

Modeling and Verification of Probabilistic Systems

Lecture 13: Model Checking MDPs

Joost-Pieter Katoen

Lehrstuhl für Informatik 2
Software Modeling and Verification Group

<http://www-i2.informatik.rwth-aachen.de/i2/mvps11/>

June 6, 2011

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.

Overview

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

PCTL syntax

[Bianco & De Alfaro, 1995]

Probabilistic Computation Tree Logic: Syntax

PCTL consists of state- and path-formulas.

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

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

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:

$$\varphi ::= \bigcirc \Phi \quad | \quad \Phi_1 \cup \Phi_2 \quad | \quad \Phi_1 \cup^{\leq n} \Phi_2$$

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:

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

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 $s_i \alpha_i$ is a single step).
- $s \models \mathbb{P}_J(\varphi)$ if probability *under all policies* that paths starting in s fulfill φ lies in J .

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 $L : S \rightarrow 2^{AP}$.

Overview

1 PCTL Semantics

2 PCTL Model Checking

3 Complexity

4 Example: Dining Cryptographers Problem

5 Summary

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:

$$\begin{aligned} s \models a &\quad \text{iff } a \in L(s) \\ s \models \neg \Phi &\quad \text{iff not } (s \models \Phi) \\ s \models \Phi \wedge \Psi &\quad \text{iff } (s \models \Phi) \text{ and } (s \models \Psi) \\ s \models \mathbb{P}_J(\varphi) &\quad \text{iff for all policies } \mathcal{G} \text{ on } \mathcal{M}. \Pr^{\mathcal{G}}(s \models \varphi) \in J \end{aligned}$$

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

Semantics of \mathbb{P} -operator

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

In particular, we have

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

and, dually,

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

For finite MDPs we have:

$$\Pr^{\max}(s \models \varphi) = \max_{\mathcal{S}} \Pr^{\mathcal{S}}(s \models \varphi) \text{ and } \Pr^{\min}(s \models \varphi) = \min_{\mathcal{S}} \Pr^{\mathcal{S}}(s \models \varphi).$$

since for any finite MDP there exists an fm-policy that maximises or minimises φ .

Equivalence of PCTL formulas

PCTL equivalence

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

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

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

The converse, however, does not hold. For instance, for $p < 1$, we have

$\mathbb{P}_{\leq p}(\varphi) \equiv_{MC} \neg \mathbb{P}_{>p}(\varphi)$. But, $\mathbb{P}_{\leq p}(\varphi) \not\equiv_{MDP} \neg \mathbb{P}_{>p}(\varphi)$.

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

$$s \models \neg \mathbb{P}_{>p}(\varphi) \text{ iff } \text{not } \left(\Pr^{\mathcal{S}}(s \models \varphi) > p \text{ for all policies } s \in \mathcal{S} \right) \\ \text{iff } \Pr^{\mathcal{S}}(s \models \varphi) \leq p \text{ for *some* policy } \mathcal{S}.$$

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:

$$\begin{aligned} \pi \models \bigcirc \Phi &\text{ iff } s_1 \models \Phi \\ \pi \models \Phi \cup \Psi &\text{ iff } \exists k \geq 0. (\pi[k] \models \Psi \wedge \forall 0 \leq i < k. \pi[i] \models \Phi) \\ \pi \models \Phi \cup^{\leq n} \Psi &\text{ iff } \exists k \geq 0. (k \leq n \wedge \pi[k] \models \Psi \wedge \\ &\quad \forall 0 \leq i < k. \pi[i] \models \Phi) \end{aligned}$$

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

Overview

1 PCTL Semantics

2 PCTL Model Checking

3 Complexity

4 Example: Dining Cryptographers Problem

5 Summary

PCTL model checking

PCTL model checking problem

Input: a finite MDP $\mathcal{M} = (S, Act, \mathbf{P}, \iota_{\text{init}}, AP, L)$, 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 $Sat(\Phi)$.

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

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

Algorithm

Determine $x_s = Pr^{\max}(s \models \bigcirc \Phi)$ and return $Sat(\Psi) = \{s \in S \mid x_s \leq p\}$.

The case for minimal probabilities is similar and omitted here.

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

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

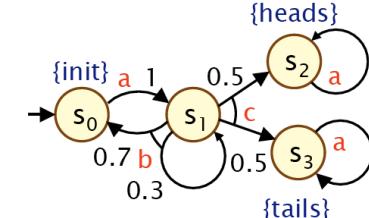
The same holds for strict upper bounds $< p$.

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

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

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:

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

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

Bounded until (1)

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

Lemma

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

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

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

The case for maximal probabilities is analogous.

Example

Bounded until (2)

Lemma

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

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

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

Algorithm

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

Until

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 $S_{=1} = Sat(\mathbb{P}_{=1}(\phi \cup \psi))$ by a graph analysis.
2. Determine $S_{=0} = Sat(\mathbb{P}_{=0}(\phi \cup \psi))$ 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.

Overview

1 PCTL Semantics

2 PCTL Model Checking

3 Complexity

4 Example: Dining Cryptographers Problem

5 Summary

Overview

1 PCTL Semantics

2 PCTL Model Checking

3 Complexity

4 Example: Dining Cryptographers Problem

5 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 $n_{\max} = \max\{n \mid \Psi_1 \leq^n \Psi_2 \text{ occurs in } \Phi\}$ with and $n_{\max} = 1$ if Φ does not contain a bounded until-operator.

Proof (sketch)

1. For each node in the parse tree, a model-checking is performed; this yields a linear complexity in $|\Phi|$.
2. The worst-case operator is until.
 - 2.1 Determining $S_{=0}$ and $S_{=1}$ can be done in linear time.
 - 2.2 Direct methods to solve linear programs with polynomial complexity

Overview

1 PCTL Semantics

2 PCTL Model Checking

3 Complexity

4 Example: Dining Cryptographers Problem

5 Summary

Dining cryptographers problem

[Chaum, 1988]

Problem statement

- ▶ Three cryptographers gather around a table for dinner.
- ▶ The waiter informs them that the meal has been paid by someone, who could be one of the cryptographers or their master.
- ▶ The cryptographers respect each other's right to make an anonymous payment, but want to find out whether the master paid or not.

Is it possible to obtain this information without revealing the identity of the cryptographer that 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).

Claim

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**).

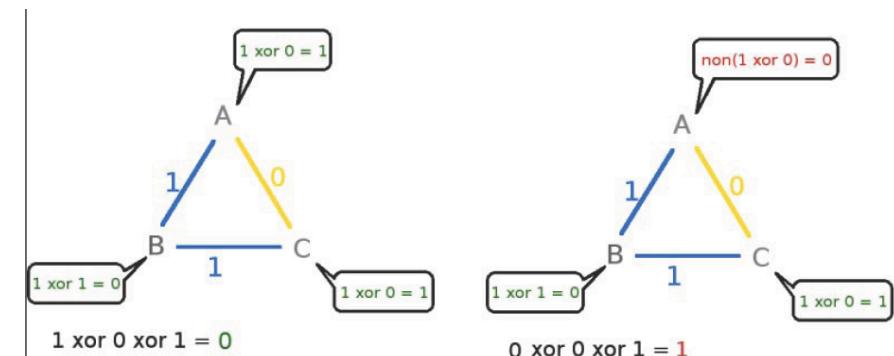
Caveat: if a cryptographer actually paid for the dinner, then it instead states the opposite (**disagree** if the coins are the same and **agree** if the coins are different).

Generalisation

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

Dining cryptographers problem



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

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

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

$\text{pay} \Rightarrow \mathbb{P}_{=1}(\Diamond(\text{done} \wedge \text{par} = N\%2)) \wedge \neg\text{pay} \Rightarrow \mathbb{P}_{=1}(\Diamond(\text{done} \wedge \text{par} \neq N\%2))$. 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.

Overview

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

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 which one 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.

Summary

- ▶ PCTL is a variant of CTL with operator $\mathbb{P}_J(\varphi)$.
- ▶ 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.

Next lecture: Monday, June 20.