Quantifying Risk in Probabilistic Systems

Jan Křetínský Tobias Meggendorfer

Fakultät für Informatik, Technische Universität München

The Goal

- Establish sensible risk measure for probabilistic systems,
- define related decision and optimization problems,
- derive their theoretical complexity bound, and
- implement practical verification / synthesis procedures to
- obtain optimal, risk averse controllers for safety-critical systems.

Motivation – Controlling a power plant

- Risk assessment / aversion imperative to safety-critical systems
- We want both good average performance and little risk of failures
- Maximizing expectation not good enough
  - Outages may occur with little probability ⇒ little impact on expectation
  - High risk, high reward behaviour incentivized
- Completely avoiding bad behaviour (worst-case) neither
  - Any reasonable plant model has *some* probability of failure
  - Only "safe" strategy: Don't produce any power
- Other typical objectives in verification also ill-suited
  - Variance: Does not distinguish between "good" and "bad" deviations
  - Value-at-Risk: "Seductive, but dangerous" – sensitive to perturbations
- Thus: Need a measure of risk suitable for a prob. context

The Conditional Value-at-Risk

- Established approach in other fields (OR / Finance)
- A.k.a.: Expected tail loss, expected shortfall, average value-at-risk

CVaR

Let X be a random variable and $p \in (0, 1)$. Then

$$\mathrm{VaR}_p(X) := \sup\{r \in \mathbb{R} \mid \mathbb{P}[X \leq r] \leq p\}$$

Let $v = \mathrm{VaR}_p(X)$ and $p' = \mathbb{P}[X < v]$.

$$\mathrm{CVaR}_p(X) := \frac{1}{p}(p' \cdot \mathbb{E}[X \mid X < v] + (p - p') \cdot v),$$

- $\mathrm{VaR}_p(X)$ (the *Value-at-Risk*): "What is a reasonable bad case?"
- $\mathrm{CVaR}_p(X)$: "What happens in the average bad case?"

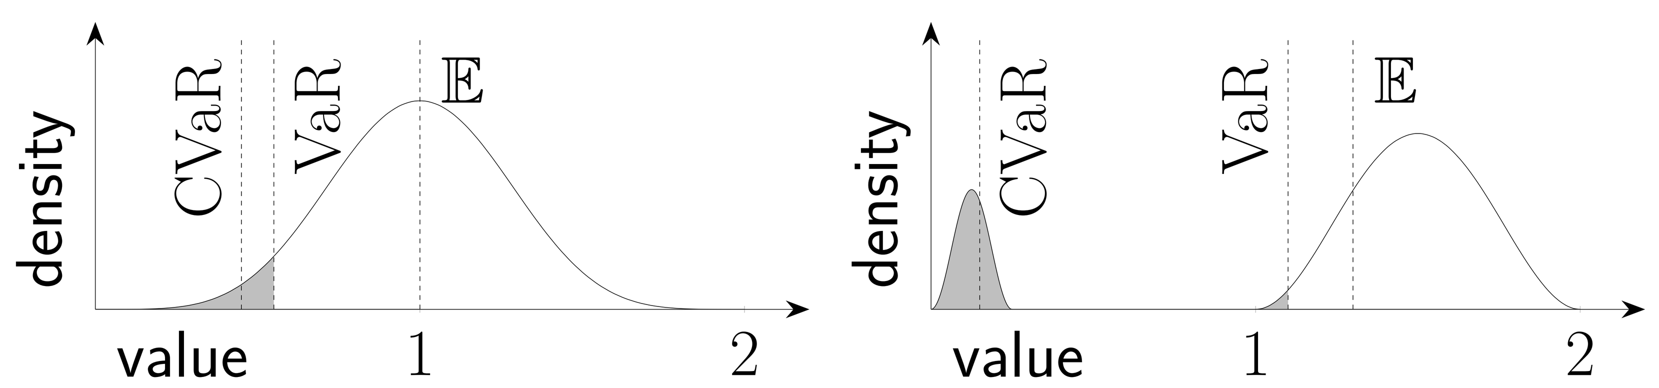

Some properties of CVaR

- Interpolation between worst-case ($p \to 0$) and expectation ($p \to 1$)
- Robust to changes in X and p caused by, e.g., modelling errors
- *Coherent* risk measure (established term in finance)

Our Contributions(1)

- Introduce CVaR both generally and in the context of MDP
- Define various related decision problems
- Derive theoretical (LP-based) decision procedures and tight complexity bounds

Markov Decision Processes (MDP)

- Standard Model for single actor in random environment
- Comprises: States, Actions, Transition Probabilities, and Rewards

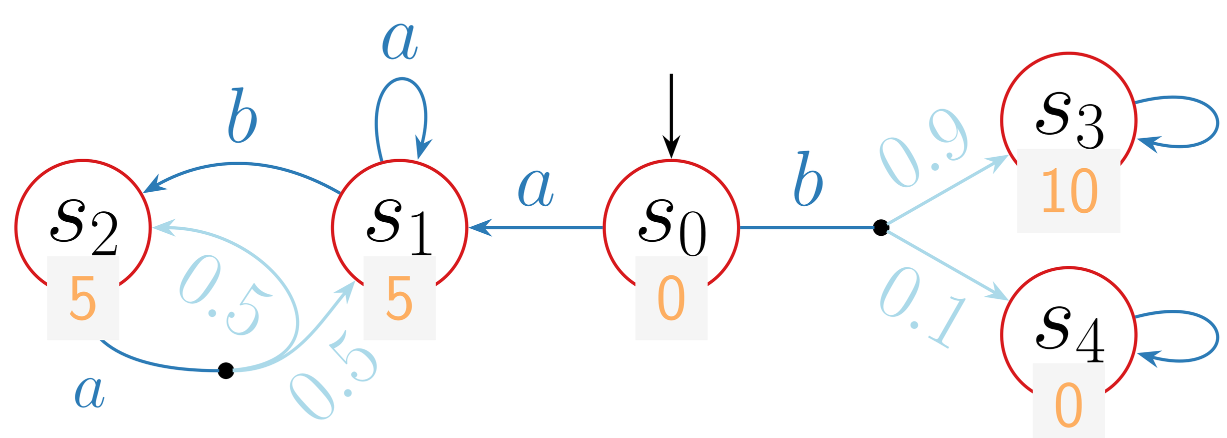

The Objectives

- Weighted reachability: Obtain first visited non-zero reward. Example: Prioritized goals.
- Mean payoff: Reward obtained "on average" per step. Example: Average energy production.

The Decision Problems

Given

MDP $\mathcal{M}$, dimensions $d \in \mathbb{N}^+$, reward function $\vec{r} : S \to \mathbb{Q}^d$, reward interpretation $\mathbf{rew}$: Run $\to \mathbb{Q}$, thresholds $\vec{e}, \vec{v}, \vec{c} \in \mathbb{Q}^d$, and probabilities $\vec{p}, \vec{q} \in [0, 1)^d$,

is there a strategy s.t.

$\mathbb{E}[\mathbf{rew}_i] \geq e_i$, $\mathrm{VaR}_{p_i}(\mathbf{rew}_i) \geq v_i$, and $\mathrm{CVaR}_{q_i}(\mathbf{rew}_i) \geq c_i$ $\forall i$?

Results

- Single dim. ($d = 1$): Everything in P; simple opt. strategies
- Weighted reach.: NP-complete (in d); simple strategies
- Mean-payoff: NP-hard, EXPSPACE (in d); complex strategies. Conjecture: NP-complete

Overall: Synthesizing risk averse controllers tractable for MDP

Future Work

- Extend to richer systems, e.g., 2-player stochastic games, and more objectives, e.g., bounded-horizon / discounted properties
- Practical implementation and (approximative) optimization
- On-the-fly reformulation
- Close complexity gap for mean-payoff
- Application and interpretation in real-life scenarios

[1] J. Křetínský and T. Meggendorfer. Conditional value-at-risk for reachability and mean payoff in markov decision processes. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science*, LICS '18, pages 609–618, New York, NY, USA, 2018. ACM. ISBN 978-1-4503-5583-4. doi:10.1145/3209108.3209176.