dots s_n}_{s_j \models \Phi} \models \Psi \quad \begin{matrix} s_j \models \Psi \\ \forall j < n \end{matrix}$$

Probabilistic Computation Tree Logic

- ▶ PCTL *state formulas* over the set AP obey the grammar:

$$\Phi ::= \text{true} \mid a \mid \Phi_1 \wedge \Phi_2 \mid \neg \Phi \mid \mathbb{P}_J(\varphi)$$

where $a \in AP$, φ is a path formula and $J \subseteq [0, 1]$, $J \neq \emptyset$.

- ▶ PCTL *path formulae* are formed according to the following grammar:

$$\varphi ::= \bigcirc \Phi \mid \Phi_1 \cup \Phi_2 \mid \Phi_1 \cup^{\leq n} \Phi_2 \quad \text{where } n \in \mathbb{N}.$$

Intuitive semantics

- ▶ $s_0 \alpha_0 s_1 \alpha_1 s_2 \alpha_2 \dots \models \Phi \cup^{\leq n} \Psi$ if Φ holds until Ψ holds within n steps (where $s_i \alpha_{i+1}$ is a single step).
- ▶ $s \models \mathbb{P}_J(\varphi)$ if the probability under all policies that paths starting in s fulfill φ lies in J .

Overview

- 1 PCTL Semantics
- 2 PCTL Model Checking
- 3 Complexity
- 4 Example: Dining Cryptographers Problem
- 5 Fairness
- 6 Summary

Markov decision process (MDP)

Markov decision process

An MDP \mathcal{M} is a tuple $(S, Act, \mathbf{P}, \iota_{\text{init}}, AP, L)$ where

- ▶ S is a countable set of states with initial distribution $\iota_{\text{init}} : S \rightarrow [0, 1]$
- ▶ Act is a finite set of actions
- ▶ $\mathbf{P} : S \times Act \times S \rightarrow [0, 1]$, transition probability function such that:

$$\text{for all } s \in S \text{ and } \alpha \in Act : \sum_{s' \in S} \mathbf{P}(s, \alpha, s') \in \{0, 1\}$$

- ▶ AP is a set of atomic propositions and labeling $L : S \rightarrow 2^{AP}$.

Assumption: in each state at least one action is enabled.

PCTL semantics (1)

Notation

$\mathcal{M}, s \models \Phi$ if and only if state-formula Φ holds in state s of (possibly infinite) MDP \mathcal{M} . As \mathcal{M} is known from the context we simply write $s \models \Phi$.

Satisfaction relation for state formulas

The satisfaction relation \models is defined for PCTL state formulas by:

$$s \models a \quad \text{iff} \quad a \in L(s)$$

$$s \models \neg \Phi \quad \text{iff} \quad \text{not } (s \models \Phi)$$

$$s \models \Phi \wedge \Psi \quad \text{iff} \quad (s \models \Phi) \text{ and } (s \models \Psi)$$

$$s \models \mathbb{P}_J(\varphi) \quad \text{iff} \quad \text{for all policies } \mathfrak{G} \text{ on } \mathcal{M}. Pr^{\mathfrak{G}}(s \models \varphi) \in J$$

where $Pr^{\mathfrak{G}}(s \models \varphi) = Pr_s^{\mathfrak{G}}\{\pi \in Paths(s) \mid \pi \models \varphi\}$.

Semantics of \mathbb{P} -operator

The probabilistic operator $\mathbb{P}_J(\cdot)$ imposes probability bounds for *all* policies.

In particular, we have *for upper bounds*

$$\begin{array}{c}
 s \models \mathbb{P}_{\leq p}(\varphi) \text{ iff } \underbrace{Pr^{\text{max}}(s \models \varphi)}_{\substack{\downarrow \\ Pr^{\text{max}}(s \models \Diamond G)}} \leq p \text{ iff } \sup_{\mathcal{G}} Pr^{\mathcal{G}}(s \models \varphi) \leq p \\
 < p \qquad \qquad \qquad < p \qquad \qquad \qquad < p \\
 \varphi = \Diamond G \qquad \qquad \qquad Pr^{\text{max}}(s \models \Diamond G)
 \end{array}$$

Semantics of \mathbb{P} -operator

The probabilistic operator $\mathbb{P}_J(\cdot)$ imposes probability bounds for *all* policies. In particular, we have

$$s \models \mathbb{P}_{\leq p}(\varphi) \text{ iff } Pr^{\max}(s \models \varphi) \leq p \text{ iff } \sup_{\mathcal{G}} Pr^{\mathcal{G}}(s \models \varphi) \leq p$$

and, dually, for lower bounds

$$\underline{s \models \mathbb{P}_{\geq p}(\varphi)} \text{ iff } Pr^{\min}(s \models \varphi) \geq p \text{ iff } \inf_{\mathcal{G}} Pr^{\mathcal{G}}(s \models \varphi) \geq p.$$

iff

$$\forall \text{ policies } \sigma. Pr^{\sigma}(s \models \varphi) \geq p$$

Semantics of \mathbb{P} -operator

The probabilistic operator $\mathbb{P}_J(\cdot)$ imposes probability bounds for *all* policies. In particular, we have

$$s \models \mathbb{P}_{\leq p}(\varphi) \text{ iff } Pr^{\max}(s \models \varphi) \leq p \text{ iff } \sup_{\mathfrak{G}} Pr^{\mathfrak{G}}(s \models \varphi) \leq p$$

and, dually,

$$s \models \mathbb{P}_{\geq p}(\varphi) \text{ iff } Pr^{\min}(s \models \varphi) \geq p \text{ iff } \inf_{\mathfrak{G}} Pr^{\mathfrak{G}}(s \models \varphi) \geq p.$$

For finite MDPs we have:

$$Pr^{\max}(s \models \varphi) = \max_{\mathfrak{G}} Pr^{\mathfrak{G}}(s \models \varphi) \text{ and } Pr^{\min}(s \models \varphi) = \min_{\mathfrak{G}} Pr^{\mathfrak{G}}(s \models \varphi)$$

as for any finite MDP an *fm-policy* exists that maximises or minimises φ .

PCTL semantics (2)

Satisfaction relation for path formulas

Let $\pi = s_0 \alpha_0 s_1 \alpha_1 s_2 \alpha_2 \dots$ be an infinite path in (possibly infinite) MDP \mathcal{M} . Recall that $\pi[i] = s_i$ denotes the $(i+1)$ -st state along π .

The satisfaction relation \models is defined for state formulas by:

$$\pi \models \bigcirc \Phi \quad \text{iff} \quad s_1 \models \Phi$$

$$\pi \models \Phi \cup \Psi \quad \text{iff} \quad \exists k \geq 0. (\pi[k] \models \Psi \wedge \forall 0 \leq i < k. \pi[i] \models \Phi)$$

$$\pi \models \Phi \cup^{\leq n} \Psi \quad \text{iff} \quad \exists k \geq 0. (k \leq n \wedge \pi[k] \models \Psi \wedge \forall 0 \leq i < k. \pi[i] \models \Phi)$$

There is indeed no difference with the PCTL semantics for DTMC paths.

Equivalence of PCTL formulas

PCTL equivalence

$\Phi \equiv_{\text{MDP}} \Psi$ if and only if for all MDPs \mathcal{M} , it holds: $\text{Sat}_{\mathcal{M}}(\Phi) = \text{Sat}_{\mathcal{M}}(\Psi)$.

$\Phi \equiv_{\text{MC}} \Psi$ if and only if for all DTMCs \mathcal{D} , it holds: $\text{Sat}_{\mathcal{D}}(\Phi) = \text{Sat}_{\mathcal{D}}(\Psi)$.

Since any DTMC is an MDP, it follows: $\Phi \equiv_{\text{MDP}} \Psi$ implies $\Phi \equiv_{\text{MC}} \Psi$.

The converse, however, does not hold. For instance, for $p < 1$, we have $\mathbb{P}_{\leq p}(\varphi) \equiv_{\text{MC}} \neg \mathbb{P}_{> p}(\varphi)$. But, $\mathbb{P}_{\leq p}(\varphi) \not\equiv_{\text{MDP}} \neg \mathbb{P}_{> p}(\varphi)$.

$s \models \mathbb{P}_{\leq p}(\varphi)$ iff $Pr^{\mathfrak{G}}(s \models \varphi) \leq p$ for *all* policies \mathfrak{G} , but

$s \models \neg \mathbb{P}_{> p}(\varphi)$ iff not $s \models \mathbb{P}_{> p}(\varphi)$

iff not $\left(Pr^{\mathfrak{G}}(s \models \varphi) > p \text{ for all policies } \mathfrak{G} \right)$

iff $Pr^{\mathfrak{G}}(s \models \varphi) \leq p$ for *some* policy \mathfrak{G} .

Overview

- 1 PCTL Semantics
- 2 PCTL Model Checking**
- 3 Complexity
- 4 Example: Dining Cryptographers Problem
- 5 Fairness
- 6 Summary

PCTL model checking

PCTL model checking problem

Input: a finite MDP $\mathcal{M} = (S, Act, \mathbf{P}, \ell_{\text{init}}, AP, L)$, state $s \in S$, and PCTL state formula ϕ

Output: yes, if $s \models \phi$; no, otherwise.

Basic algorithm

In order to check whether $s \models \phi$ do:

1. Compute the **satisfaction set** $Sat(\phi) = \{s \in S \mid s \models \phi\}$.
2. This is done **recursively** by a bottom-up traversal of ϕ 's parse tree.
 - ▶ The nodes of the parse tree represent the subformulae of ϕ .
 - ▶ For each node, i.e., for each subformula ψ of ϕ , determine $Sat(\psi)$.
 - ▶ Determine $Sat(\psi)$ as function of the satisfaction sets of its children:
e.g., $Sat(\psi_1 \wedge \psi_2) = Sat(\psi_1) \cap Sat(\psi_2)$ and $Sat(\neg\psi) = S \setminus Sat(\psi)$.
3. Check whether state s belongs to $Sat(\phi)$.

Core model checking algorithm

Propositional formulas

$Sat(\cdot)$ is defined by structural induction as for PCTL on DTMCs.

Probabilistic operator \mathbb{P}

In order to determine whether $s \in Sat(\mathbb{P}_{\leq p}(\varphi))$, the probability $Pr^{\max}(s \models \varphi)$ needs to be established. Then

$$Sat(\mathbb{P}_{\leq p}(\varphi)) = \{s \in S \mid Pr^{\max}(s \models \varphi) \leq p\}.$$

Handwritten note: $0 \leq v \leq 1$

The same holds for strict upper bounds $< p$.

Similarly, lower bounds amount to determining $Pr^{\min}(s \models \varphi)$, e.g.,



$$Sat(\mathbb{P}_{> p}(\varphi)) = \{s \in S \mid Pr^{\min}(s \models \varphi) > p\}.$$

The next-step operator

Recall that: $s \models \mathbb{P}_{\leq p}(\bigcirc \phi)$ if and only if $Pr^{\max}(s \models \bigcirc \phi) \leq p$.

Lemma

$$Pr^{\max}(s \models \bigcirc \phi) = \max \left\{ \sum_{t \in \text{Sat}(\phi)} \mathbf{P}(s, \alpha, t) \mid \alpha \in \text{Act}(s) \right\}.$$



The next-step operator

Recall that: $s \models \mathbb{P}_{\leq p}(\bigcirc \phi)$ if and only if $Pr^{\max}(s \models \bigcirc \phi) \leq p$.

Lemma

$$Pr^{\max}(s \models \bigcirc \phi) = \max \left\{ \sum_{t \in \text{Sat}(\phi)} \mathbf{P}(s, \alpha, t) \mid \alpha \in \text{Act}(s) \right\}.$$


Algorithm

Determine $x_s = Pr^{\max}(s \models \bigcirc \phi)$ and return $\text{Sat}(\mathbb{P}_{\leq p}(\bigcirc \phi)) = \{s \in S \mid x_s \leq p\}$.

The case for minimal probabilities is similar and omitted here.

Example

Consider MDP:




and PCTL-formula:

$$\mathbb{P}_{\geq \frac{1}{2}} (\bigcirc \underbrace{heads})$$

1. $Sat(heads) = \{s_2\}$

Example

Consider MDP:




and PCTL-formula:

$$\mathbb{P}_{\geq \frac{1}{2}} (\bigcirc heads)$$


1. $Sat(heads) = \{s_2\}$
2. $x_{s_1} = Pr^{\min}(s_1 \models \bigcirc heads) = \min(0, 0.5) = 0$
3. Applying that to all states yields:

$$(Pr^{\min}(s \models \bigcirc \Phi))_{s \in S} = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0.7 & 0.3 & 0 & 0 \\ 0 & 0 & 0.5 & 0.5 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$



Example

Consider MDP:



and PCTL-formula:


$$\mathbb{P}_{\geq \frac{1}{2}} (\bigcirc heads)$$

1. $Sat(heads) = \{s_2\}$
2. $x_{s_1} = Pr^{\min}(s_1 \models \bigcirc heads) = \min(0, 0.5) = 0$
3. Applying that to all states yields:

$$(Pr^{\min}(s \models \bigcirc \Phi))_{s \in S} = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0.7 & 0.3 & 0 & 0 \\ 0 & 0 & 0.5 & 0.5 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0.5 \\ 1 \\ 0 \end{pmatrix}$$

Example

Consider MDP:



and PCTL-formula:

$$\mathbb{P}_{\geq \frac{1}{2}} (\bigcirc heads)$$

$$\mathbb{P}_{\leq \frac{1}{2}} (\bigcirc heads)$$

1. $Sat(heads) = \{s_2\}$
2. $x_{s_1} = Pr^{\min}(s_1 \models \bigcirc heads) = \min(0, 0.5) = 0$
3. Applying that to all states yields:

$$(Pr^{\min}(s \models \bigcirc \Phi))_{s \in S} = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0.7 & 0.3 & 0 & 0 \\ 0 & 0 & 0.5 & 0.5 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0.5 \\ 1 \\ 0 \end{pmatrix}$$

4. Thus: $Sat(\mathbb{P}_{\geq 0.5}(\bigcirc heads)) = \{s_2\}$.

$$\begin{aligned} \min &= (0, 0, 1, 0) \\ \max &= (0, 1/2, 1, 0) \end{aligned}$$

Bounded until (1)

Recall that: $s \models \mathbb{P}_{\geq p}(\phi \text{ U}^{\leq n} \psi)$ if and only if $Pr^{\min}(s \models \phi \text{ U}^{\leq n} \psi) \geq p$.

Lemma

Let $S_{=1} = \text{Sat}(\psi)$, $S_{=0} = S \setminus (\text{Sat}(\phi) \cup \text{Sat}(\psi))$, and $S_? = S \setminus (S_{=0} \cup S_{=1})$.

Then: $Pr^{\min}(s \models \phi \text{ U}^{\leq n} \psi)$ equals

$$\begin{cases} 1 & \text{if } s \in S_{=1} \\ 0 & \text{if } s \in S_{=0} \\ 0 & \text{if } s \in S_? \wedge n=0 \\ \min\left\{\sum_{s' \in S} \mathbf{P}(s, \alpha, s') \cdot Pr^{\min}(s' \models \phi \text{ U}^{\leq n-1} \psi) \mid \alpha \in \text{Act}(s)\right\} & \text{otherwise} \end{cases}$$

The case for maximal probabilities is analogous.

Bounded until (2)

Lemma

Let $S_{=1} = \text{Sat}(\psi)$, $S_{=0} = S \setminus (\text{Sat}(\phi) \cup \text{Sat}(\psi))$, and $S_? = S \setminus (S_{=0} \cup S_{=1})$.


Then: $Pr^{\min}(s \models \phi \text{ U}^{\leq n} \psi)$ equals

$$\begin{cases} 1 & \text{if } s \in S_{=1} \\ 0 & \text{if } s \in S_{=0} \\ 0 & \text{if } s \in S_? \wedge n=0 \\ \min\left\{\sum_{s' \in S} \mathbf{P}(s, \alpha, s') \cdot Pr^{\min}(s' \models \phi \text{ U}^{\leq n-1} \psi) \mid \alpha \in \text{Act}(s)\right\} & \text{otherwise} \end{cases}$$

Algorithm

1. Let $\mathbf{P}_{\phi, \psi}$ be the probability matrix of $\mathcal{M}[S_{=0} \cup S_{=1}]^1$.
2. Then $(Pr^{\min}(s \models \phi \text{ U}^{\leq 0} \psi))_{s \in S} = \mathbf{b}_{\psi}$
3. And $(Pr^{\min}(s \models \phi \text{ U}^{\leq i+1} \psi))_{s \in S} = \mathbf{P}_{\phi, \psi} \cdot (Pr^{\min}(s \models \phi \text{ U}^{\leq i} \psi))_{s \in S}$.
4. This requires n matrix-vector multiplications and n minimum operators.

Example



$$\varphi = \mathbb{P} < 0.95 \quad (\Diamond^{\leq 3} \text{init})$$

$$S_{\text{at}}(\text{init}) = \{0\} \quad S_{=1} = \{0\}$$

$$S_{?} = \{1, 2, 3\} \quad S_{=0} = \emptyset$$

$$\begin{aligned}
 (\Pr^{\max}(\Diamond^{\leq 0} \text{init})) &= (1, 0, 0, 0) \\
 (\Pr^{\max}(\Diamond^{\leq 1} \text{init})) &= (1, \underbrace{\frac{7}{10}}_b, 0, 0) \\
 &= (1, \frac{7}{10} + \frac{7}{10} \cdot \frac{3}{10}, 0, 0) \\
 &= (1, \underbrace{\frac{7}{10} + \frac{7}{10} \cdot \frac{3}{10} + \frac{7}{10} \cdot (\frac{3}{10})^2}_{0.93}, 0, 0)
 \end{aligned}$$

$\max \left(\underbrace{\frac{7}{10}}_b, \underbrace{0}_c \right)$
 $\text{Sat}(\varphi) = \{1, 2, 3\}$

Until

 $\geq 0 \quad \geq 1$

Recall that: $s \models \mathbb{P}_{\geq p}(\phi \cup \psi)$ if and only if $Pr^{\min}(s \models \phi \cup \psi) \geq p$.

Algorithm

1. Determine $S_{=1} = \text{Sat}(\mathbb{P}_{=1}(\phi \cup \psi))$ by a graph analysis.
2. Determine $S_{=0} = \text{Sat}(\mathbb{P}_{=0}(\phi \cup \psi))$ by a graph analysis.
3. Then solve a linear program (or use value iteration) over all remaining states.

Importance of pre-computation

1. Determining $S_{=0}$ ensures **unique** solution of linear program.
2. Determining $S_{=1}$ **reduces** the number of variables in the linear program.
3. Gives **exact** results for the states in $S_{=1}$ and $S_{=0}$ (i.e., no round-off).
4. For **qualitative** properties, no further computation is needed.

Precomputations

Qualitative reachability

1. Determine all states for which probability is zero

$$S=0$$

- 1.1 minimum: $\{s \in S \mid Pr^{\min}(s \models \phi \cup \psi) = 0\}$

- 1.2 maximum: $\{s \in S \mid Pr^{\max}(s \models \phi \cup \psi) = 0\}$

2. Determine all states for which probability is one

$$S=1$$

- 2.1 minimum: $\{s \in S \mid Pr^{\min}(s \models \phi \cup \psi) = 1\}$

- 2.2 maximum: $\{s \in S \mid Pr^{\max}(s \models \phi \cup \psi) = 1\}$

3. Then solve a linear program (or use value iteration) over all remaining states.

The first case has been treated in the previous lecture (for $\diamond G$).

Qualitative reachability

- ▶ Goal is to compute $\{s \in S \mid Pr^{\max}(s \models \Diamond G) = 1\}$
- ▶ First make all states in G absorbing, i.e., $\mathbf{P}(s, \alpha_s, s) = 1$
- ▶ Iteratively remove state t for which $Pr^{\max}(t \models \Diamond G) < 1$

Sketch of algorithm


1. Let $U_0 = S \setminus \underbrace{Sat(\exists \Diamond G)}_{Sat(\neg \exists \Diamond G)}$; this can be done by a graph analysis

Qualitative reachability

- ▶ Goal is to compute $\{s \in S \mid Pr^{\max}(s \models \Diamond G) = 1\}$
- ▶ First make all states in G absorbing, i.e., $\mathbf{P}(s, \alpha_s, s) = 1$
- ▶ Iteratively remove state t for which $Pr^{\max}(t \models \Diamond G) < 1$

Sketch of algorithm

1. Let $U_0 = S \setminus Sat(\exists \Diamond G)$; this can be done by a graph analysis
2. Remove all actions α from state u for which $Post(u, \alpha) \cap U_0 \neq \emptyset$



Qualitative reachability

- ▶ Goal is to compute $\{s \in S \mid Pr^{\max}(s \models \Diamond G) = 1\}$
- ▶ First make all states in G absorbing, i.e., $\mathbf{P}(s, \alpha_s, s) = 1$
- ▶ Iteratively remove state t for which $Pr^{\max}(t \models \Diamond G) < 1$

Sketch of algorithm

1. Let $U_0 = S \setminus \text{Sat}(\exists \Diamond G)$; this can be done by a graph analysis
2. Remove all actions α from state u for which $\text{Post}(u, \alpha) \cap U_0 \neq \emptyset$
3. If after removal of actions $\text{Act}(u) = \emptyset$, then remove state u
4. Repeat this procedure for all states, yielding the new MDP \mathcal{M}'
5. As this may yield new states from which G is unreachable, repeat the above steps until all states can reach G

This procedure is quadratic in the size of the MDP.

Algorithm

Algorithm 45 Computing the set of states s with $Pr^{\max}(s \models \Diamond B) = 1$


Input: MDP \mathcal{M} with finite state space S , $B \subseteq S$ for $s \in B$: $Act(s) = \{\alpha_s\}$ and $P(s, \alpha_s, s) = 1$
 (i.e., B is absorbing)

Output: $\{s \in S \mid Pr^{\max}(s \models \Diamond B) = 1\}$

```

 $U := \{s \in S \mid s \not\models \exists \Diamond B\};$ 
repeat
   $R := U;$ 
  while  $R \neq \emptyset$  do
    let  $u \in R;$ 
     $R := R \setminus \{u\};$ 
    for all  $(t, \alpha) \in Pre(u)$  such that  $t \notin U$  do
      remove  $\alpha$  from  $Act(t);$ 
      if  $Act(t) = \emptyset$  then
         $R := R \cup \{t\};$ 
         $U := U \cup \{t\};$ 
      fi
    od
    (* all incoming edges of  $u$  have been removed *)
    remove  $u$  and its outgoing edges from  $\mathcal{M}$ 
  od
  (* determine the states  $s$  that cannot reach  $B$  in the modified MDP *)
   $U := \{s \in S \setminus U \mid s \not\models \exists \Diamond B\};$ 
until  $U = \emptyset$ 
(* all states can reach  $B$  in the generated sub-MDP of  $\mathcal{M}$  *)
return all states in the remaining MDP

```



$$U_0 = \{0\}$$

$$R = \{0\}$$

1st iteration: $0 \xrightarrow{\alpha} 0 \rightarrow$ remove α from 0
 $\text{Act}(0) = \emptyset$

remove 0

$$U_1 = \{4\}$$

$4 \xrightarrow{\beta} 0 \rightarrow$ remove β from 4

2nd iteration: $6 \xrightarrow{\alpha} 4 \rightarrow$ remove α from 6
 $4 \xrightarrow{\alpha} 4 \rightarrow$ remove α from 4

$$U_2 = \emptyset$$

$$Pr^{\max}(s \models \Diamond 2) = 1 =$$

$$\{1, 2, 3, 5, 6\}$$

Overview

- 1 PCTL Semantics
- 2 PCTL Model Checking
- 3 Complexity**
- 4 Example: Dining Cryptographers Problem
- 5 Fairness
- 6 Summary

Time complexity

Let $|\Phi|$ be the **size** of Φ , i.e., the number of logical and temporal operators in Φ .

Time complexity of PCTL model checking of MDPs

For finite MDP \mathcal{M} and PCTL state-formula Φ , the PCTL model-checking problem can be solved in time

$$\mathcal{O}(\text{poly}(\text{size}(\mathcal{M})) \cdot n_{\max} \cdot |\Phi|)$$

where $n_{\max} = \max\{n \mid \psi_1 \text{ U}^{\leq n} \psi_2 \text{ occurs in } \Phi\}$ with $\max \emptyset = 1$.


Overview

- 1 PCTL Semantics
- 2 PCTL Model Checking
- 3 Complexity
- 4 Example: Dining Cryptographers Problem**
- 5 Fairness
- 6 Summary

Dining cryptographers problem

Dining cryptographer's protocol

1. Each cryptographer flips an unbiased coin and only informs the cryptographer on the right of the outcome.



Dining cryptographers problem

Dining cryptographer's protocol


1. Each cryptographer **flips an unbiased coin** and only informs the cryptographer **on the right** of the outcome.
2. Each cryptographer states whether the two coins that it can see—the one it flipped and the one the left-hand neighbour flipped—are the same (**agree**) or different (**disagree**).

Caveat: if a cryptographer actually paid for the dinner, then it instead states the opposite (**disagree** if the coins are the same and **agree** if the coins are different).

Claim

An odd number of **agrees** indicates that the master paid, while an even number indicates that a cryptographer paid.

Dining cryptographers problem



Example scenario in which master paid (left) or cryptographer A paid (right) and provides a misleading vote.

Dining cryptographers problem

Dining cryptographer's protocol

1. Each cryptographer **flips an unbiased coin** and only informs the cryptographer **on the right** of the outcome.
2. Each cryptographer states whether the two coins that it can see—the one it flipped and the one the left-hand neighbour flipped—are the same (**agree**) or different (**disagree**).

Caveat: if a cryptographer actually paid for the dinner, then it instead states the opposite (**disagree** if the coins are the same and **agree** if the coins are different).

Generalisation

The dining cryptographer's protocol can be generalised to any number N of cryptographers. Then:

- ▶ if N is odd, then an odd number of **agrees** indicates that the master paid while an even number indicates that a cryptographer paid.

Dining cryptographers problem

Dining cryptographer's protocol

1. Each cryptographer **flips an unbiased coin** and only informs the cryptographer **on the right** of the outcome.
2. Each cryptographer states whether the two coins that it can see—the one it flipped and the one the left-hand neighbour flipped—are the same (**agree**) or different (**disagree**).

Caveat: if a cryptographer actually paid for the dinner, then it instead states the opposite (**disagree** if the coins are the same and **agree** if the coins are different).

Generalisation

The dining cryptographer's protocol can be generalised to any number N of cryptographers. Then:

- ▶ if N is odd, then an odd number of **agrees** indicates that the master paid while an even number indicates that a cryptographer paid.
- ▶ if N is even, then an even number of **agrees** indicates that the master paid while an odd number indicates that a cryptographer paid.

MDP generation times

N:	Model:		Construction time (s):
	States:	Transitions:	
3	286	585	0.001
4	1,733	4,580	0.01
5	9,876	32,315	0.03
6	54,055	211,566	0.07
7	287,666	1,312,045	0.11
8	1,499,657	7,813,768	0.22
9	7,695,856	45,103,311	0.34
10	39,005,611	253,985,650	0.52
15	115,553,171,626	1,128,594,416,085	3.27
20	304,287,522,253,461	3,962,586,180,540,340	13.48

symbolically

The number of states and transitions in the MDP representing the model for the dining cryptographers problem with N cryptographers.

Checking correctness

N:	master pays:		cryptographers pay:	
	time:	iterations:	time:	iterations:
3	0.028	7	0.008	7
4	0.061	9	0.032	9
5	0.141	11	0.085	11
6	0.322	13	0.292	13
7	0.778	15	0.563	15
8	1.467	17	2.25	17
9	2.67	19	4.14	19
10	6.30	21	7.63	21
15	56.9	31	185	31
20	268	41	954	41

$pay \Rightarrow \mathbb{P}_{=1}(\Diamond(done \wedge par = N \% 2)) \wedge \neg pay \Rightarrow \mathbb{P}_{=1}(\Diamond(done \wedge par \neq N \% 2)).$
 That is: if the master paid, the parity of the number of **agrees** matches the parity of N and, if a cryptographer paid, it does not.


Checking anonymity

N:	minimum:			maximum:		
	time:	iterations:	probability:	time:	iterations:	probability:
3	0.099	8	0.25	0.004	8	0.25
4	0.041	10	0.125	0.006	10	0.125
5	0.172	12	0.0625	0.032	12	0.0625
6	0.231	14	0.03125	0.044	14	0.03125
7	0.595	16	0.015625	0.301	16	0.015625
8	1.111	18	0.0078125	0.540	18	0.0078125
9	2.12	20	0.00390625	1.31	20	0.00390625
10	3.53	22	0.001953125	2.67	22	0.001953125
15	45.1	32	6.103515625E-5	36.8	32	6.103515625E-5

To verify anonymity – when a cryptographer pays then no cryptographer can tell who has paid – we check that any possible combination of **agree** and **disagree** has the same likelihood no matter which of the cryptographers pays. This needs to be checked for all 2^N possible outcomes. Above the results are listed for one possible outcome.

Overview

- 1 PCTL Semantics
- 2 PCTL Model Checking
- 3 Complexity
- 4 Example: Dining Cryptographers Problem
- 5 Fairness**
- 6 Summary



$$\sigma_{\alpha}(s) = \alpha$$

$$\sigma_{\beta}(s) = \beta$$



ignores for an infinite path, infinitely often the other option.

⇒ this policy is perhaps not fair

$$\underline{\text{fair}} = \boxed{\Box \Diamond \text{yellow} \Rightarrow \Box \Diamond \text{green}} \quad (*)$$

σ_{β} is not fair w.r.t. (*) LTL formula

$$Pr_s^{\sigma_{\beta}} \{ \pi \models \underline{\text{fair}} \} = 0 < 1.$$

fairness assumption is defined by an LTL-formula fair.

Def. (fair policy)

MDP M and fair is LTL-formula.

A policy σ is fair if $\forall s \in M.$

$$Pr_s^{\sigma} \{ \pi \in \text{Paths}(s) \mid \pi \models \underline{\text{fair}} \} = 1.$$

Fairness

- ▶ A policy \mathcal{G} is **fair** if for every state s , the probability under \mathcal{G} of all fair paths from s is one

Fairness

- ▶ A policy \mathcal{G} is **fair** if for every state s , the probability under \mathcal{G} of all fair paths from s is one
- ▶ A fairness assumption is **realizable** in MDP \mathcal{M} if there is some fair policy for \mathcal{M}



in the example the policy that
alternates going left + going right
establishes fair, so fair is
realizable.

Fairness

- ▶ A policy \mathcal{G} is **fair** if for every state s , the probability under \mathcal{G} of all fair paths from s is one
- ▶ A fairness assumption is **realizable** in MDP \mathcal{M} if there is some fair policy for \mathcal{M}
- ▶ Realizable fairness assumptions are **irrelevant** for maximal reachability probabilities (i.e., safety)

Theorem

M finite MDP

$G \subseteq S$

fair is a realizable fairness assumption for M

Then:

$$(1) \sup_{\substack{\text{fair policy} \\ \sigma \in M}} \Pr^{\sigma}(s \models \Diamond G) = \underbrace{\Pr^{\max}(s \models \Diamond G)}_{= \sup_{\substack{\text{all policy} \\ \sigma' \in M}} \Pr^{\sigma'}(s \models \Diamond G)}$$

(2) there exists a finite memory policy that maximises the reachability probabilities.


So: for maximum probabilities,

imposing fairness does not result in any difference.

Fairness

- ▶ A policy \mathcal{G} is **fair** if for every state s , the probability under \mathcal{G} of all fair paths from s is one
- ▶ A fairness assumption is **realizable** in MDP \mathcal{M} if there is some fair policy for \mathcal{M}
- ▶ Realizable fairness assumptions are **irrelevant** for maximal reachability probabilities (i.e., safety)
- ▶ They are **relevant** for minimal reachability probabilities (i.e., liveness)

Fair policies may not achieve minimal reachability probabilities:



$$\begin{aligned} \text{fair} &= \Box \Diamond a \Rightarrow \Box \Diamond b \\ &= \Box \Diamond u \Rightarrow \Box \Diamond t \end{aligned}$$

any fair policy has to take action α ∞ often

then $\inf_{\sigma} \Pr^{\sigma}(s \models \Box b) = 1$ \leftarrow

fair policies σ

but for ∇_{β} $\Pr^{\sigma_{\beta}}(s \models \Box b) = 0$ \searrow

Fairness

- ▶ A policy \mathcal{G} is **fair** if for every state s , the probability under \mathcal{G} of all fair paths from s is one
- ▶ A fairness assumption is **realizable** in MDP \mathcal{M} if there is some fair policy for \mathcal{M}
- ▶ Realizable fairness assumptions are **irrelevant** for maximal reachability probabilities (i.e., safety)
- ▶ They are **relevant** for minimal reachability probabilities (i.e., liveness)
- ▶ Computing **minimal** reachability probabilities under **strongly fair** policies is **reducible** to computing **maximal** reachability probabilities

Overview

- 1 PCTL Semantics
- 2 PCTL Model Checking
- 3 Complexity
- 4 Example: Dining Cryptographers Problem
- 5 Fairness
- 6 Summary**

Theorem (computing Pr^{\min} under fair policies)

MDP M

$G \subseteq S$

fair is a strong fairness assumption e.g.

" $\Box \Diamond \dots \Rightarrow \Box \Diamond \dots$ "

$$\inf_{\text{fair policy } \sigma \text{ on } M} Pr^{\sigma}(s \models \Box G) = 1 - Pr^{\max}(s \models \neg G \cup F_{\Rightarrow}^{\min})$$

= 1 in my example

where $F_{=0}^{\min} = \{ t \in S \mid Pr^{\sigma}(t \models \Box G) = 0 \text{ for some fair policy } \sigma \}$

How to compute $F_{=0}^{\min}$?

lemma $s \in F_{=0}^{\min}$ iff $Pr^{\max}(s \models \neg G \cup V) = 1$

V is the union of all end components
"BSCCs"

of M that are fair and contain no G -state.

Summary

- ▶ PCTL is a variant of CTL with operator $\Phi = \mathbb{P}_J(\varphi)$.
- ▶ PCTL model checking is performed by a recursive descent over Φ .
- ▶ Checking whether $s \models \mathbb{P}_{>p}(\varphi)$ amounts to determine $Pr^{\min}(s \models \varphi)$.
- ▶ Checking whether $s \models \mathbb{P}_{<p}(\varphi)$ amounts to determine $Pr^{\max}(s \models \varphi)$.
- ▶ The next operator amounts to a single matrix-vector multiplication and a max/min.
- ▶ The bounded-until operator $U^{\leq n}$ amounts to n matrix-vector multiplications + n minimums (or maximums).
- ▶ The until-operator amounts to solving a linear inequation system.
- ▶ The worst-case time complexity is polynomial in the size of the MDP and linear in the size of the formula.