Percentile Queries in Multi-Dimensional Markov Decision Processes

Mickael Randour[1] Jean-François Raskin[2] Ocan Sankur[2]

[1]LSV - CNRS & ENS Cachan, France

[2]ULB, Belgium

September 16, 2015 - *Highlights 2015, Prague*

3rd Highlights of Logic, Games and Automata

In a nutshell

Strategy synthesis for Markov Decision Processes (MDPs)

Finding **good** controllers for systems interacting with a *stochastic* environment.

In a nutshell

Strategy synthesis for Markov Decision Processes (MDPs)

Finding **good** controllers for systems interacting with a *stochastic* environment.

- Good? Performance evaluated through *payoff functions*.
- Usual problem is to optimize the *expected performance* or the *probability of achieving a given performance level*.
- Not sufficient for many practical applications.
  - ▷ Reason about *trade-offs* and *interplays*.
  - ▷ Several extensions, more expressive but also more complex...

In a nutshell

Strategy synthesis for Markov Decision Processes (MDPs)

Finding **good** controllers for systems interacting with a *stochastic* environment.

- Good? Performance evaluated through *payoff functions*.
- Usual problem is to optimize the *expected performance* or the *probability of achieving a given performance level*.
- Not sufficient for many practical applications.
  - Reason about *trade-offs* and *interplays*.
  - Several extensions, more expressive but also more complex...

Aim of this talk

Multi-constraint percentile queries: generalizes the problem to multiple dimensions, multiple constraints.

Advertisement

Full paper available on arXiv [RRS14]: abs/1410.4801

Featured in CAV'15 [RRS15a]

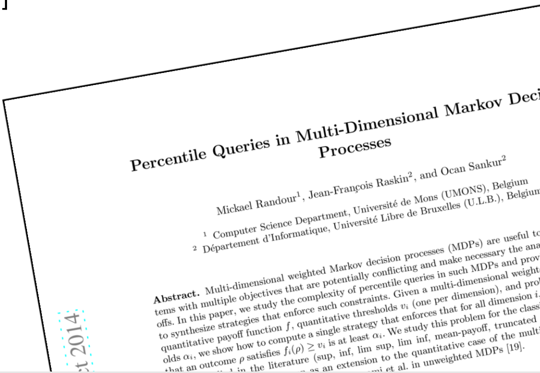
Percentile Queries in Multi-Dimensional Markov Deci
Processes

Mickael Randour[1], Jean-François Raskin[2], and Ocan Sankur[2]

[1] Computer Science Department, Université de Mons (UMONS), Belgium
[2] Département d'Informatique, Université Libre de Bruxelles (U.L.B.), Belgium

Abstract. Multi-dimensional weighted Markov decision processes (MDPs) are useful to
tems with multiple objectives that are potentially conflicting and make necessary the ana
offs. In this paper, we study the complexity of percentile queries in such MDPs and prov
to synthesize strategies that enforce such constraints. Given a multi-dimensional weight
quantitative payoff function f, quantitative thresholds v_i (one per dimension), and prob
olds α_i, we show how to compute a single strategy that enforces that for all dimension i,
that an outcome ρ satisfies $f_i(\rho) \geq v_i$ is at least α_i. We study this problem for the class
the literature (sup, inf, lim sup, lim inf, mean-payoff, truncated
as an extension to the quantitative case of the multi
et al. in unweighted MDPs [19].

Illustration: stochastic shortest path problem

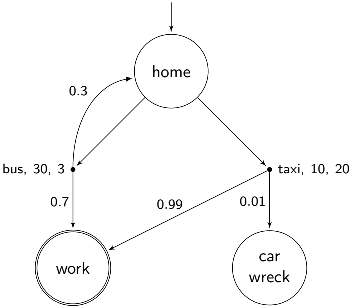

Two-dimensional weights on actions: *time* and *cost*.

Payoff: sum of weights up to work.

Often necessary to consider **trade-offs**: e.g., between the probability to reach work in due time and the risks of an expensive journey.

Illustration: stochastic shortest path problem

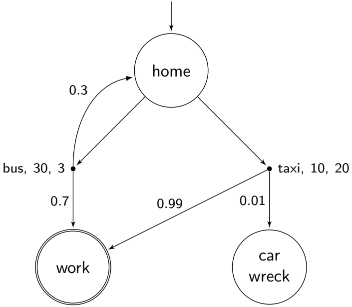

Classical problem considers only a **single percentile constraint**.

Single-constraint percentile problem

Given MDP M, initial state s_{init}, *one-dimension* payoff function f, value threshold $v \in \mathbb{Q}$, and probability threshold $\alpha \in [0, 1] \cap \mathbb{Q}$, decide if there exists a strategy σ such that $\mathbb{P}^{\sigma}_{M,s_{\text{init}}}\left[f \geq v\right] \geq \alpha$.

Illustration: stochastic shortest path problem

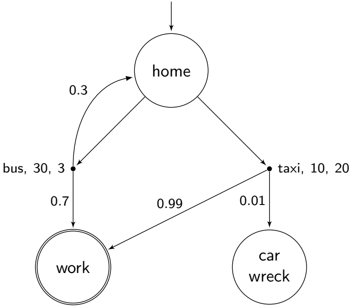

Classical problem considers only a **single percentile constraint**.

- **C1**: 80% of runs reach work in at most 40 minutes.
  - ▷ Taxi $\rightsquigarrow \leq 10$ minutes with probability $0.99 > 0.8$.

Illustration: stochastic shortest path problem

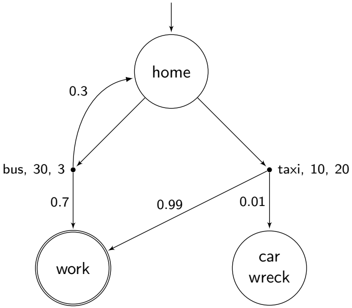

Classical problem considers only a **single percentile constraint**.

- **C1**: 80% of runs reach work in at most 40 minutes.
  - ▷ Taxi $\rightsquigarrow \leq 10$ minutes with probability $0.99 > 0.8$.
- **C2**: 50% of them cost at most 10\$ to reach work.
  - ▷ Bus $\rightsquigarrow \geq 70\%$ of the runs reach work for 3\$.

Illustration: stochastic shortest path problem

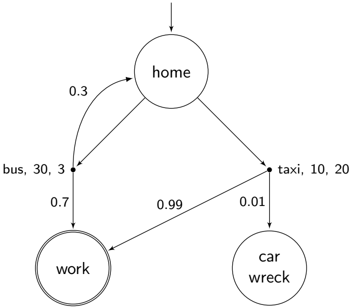

Classical problem considers only a **single percentile constraint**.

- **C1**: 80% of runs reach work in at most 40 minutes.
  - ▷ Taxi $\rightsquigarrow \leq 10$ minutes with probability $0.99 > 0.8$.
- **C2**: 50% of them cost at most 10$ to reach work.
  - ▷ Bus $\rightsquigarrow \geq 70\%$ of the runs reach work for 3$.

Taxi $\not\models$ C2, bus $\not\models$ C1. What if we want C1 $\wedge$ C2?

Illustration: stochastic shortest path problem

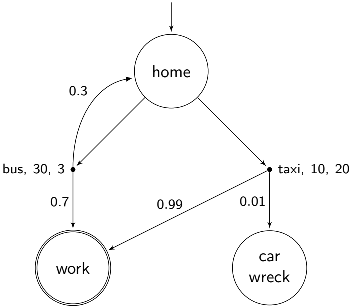

- **C1**: 80% of runs reach work in at most 40 minutes.
- **C2**: 50% of them cost at most 10$ to reach work.

Study of **multi-constraint percentile queries**.

- ▷ Sample strategy: bus once, then taxi. Requires *memory*.
- ▷ Another strategy: bus with probability 3/5, taxi with probability 2/5. Requires *randomness*.

Illustration: stochastic shortest path problem

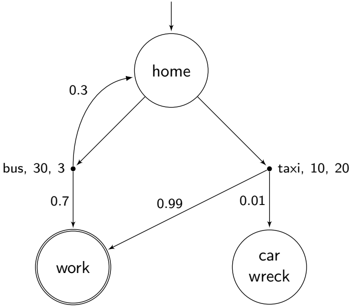

- **C1**: 80% of runs reach work in at most 40 minutes.
- **C2**: 50% of them cost at most 10$ to reach work.

Study of **multi-constraint percentile queries**.

In general, *both* memory *and* randomness are required.

$\neq$ classical problems (single constraint, expected value, etc)

Multi-constraint percentile problem

Multi-constraint percentile problem

Given d-dimensional MDP M, initial state s_{init}, payoff function f, and $q \in \mathbb{N}$ percentile constraints described by dimensions $l_i \in \{1, \ldots, d\}$, value thresholds $v_i \in \mathbb{Q}$ and probability thresholds $\alpha_i \in [0,1] \cap \mathbb{Q}$, where $i \in \{1, \ldots, q\}$, decide if there exists a strategy σ such that query $\mathcal{Q}$ holds, with

$$\mathcal{Q} := \bigwedge_{i=1}^{q} \mathbb{P}^{\sigma}_{M,s_{\text{init}}}\left[f_{l_i} \geq v_i\right] \geq \alpha_i.$$

Very general framework allowing for: multiple constraints related to $\neq$ or $=$ dimensions, $\neq$ value and probability thresholds.

- $\rightsquigarrow$ For SP, even $\neq$ targets for each constraint.
- $\rightsquigarrow$ Great flexibility in modeling applications.

Results overview (1/2)

- **Wide range of payoff functions**
  - multiple reachability,
  - mean-payoff ($\overline{\mathsf{MP}}$, $\underline{\mathsf{MP}}$),
  - discounted sum (DS).
  - $\inf$, $\sup$, $\liminf$, $\limsup$,
  - shortest path (SP),

Results overview (1/2)

- **Wide range of payoff functions**
  - multiple reachability,
  - $\inf$, $\sup$, $\liminf$, $\limsup$,
  - mean-payoff ($\overline{\mathsf{MP}}$, $\underline{\mathsf{MP}}$),
  - shortest path (SP),
  - discounted sum (DS).
- **Several variants**:
  - multi-dim. multi-constraint,
  - single-dim. multi-constraint,
  - single-constraint.

Results overview (1/2)

- **Wide range of payoff functions**
  - multiple reachability,
  - inf, sup, lim inf, lim sup,
  - mean-payoff ($\overline{\mathsf{MP}}$, $\underline{\mathsf{MP}}$),
  - shortest path (SP),
  - discounted sum (DS).
- **Several variants**:
  - multi-dim. multi-constraint,
  - single-dim. multi-constraint,
  - single-constraint.
- For each one:
  - algorithms,
  - lower bounds,
  - memory requirements.

$\leadsto$ **Complete picture** for this new framework.

Results overview (2/2)

	Single-constraint	Single-dim. Multi-constraint	Multi-dim. Multi-constraint
Reachability	P [Put94]	P(M)·E($\mathcal{Q}$) [EKVY08], PSPACE-h	—
$f \in \mathcal{F}$	P [CH09]	P	P(M)·E($\mathcal{Q}$) PSPACE-h.
$\overline{\mathsf{MP}}$	P [Put94]	P	P
$\underline{\mathsf{MP}}$	P [Put94]	P(M)·E($\mathcal{Q}$)	P(M)·E($\mathcal{Q}$)
SP	P(M)·P_{ps}($\mathcal{Q}$) [HK15] PSPACE-h. [HK15]	P(M)·P_{ps}($\mathcal{Q}$) (one target) PSPACE-h. [HK15]	P(M)·E($\mathcal{Q}$) PSPACE-h. [HK15]
ε-gap DS	$\mathsf{P}_{ps}(M, \mathcal{Q}, \varepsilon)$ NP-h.	$\mathsf{P}_{ps}(M, \varepsilon)$·E($\mathcal{Q}$) NP-h.	$\mathsf{P}_{ps}(M, \varepsilon)$·E($\mathcal{Q}$) PSPACE-h.

- ▷ $\mathcal{F} = \{\inf, \sup, \liminf, \limsup\}$
- ▷ M = model size, $\mathcal{Q}$ = query size
- ▷ P(x), E(x) and $\mathsf{P}_{ps}(x)$ resp. denote polynomial, exponential and pseudo-polynomial time in parameter x.

All results without reference are new.

Results overview (2/2)

	Single-constraint	Single-dim. Multi-constraint	Multi-dim. Multi-constraint
Reachability	P [Put94]	$\mathsf{P}(M)\cdot\mathsf{E}(\mathcal{Q})$ [EKVY08], PSPACE-h	—
$f \in \mathcal{F}$	P [CH09]	P	$\mathsf{P}(M)\cdot\mathsf{E}(\mathcal{Q})$ PSPACE-h.
$\overline{\mathsf{MP}}$	P [Put94]	P	P
$\underline{\mathsf{MP}}$	P [Put94]	$\mathsf{P}(M)\cdot\mathsf{E}(\mathcal{Q})$	$\mathsf{P}(M)\cdot\mathsf{E}(\mathcal{Q})$
SP	$\mathsf{P}(M)\cdot\mathsf{P}_{ps}(\mathcal{Q})$ [HK15] PSPACE-h. [HK15]	$\mathsf{P}(M)\cdot\mathsf{P}_{ps}(\mathcal{Q})$ (one target) PSPACE-h. [HK15]	$\mathsf{P}(M)\cdot\mathsf{E}(\mathcal{Q})$ PSPACE-h. [HK15]
ε-gap DS	$\mathsf{P}_{ps}(M, \mathcal{Q}, \varepsilon)$ NP-h.	$\mathsf{P}_{ps}(M, \varepsilon)\cdot\mathsf{E}(\mathcal{Q})$ NP-h.	$\mathsf{P}_{ps}(M, \varepsilon)\cdot\mathsf{E}(\mathcal{Q})$ PSPACE-h.

In most cases, only **polynomial in the model size**.

- ▷ In practice, the query size can often be bounded while the model can be very large.

Some related work

- **Same philosophy** (i.e., beyond uni-dimensional $\mathbb{E}$ or $\mathbb{P}$ maximization), $\neq$ approaches.
  - ▷ Beyond worst-case synthesis: $\mathbb{E}$ + worst-case [BFRR14b].
  - ▷ Survey of recent extensions in VMCAI'15 [RRS15b].

Some related work

- **Same philosophy** (i.e., beyond uni-dimensional $\mathbb{E}$ or $\mathbb{P}$ maximization), $\neq$ approaches.
  - ▷ Beyond worst-case synthesis: $\mathbb{E}$ + worst-case [BFRR14b].
  - ▷ Survey of recent extensions in VMCAI'15 [RRS15b].

- Multi-dim. MDPs: DS [CMH06], MP [BBC+14, FKR95].

Some related work

- **Same philosophy** (i.e., beyond uni-dimensional $\mathbb{E}$ or $\mathbb{P}$ maximization), $\neq$ approaches.
  - ▷ Beyond worst-case synthesis: $\mathbb{E}$ + worst-case [BFRR14b].
  - ▷ Survey of recent extensions in VMCAI'15 [RRS15b].
- Multi-dim. MDPs: DS [CMH06], MP [BBC+14, FKR95].
- Many related works for each particular payoff: MP [Put94], SP [UB13, HK15], DS [Whi93, WL99, BCF+13], etc.
  - ▷ All with a *single* constraint.

Some related work

- **Same philosophy** (i.e., beyond uni-dimensional $\mathbb{E}$ or $\mathbb{P}$ maximization), $\neq$ approaches.
  - ▷ Beyond worst-case synthesis: $\mathbb{E}$ + worst-case [BFRR14b].
  - ▷ Survey of recent extensions in VMCAI'15 [RRS15b].
- Multi-dim. MDPs: DS [CMH06], MP [BBC$^+$14, FKR95].
- Many related works for each particular payoff: MP [Put94], SP [UB13, HK15], DS [Whi93, WL99, BCF$^+$13], etc.
  - ▷ All with a *single* constraint.
- Multi-constraint percentile queries for LTL [EKVY08].
  - ▷ Closest to our work.
  - ▷ We use *multiple reachability*.

Some related work

- **Same philosophy** (i.e., beyond uni-dimensional $\mathbb{E}$ or $\mathbb{P}$ maximization), $\neq$ approaches.
  - ▷ Beyond worst-case synthesis: $\mathbb{E}$ + worst-case [BFRR14b].
  - ▷ Survey of recent extensions in VMCAI'15 [RRS15b].
- Multi-dim. MDPs: DS [CMH06], MP [BBC+14, FKR95].
- Many related works for each particular payoff: MP [Put94], SP [UB13, HK15], DS [Whi93, WL99, BCF+13], etc.
  - ▷ All with a *single* constraint.
- Multi-constraint percentile queries for LTL [EKVY08].
  - ▷ Closest to our work.
  - ▷ We use *multiple reachability*.
- Recent work on percentile queries + $\mathbb{E}$ for MP [CKK15].

Summary

	Single-constraint	Single-dim. Multi-constraint	Multi-dim. Multi-constraint
Reachability	P [Put94]	$\mathsf{P}(M)\cdot\mathsf{E}(\mathcal{Q})$ [EKVY08], PSPACE-h	—
$f \in \mathcal{F}$	P [CH09]	P	$\mathsf{P}(M)\cdot\mathsf{E}(\mathcal{Q})$ PSPACE-h.
$\overline{\mathsf{MP}}$	P [Put94]	P	P
$\underline{\mathsf{MP}}$	P [Put94]	$\mathsf{P}(M)\cdot\mathsf{E}(\mathcal{Q})$	$\mathsf{P}(M)\cdot\mathsf{E}(\mathcal{Q})$
SP	$\mathsf{P}(M)\cdot\mathsf{P}_{ps}(\mathcal{Q})$ [HK15] PSPACE-h. [HK15]	$\mathsf{P}(M)\cdot\mathsf{P}_{ps}(\mathcal{Q})$ (one target) PSPACE-h. [HK15]	$\mathsf{P}(M)\cdot\mathsf{E}(\mathcal{Q})$ PSPACE-h. [HK15]
ε-gap DS	$\mathsf{P}_{ps}(M, \mathcal{Q}, \varepsilon)$ NP-h.	$\mathsf{P}_{ps}(M, \varepsilon)\cdot\mathsf{E}(\mathcal{Q})$ NP-h.	$\mathsf{P}_{ps}(M, \varepsilon)\cdot\mathsf{E}(\mathcal{Q})$ PSPACE-h.

- $\mathcal{F} = \{\inf, \sup, \liminf, \limsup\}$
- M = model size, $\mathcal{Q}$ = query size
- $\mathsf{P}(x)$, $\mathsf{E}(x)$ and $\mathsf{P}_{ps}(x)$ resp. denote polynomial, exponential and pseudo-polynomial time in parameter x.

Thank you! Any question?

References I

Tomáš Brázdil, Václav Brozek, Krishnendu Chatterjee, Vojtech Forejt, and Antonín Kucera.
Markov decision processes with multiple long-run average objectives.
LMCS, 10(13):1–29, 2014.

Tomás Brázdil, Taolue Chen, Vojtech Forejt, Petr Novotný, and Aistis Simaitis.
Solvency Markov decision processes with interest.
In Proc. of FSTTCS, LIPIcs 24, pages 487–499. Schloss Dagstuhl - LZI, 2013.

Véronique Bruyère, Emmanuel Filiot, Mickael Randour, and Jean-François Raskin.
Expectations or guarantees? I want it all! A crossroad between games and MDPs.
In Proc. of SR, EPTCS 146, pages 1–8, 2014.

Véronique Bruyère, Emmanuel Filiot, Mickael Randour, and Jean-François Raskin.
Meet your expectations with guarantees: Beyond worst-case synthesis in quantitative games.
In Proc. of STACS, LIPIcs 25, pages 199–213. Schloss Dagstuhl - LZI, 2014.

Krishnendu Chatterjee, Laurent Doyen, Mickael Randour, and Jean-François Raskin.
Looking at mean-payoff and total-payoff through windows.
Inf. Comput., 242:25–52, 2015.

Krishnendu Chatterjee and Thomas A. Henzinger.
Probabilistic systems with limsup and liminf objectives.
In Margaret Archibald, Vasco Brattka, Valentin Goranko, and Benedikt Löwe, editors, Infinity in Logic and Computation, LNCS 5489, pages 32–45. Springer, 2009.

References II

Krishnendu Chatterjee, Zuzana Komárková, and Jan Kretínský.
Unifying two views on multiple mean-payoff objectives in Markov decision processes.
In Proc. of LICS. IEEE Computer Society, 2015.

Krishnendu Chatterjee, Rupak Majumdar, and Thomas A. Henzinger.
Markov decision processes with multiple objectives.
In Proc. of STACS, LNCS 3884, pages 325–336. Springer, 2006.

Krishnendu Chatterjee, Mickael Randour, and Jean-François Raskin.
Strategy synthesis for multi-dimensional quantitative objectives.
Acta Informatica, 51(3-4):129–163, 2014.

Kousha Etessami, Marta Z. Kwiatkowska, Moshe Y. Vardi, and Mihalis Yannakakis.
Multi-objective model checking of Markov decision processes.
LMCS, 4(4), 2008.

Jerzy A. Filar, Dmitry Krass, and Kirsten W. Ross.
Percentile performance criteria for limiting average Markov decision processes.
Automatic Control, IEEE Transactions on, 40(1):2–10, 1995.

Christoph Haase and Stefan Kiefer.
The odds of staying on budget.
In Proc. of ICALP, LNCS 9135, pages 234–246. Springer, 2015.

Martin L. Puterman.
Markov Decision Processes: Discrete Stochastic Dynamic Programming.
John Wiley & Sons, Inc., New York, NY, USA, 1st edition, 1994.

References III

Mickael Randour, Jean-François Raskin, and Ocan Sankur.
Percentile queries in multi-dimensional Markov decision processes.
CoRR, abs/1410.4801, 2014.

Mickael Randour, Jean-François Raskin, and Ocan Sankur.
Percentile queries in multi-dimensional Markov decision processes.
In Proc. of CAV, LNCS 9206, pages 123–139. Springer, 2015.

Mickael Randour, Jean-François Raskin, and Ocan Sankur.
Variations on the stochastic shortest path problem.
In Proc. of VMCAI, LNCS 8931, pages 1–18. Springer, 2015.

Michael Ummels and Christel Baier.
Computing quantiles in Markov reward models.
In Proc. of FOSSACS, LNCS 7794, pages 353–368. Springer, 2013.

Douglas J. White.
Minimizing a threshold probability in discounted Markov decision processes.
J. of Math. Anal. and App., 173(2):634 – 646, 1993.

Congbin Wu and Yuanlie Lin.
Minimizing risk models in Markov decision processes with policies depending on target values.
J. of Math. Anal. and App., 231(1), 1999.

Markov decision processes

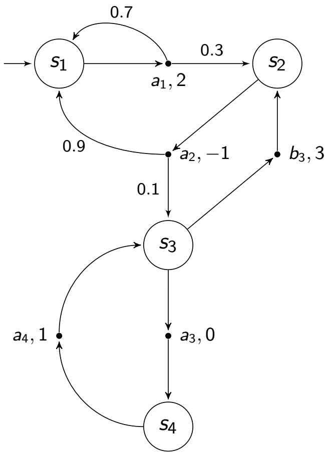

- **MDP** $M = (S, A, \delta, w)$
  - ▷ finite sets of states S and actions A
  - ▷ probabilistic transition $\delta \colon S \times A \to \mathcal{D}(S)$
  - ▷ weight function $w \colon A \to \mathbb{Z}^d$
- **Run** (or play): $\rho = s_1 a_1 \ldots a_{n-1} s_n \ldots$ such that $\delta(s_i, a_i, s_{i+1}) > 0$ for all $i \geq 1$
  - ▷ set of runs $\mathcal{R}(M)$
  - ▷ set of histories (finite runs) $\mathcal{H}(M)$
- **Strategy** $\sigma \colon \mathcal{H}(M) \to \mathcal{D}(A)$
  - ▷ $\forall\, h$ ending in s, $\text{Supp}(\sigma(h)) \in A(s)$

Markov decision processes

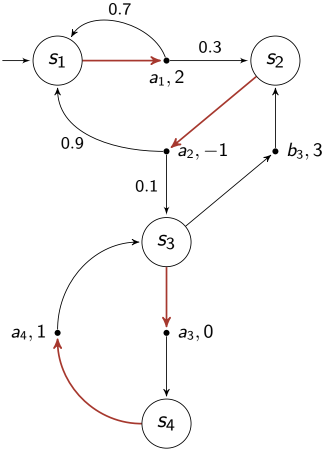

Sample *pure memoryless* strategy σ

Sample run $\rho = s_1 a_1 s_2 a_2 s_1 a_1 s_2 a_2 (s_3 a_3 s_4 a_4)^\omega$

Other possible run $\rho' = s_1 a_1 s_2 a_2 (s_3 a_3 s_4 a_4)^\omega$

- Strategies may use
  - finite or infinite **memory**
  - **randomness**
- **Payoff functions** map runs to numerical values
  - truncated sum up to $T = \{s_3\}$: $\mathsf{TS}^T(\rho) = 2$, $\mathsf{TS}^T(\rho') = 1$
  - mean-payoff: $\underline{\mathsf{MP}}(\rho) = \underline{\mathsf{MP}}(\rho') = 1/2$
  - many more

Markov chains

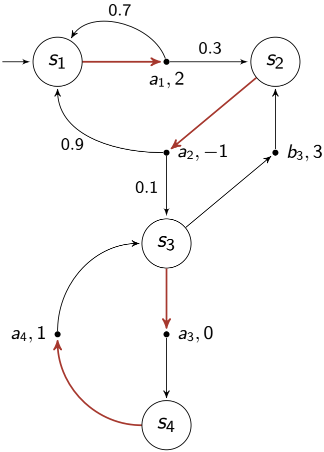

Once initial state s_{init} and strategy σ fixed, fully stochastic process

$\rightsquigarrow$ **Markov chain** (**MC**)

Markov chains

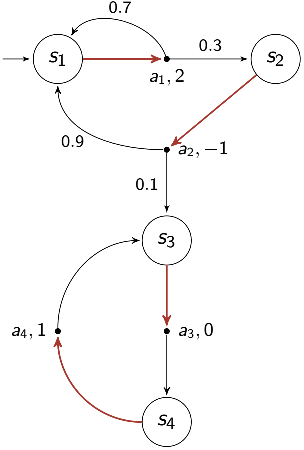

Once initial state s_{init} and strategy σ fixed, fully stochastic process

$\rightsquigarrow$ **Markov chain** (**MC**)

State space = product of the MDP and the memory of σ

Markov chains

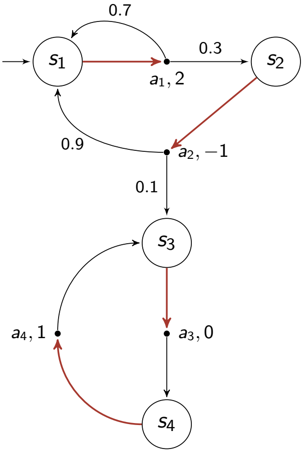

Once initial state s_{init} and strategy σ fixed, fully stochastic process

$\rightsquigarrow$ **Markov chain** (**MC**)

State space = product of the MDP and the memory of σ

- Event $\mathcal{E} \subseteq \mathcal{R}(M)$
  - ▷ probability $\mathbb{P}^{\sigma}_{M,s_{\mathrm{init}}}(\mathcal{E})$
- Measurable $f \colon \mathcal{R}(M) \to (\mathbb{R} \cup \{-\infty, \infty\})^d$
  - ▷ expected value $\mathbb{E}^{\sigma}_{M,s_{\mathrm{init}}}(f)$

Results overview: sketches

	Single-constraint	Single-dim. Multi-constraint	Multi-dim. Multi-constraint
Reachability	P [Put94]	$P(M) \cdot E(\mathcal{Q})$ [EKVY08], PSPACE-h	—
$f \in \mathcal{F}$	P [CH09]	P	$P(M) \cdot E(\mathcal{Q})$ PSPACE-h.
$\overline{\mathsf{MP}}$	P [Put94]	P	P
$\underline{\mathsf{MP}}$	P [Put94]	$P(M) \cdot E(\mathcal{Q})$	$P(M) \cdot E(\mathcal{Q})$
SP	$P(M) \cdot P_{ps}(\mathcal{Q})$ [HK15] PSPACE-h. [HK15]	$P(M) \cdot P_{ps}(\mathcal{Q})$ (one target) PSPACE-h. [HK15]	$P(M) \cdot E(\mathcal{Q})$ PSPACE-h. [HK15]
ε-gap DS	$P_{ps}(M, \mathcal{Q}, \varepsilon)$ NP-h.	$P_{ps}(M, \varepsilon) \cdot E(\mathcal{Q})$ NP-h.	$P_{ps}(M, \varepsilon) \cdot E(\mathcal{Q})$ PSPACE-h.

- ▷ $\mathcal{F} = \{\inf, \sup, \liminf, \limsup\}$
- ▷ M = model size, $\mathcal{Q}$ = query size
- ▷ $P(x)$, $E(x)$ and $P_{ps}(x)$ resp. denote polynomial, exponential and pseudo-polynomial time in parameter x.

All results without reference are new.

Results overview: sketches

	Single-constraint	Single-dim. Multi-constraint	Multi-dim. Multi-constraint
Reachability	P [Put94]	P(M)·E($\mathcal{Q}$) [EKVY08], PSPACE-h	—
$f \in \mathcal{F}$	P [CH09]	P	P(M)·E($\mathcal{Q}$) PSPACE-h.
$\overline{\mathsf{MP}}$	P [Put94]	P	P
$\underline{\mathsf{MP}}$	P [Put94]	P(M)·E($\mathcal{Q}$)	P(M)·E($\mathcal{Q}$)
SP	P(M)·P_{ps}($\mathcal{Q}$) [HK15] PSPACE-h. [HK15]	P(M)·P_{ps}($\mathcal{Q}$) (one target) PSPACE-h. [HK15]	P(M)·E($\mathcal{Q}$) PSPACE-h. [HK15]
ε-gap DS	$\mathsf{P}_{ps}(M, \mathcal{Q}, \varepsilon)$ NP-h.	$\mathsf{P}_{ps}(M, \varepsilon)$·E($\mathcal{Q}$) NP-h.	$\mathsf{P}_{ps}(M, \varepsilon)$·E($\mathcal{Q}$) PSPACE-h.

In most cases, only **polynomial in the model size**.

- ▷ In practice, the query size can often be bounded while the model can be very large.

Results overview: sketches

	Single-constraint	Single-dim. Multi-constraint	Multi-dim. Multi-constraint
Reachability	P [Put94]	P(M)·E($\mathcal{Q}$) [EKVY08], PSPACE-h	—
$f \in \mathcal{F}$	P [CH09]	P	P(M)·E($\mathcal{Q}$) PSPACE-h.
$\overline{\mathsf{MP}}$	P [Put94]	P	P
$\underline{\mathsf{MP}}$	P [Put94]	P(M)·E($\mathcal{Q}$)	P(M)·E($\mathcal{Q}$)
SP	P(M)·P$_{ps}$($\mathcal{Q}$) [HK15] PSPACE-h. [HK15]	P(M)·P$_{ps}$($\mathcal{Q}$) (one target) PSPACE-h. [HK15]	P(M)·E($\mathcal{Q}$) PSPACE-h. [HK15]
ε-gap DS	P$_{ps}$($M, \mathcal{Q}, \varepsilon$) NP-h.	P$_{ps}$(M, ε)·E($\mathcal{Q}$) NP-h.	P$_{ps}$(M, ε)·E($\mathcal{Q}$) PSPACE-h.

No time to discuss every result!

Results overview: sketches

	Single-constraint	Single-dim. Multi-constraint	Multi-dim. Multi-constraint
Reachability	P [Put94]	$\mathsf{P}(M)\cdot\mathsf{E}(\mathcal{Q})$ [EKVY08], PSPACE-h	—
$f \in \mathcal{F}$	P [CH09]	P	$\mathsf{P}(M)\cdot\mathsf{E}(\mathcal{Q})$ PSPACE-h.
$\overline{\mathsf{MP}}$	P [Put94]	P	P
$\underline{\mathsf{MP}}$	P [Put94]	$\mathsf{P}(M)\cdot\mathsf{E}(\mathcal{Q})$	$\mathsf{P}(M)\cdot\mathsf{E}(\mathcal{Q})$
SP	$\mathsf{P}(M)\cdot\mathsf{P}_{ps}(\mathcal{Q})$ [HK15] PSPACE-h. [HK15]	$\mathsf{P}(M)\cdot\mathsf{P}_{ps}(\mathcal{Q})$ (one target) PSPACE-h. [HK15]	$\mathsf{P}(M)\cdot\mathsf{E}(\mathcal{Q})$ PSPACE-h. [HK15]
ε-gap DS	$\mathsf{P}_{ps}(M, \mathcal{Q}, \varepsilon)$ NP-h.	$\mathsf{P}_{ps}(M, \varepsilon)\cdot\mathsf{E}(\mathcal{Q})$ NP-h.	$\mathsf{P}_{ps}(M, \varepsilon)\cdot\mathsf{E}(\mathcal{Q})$ PSPACE-h.

Four groups of results

1. **Reachability**. Algorithm based on multi-objective linear programming (LP) in [EKVY08]. We refine the complexity analysis, provide LBs and tractable subclasses.
   - ▷ **Useful tool** for many payoff functions!

Results overview: sketches

	Single-constraint	Single-dim. Multi-constraint	Multi-dim. Multi-constraint
Reachability	P [Put94]	P(M)·E($\mathcal{Q}$) [EKVY08], PSPACE-h	—
$f \in \mathcal{F}$	P [CH09]	P	P(M)·E($\mathcal{Q}$) PSPACE-h.
$\overline{\mathsf{MP}}$	P [Put94]	P	P
$\underline{\mathsf{MP}}$	P [Put94]	P(M)·E($\mathcal{Q}$)	P(M)·E($\mathcal{Q}$)
SP	P(M)·P$_{ps}$($\mathcal{Q}$) [HK15] PSPACE-h. [HK15]	P(M)·P$_{ps}$($\mathcal{Q}$) (one target) PSPACE-h. [HK15]	P(M)·E($\mathcal{Q}$) PSPACE-h. [HK15]
ε-gap DS	P$_{ps}$($M, \mathcal{Q}, \varepsilon$) NP-h.	P$_{ps}$(M, ε)·E($\mathcal{Q}$) NP-h.	P$_{ps}$(M, ε)·E($\mathcal{Q}$) PSPACE-h.

Four groups of results

2 $\mathcal{F}$ **and** $\overline{\mathsf{MP}}$. Easiest cases.

- inf and sup: reduction to *multiple reachability*.
- lim inf, lim sup and $\overline{\mathsf{MP}}$: *maximal end-component* (MEC) decomposition + reduction to multiple reachability.

Results overview: sketches

	Single-constraint	Single-dim. Multi-constraint	Multi-dim. Multi-constraint
Reachability	P [Put94]	$\mathsf{P}(\mathcal{M}){\cdot}\mathsf{E}(\mathcal{Q})$ [EKVY08], PSPACE-h	—
$f \in \mathcal{F}$	P [CH09]	P	$\mathsf{P}(\mathcal{M}){\cdot}\mathsf{E}(\mathcal{Q})$ PSPACE-h.
$\overline{\mathsf{MP}}$	P [Put94]	P	P
$\underline{\mathsf{MP}}$	P [Put94]	$\mathsf{P}(\mathcal{M}){\cdot}\mathsf{E}(\mathcal{Q})$	$\mathsf{P}(\mathcal{M}){\cdot}\mathsf{E}(\mathcal{Q})$
SP	$\mathsf{P}(\mathcal{M}){\cdot}\mathsf{P}_{ps}(\mathcal{Q})$ [HK15] PSPACE-h. [HK15]	$\mathsf{P}(\mathcal{M}){\cdot}\mathsf{P}_{ps}(\mathcal{Q})$ (one target) PSPACE-h. [HK15]	$\mathsf{P}(\mathcal{M}){\cdot}\mathsf{E}(\mathcal{Q})$ PSPACE-h. [HK15]
ε-gap DS	$\mathsf{P}_{ps}(\mathcal{M}, \mathcal{Q}, \varepsilon)$ NP-h.	$\mathsf{P}_{ps}(\mathcal{M}, \varepsilon){\cdot}\mathsf{E}(\mathcal{Q})$ NP-h.	$\mathsf{P}_{ps}(\mathcal{M}, \varepsilon){\cdot}\mathsf{E}(\mathcal{Q})$ PSPACE-h.

Four groups of results

3. $\underline{\mathsf{MP}}$. Technically involved.
   - ▷ Inside MECs: (a) strategies satisfying *maximal subsets of constraints*, (b) combine them linearly.
   - ▷ Overall: write an LP combining multiple reachability toward MECs and those linear combinations equations.

Results overview: sketches

	Single-constraint	Single-dim. Multi-constraint	Multi-dim. Multi-constraint
Reachability	P [Put94]	$\mathsf{P}(M)\cdot\mathsf{E}(\mathcal{Q})$ [EKVY08], PSPACE-h	—
$f \in \mathcal{F}$	P [CH09]	P	$\mathsf{P}(M)\cdot\mathsf{E}(\mathcal{Q})$ PSPACE-h.
$\overline{\mathsf{MP}}$	P [Put94]	P	P
$\underline{\mathsf{MP}}$	P [Put94]	$\mathsf{P}(M)\cdot\mathsf{E}(\mathcal{Q})$	$\mathsf{P}(M)\cdot\mathsf{E}(\mathcal{Q})$
SP	$\mathsf{P}(M)\cdot\mathsf{P}_{ps}(\mathcal{Q})$ [HK15] PSPACE-h. [HK15]	$\mathsf{P}(M)\cdot\mathsf{P}_{ps}(\mathcal{Q})$ (one target) PSPACE-h. [HK15]	$\mathsf{P}(M)\cdot\mathsf{E}(\mathcal{Q})$ PSPACE-h. [HK15]
ε-gap DS	$\mathsf{P}_{ps}(M, \mathcal{Q}, \varepsilon)$ NP-h.	$\mathsf{P}_{ps}(M, \varepsilon)\cdot\mathsf{E}(\mathcal{Q})$ NP-h.	$\mathsf{P}_{ps}(M, \varepsilon)\cdot\mathsf{E}(\mathcal{Q})$ PSPACE-h.

Four groups of results

4. **SP and DS**. Based on *unfoldings* and multiple reachability.
   - ▷ Need finite and bounded unfoldings.
   - ▷ For SP, we bound the size of the unfolding by *node merging*.
   - ▷ For DS, we can only *approximate* the answer in general. Need to analyze the cumulative error due to necessary *roundings*.