Controlling probabilistic systems under partial observation a verification perspective

Nathalie Bertrand, Inria Rennes, France

Journées du GDR IM
15 mars 2017, Montpellier

Partially observable probabilistic systems

Why probabilities?

randomized algorithms, unpredictable behaviours, abstraction of non-determinism

Why partial observation?

abstraction of large systems, security concerns

Partially observable probabilistic systems

Why probabilities?

randomized algorithms, unpredictable behaviours, abstraction of non-determinism

Why partial observation?

abstraction of large systems, security concerns

this talk: known automaton-like model

- language-theoretic questions: languages defined by prob. automata
- monitoring issues: fault diagnosis, supervision, etc.
- **control problems**: optimization for a given objective

Outline

Probabilistic automata

Partially observable MDP

Discussion

Motivating example for probabilistic automata (PA)

Planning holidays *in advance*:

1. choose an airline type (`lowcost`/`highcost`);
2. book accommodation (`internet`/`phone`);
3. choose tour (`seeall`/`missnothing`).

each action fails with some probability

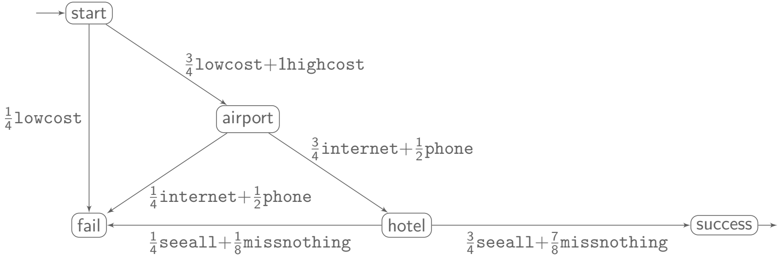

success probability of plan $\texttt{lowcost} \cdot \texttt{internet} \cdot \texttt{seeall}$ is $\frac{27}{64}$.

Control strategies in PA

Strategies are words
what is the probability to reach a final state after word w?

The acceptance probability of $w = a_1 \ldots a_n$ by $\mathcal{A}$ is:

$$\mathbf{Pr}_{\mathcal{A}}(w) = \sum_{q \in Q} \pi_{\mathbf{0}}[q] \sum_{q' \in F} \left(\prod_{i=1}^{n} \mathbf{P}_{a_i} \right) [q, q'] = \pi_0 \mathbf{P}_w \mathbf{1}_F^T$$

Control strategies in PA

Strategies are words
what is the probability to reach a final state after word w?

The acceptance probability of $w = a_1 \dots a_n$ by $\mathcal{A}$ is:

$$\mathbf{Pr}_{\mathcal{A}}(w) = \sum_{q \in Q} \pi_{\mathbf{0}}[q] \sum_{q' \in F} \left(\prod_{i=1}^{n} \mathbf{P}_{a_i} \right) [q, q'] = \pi_0 \mathbf{P}_w \mathbf{1}_F^T$$

Optimal strategies may not exist

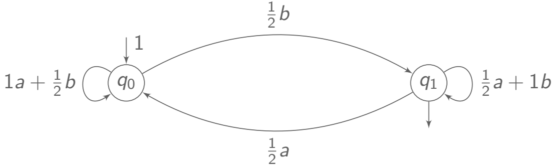

$$\mathbf{Pr}_{\mathcal{A}}(a_1 \dots a_n) = \sum_{i=1}^{n} 2^{i-n-1} \cdot \mathbf{1}_{a_i = b}$$

Control strategies in PA

Strategies are words

what is the probability to reach a final state after word w?

The acceptance probability of $w = a_1 \dots a_n$ by $\mathcal{A}$ is:

$$\mathbf{Pr}_{\mathcal{A}}(w) = \sum_{q \in Q} \pi_{\mathbf{0}}[q] \sum_{q' \in F} \left(\prod_{i=1}^{n} \mathbf{P}_{a_i} \right) [q, q'] = \pi_0 \mathbf{P}_w \mathbf{1}_F^T$$

Optimal strategies may not exist

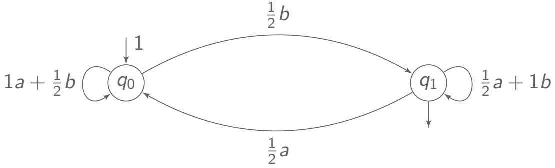

$$\mathbf{Pr}_{\mathcal{A}}(a_1 \dots a_n) = \sum_{i=1}^{n} 2^{i-n-1} \cdot \mathbf{1}_{a_i = b}$$

$\longrightarrow$ Find good enough strategies, *i.e.* that guarantee a given probability

Existence of good-enough strategies

$$L_{\bowtie\theta}(\mathcal{A}) = \{w \in A^* \mid \mathbf{Pr}_{\mathcal{A}}(w) \bowtie \theta\}$$

The problem, given a PA $\mathcal{A}$ of telling whether $L_{\geq\frac{1}{2}}(\mathcal{A}) \neq \emptyset$ is undecidable. Paz'71

Existence of good-enough strategies

$$L_{\bowtie\theta}(\mathcal{A}) = \{w \in A^* \mid \mathbf{Pr}_{\mathcal{A}}(w) \bowtie \theta\}$$

The problem, given a PA $\mathcal{A}$ of telling whether $L_{\geq\frac{1}{2}}(\mathcal{A}) \neq \emptyset$ is undecidable. Paz'71

Undecidability is robust

refined emptiness assuming that for $\epsilon > 0$ either $\exists w\ \mathbf{Pr}_{\mathcal{A}}(w) \geq 1 - \varepsilon$ or $\forall w\ \mathbf{Pr}_{\mathcal{A}}(w) < \varepsilon$, decide which is the case Condon *et al.*'03

value one problem does there exist $(w_n)_{n\in\mathbb{N}}$ such that $\limsup_n \mathbf{Pr}_{\mathcal{A}}(w_n) = 1$? Gimbert and Oualhadj'10

parametric probability values does there exist a valuation of probabilities such that $\mathcal{A}$ has value one? Fijalkow *et al.*'14

Anything decidable?

Almost-sure language: $L_{=1}(\mathcal{A})$

Emptiness of almost-sure language is PSPACE-complete.

equivalent to universality problem for NFA

Anything decidable?

Almost-sure language: $L_{=1}(\mathcal{A})$

Emptiness of almost-sure language is PSPACE-complete.

equivalent to universality problem for NFA

Quantitative language equivalence
Input: $\mathcal{A}$ and $\mathcal{A}'$ PA
Output: yes iff $\forall w \in A^* \ \mathbf{Pr}_{\mathcal{A}}(w) = \mathbf{Pr}_{\mathcal{A}'}(w)$

Quantitative language equivalence is decidable in PTIME.

Schützenberger'61, Tzeng'92

linear algebra argument
polynomial bound on length of counterexample to equivalence

Recap on Probabilistic Automata

Partial observation: the plan must be decided in advance!

Recap on Probabilistic Automata

Partial observation: the plan must be decided in advance!

- model of system is known
- the effect of a plan can be computed:
  after word w, probability distribution over states
- yet most optimization problems are undecidable

Recap on Probabilistic Automata

Partial observation: the plan must be decided in advance!

- model of system is known
- the effect of a plan can be computed:
  after word w, probability distribution over states
- yet most optimization problems are undecidable

What if the system provides feedback, and we can update the plan?
partially observable Markov decision processes

Outline

Example of partially observable MDP (POMDP)

McCallum maze: robot with limited sensor abilities, and imperfect moves

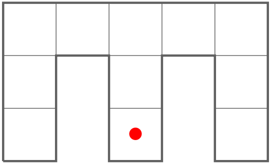

Example of partially observable MDP (POMDP)

McCallum maze: robot with limited sensor abilities, and imperfect moves

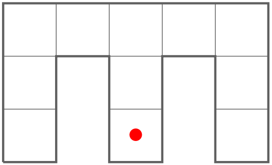

- random initial position
- robot only sees walls surrounding it, not the precise cell

  observations $\Omega = \{\{L, U\}, \{U, D\}, \{U, R\}, \{L, D, R\} \cdots\}$
- actions $A = \{N, W, S, E\}$ are not implemented accurately

  action N leads to north with probability $\frac{2}{3}$ and others with $\frac{1}{3}$

Example of partially observable MDP (POMDP)

McCallum maze: robot with limited sensor abilities, and imperfect moves

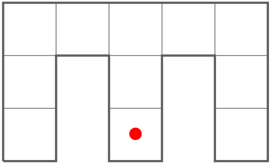

- random initial position
- robot only sees walls surrounding it, not the precise cell
  observations $\Omega = \{\{L, U\}, \{U, D\}, \{U, R\}, \{L, D, R\} \cdots\}$
- actions $A = \{N, W, S, E\}$ are not implemented accurately
  action N leads to north with probability $\frac{2}{3}$ and others with $\frac{1}{3}$

Reachability objective: move to target cell ●
Optimization: minimum expected time

Strategies

Strategy: maps *history* $\rho \in (A\Omega)^*$ with distribution over actions;

$$\nu : (A\Omega)^* \to \mathsf{Dist}(A)$$

$\nu(\rho, a)$: probability that a is chosen given history ρ

Strategies

Strategy: maps *history* $\rho \in (A\Omega)^*$ with distribution over actions;

$$\nu : (A\Omega)^* \to \mathsf{Dist}(A)$$

$\nu(\rho, a)$: probability that a is chosen given history ρ

- **pure** strategy: all distributions are Dirac
- **belief-based** strategy: based on set of current possible states

Strategies

Strategy: maps *history* $\rho \in (A\Omega)^*$ with distribution over actions;

$$\nu : (A\Omega)^* \to \mathsf{Dist}(A)$$

$\nu(\rho, a)$: probability that a is chosen given history ρ

- **pure** strategy: all distributions are Dirac
- **belief-based** strategy: based on set of current possible states

word in PA $\Longleftrightarrow$ pure strategy in POMDP with $|\Omega| = 1$

Consequence: all hardness results lift from PA to POMDP

Infinite horizon objectives

Objectives Reachability F visited at least once:

$$\Diamond F = \{q_0 q_1 q_2 \cdots \in S^\omega \mid \exists n,\ q_n \in F\}$$

Safety always stay in F:

$$\Box F = \{q_0 q_1 q_2 \cdots \in S^\omega \mid \forall n,\ q_n \in F\}$$

Büchi F visited an infinite number of times:

$$\Box\Diamond F = \{q_0 q_1 q_2 \cdots \in S^\omega \mid \forall m\, \exists n \geq m,\ q_n \in F\}$$

Goal: For φ an objective, evaluate $\sup_\nu \mathbb{P}^\nu(\mathcal{M} \models \varphi)$.

Infinite horizon objectives

Objectives Reachability F visited at least once:

$$\Diamond F = \{q_0 q_1 q_2 \cdots \in S^\omega \mid \exists n,\ q_n \in F\}$$

Safety always stay in F:

$$\Box F = \{q_0 q_1 q_2 \cdots \in S^\omega \mid \forall n,\ q_n \in F\}$$

Büchi F visited an infinite number of times:

$$\Box\Diamond F = \{q_0 q_1 q_2 \cdots \in S^\omega \mid \forall m\, \exists n \geq m,\ q_n \in F\}$$

Goal: For φ an objective, evaluate $\sup_\nu \mathbb{P}^\nu(\mathcal{M} \models \varphi)$.

Pure strategies suffice!

For every strategy ν, there exists a pure strategy ν' such that $\mathbb{P}^\nu(\mathcal{M} \models \varphi) \leq \mathbb{P}^{\nu'}(\mathcal{M} \models \varphi)$. Chatterjee *et al.*'15

Undecidability results

Undecidability of qualitative objectives...

... beyond the ones already mentioned for PA

Undecidability results

Undecidability of qualitative objectives...

... beyond the ones already mentioned for PA

positive repeated reachability does there exist ν such that
$\mathbb{P}^{\nu}(\mathcal{M} \models \Box\Diamond F) > 0$? Baier *et al.*'08

combined objectives does there exist ν such that Bertrand *et al.*'14
$\mathbb{P}^{\nu}(\mathcal{M} \models \Box\Diamond F_1) = 1$ and $\mathbb{P}^{\nu}(\mathcal{M} \models \Box F_2) > 0$?

Undecidability results

Undecidability of qualitative objectives...

... beyond the ones already mentioned for PA

positive repeated reachability does there exist ν such that

$$\mathbb{P}^{\nu}(\mathcal{M} \models \square\lozenge F) > 0?$$

Baier *et al.*'08

combined objectives does there exist ν such that

Bertrand *et al.*'14

$$\mathbb{P}^{\nu}(\mathcal{M} \models \square\lozenge F_1) = 1 \text{ and } \mathbb{P}^{\nu}(\mathcal{M} \models \square F_2) > 0?$$

Proof of first statement: reduction from the value one problem for PA

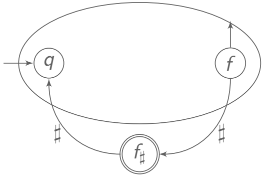

pure strategies in $\mathcal{M}$:

$\nu_{\mathbf{w}} = w_1 \sharp\sharp w_2 \sharp\sharp w_3 \cdots$

$$\mathrm{val}(\mathcal{A}) = 1 \iff \exists (w_i)_{i\in\mathbb{N}} \prod_i \mathbb{P}_{\mathcal{A}}(w_i) > 0$$

$$\iff \exists \nu_{\mathbf{w}}\ \mathbb{P}^{\nu_{\mathbf{w}}}(\mathcal{M} \models \square\lozenge f_\sharp) > 0$$

Decidability results

Good news: decidable problems for PA remain decidable

Decidability results

Good news: decidable problems for PA remain decidable

almost-sure safety existence of ν s.t. $\mathbb{P}^{\nu}(\mathcal{M} \models \Box F) = 1$

positive safety existence of ν s.t. $\mathbb{P}^{\nu}(\mathcal{M} \models \Box F) > 0$

almost-sure repeated reachability existence of ν s.t. $\mathbb{P}^{\nu}(\mathcal{M} \models \Box\Diamond F) = 1$

are all EXPTIME-complete.

Decidability results

Good news: decidable problems for PA remain decidable

almost-sure safety existence of ν s.t. $\mathbb{P}^{\nu}(\mathcal{M} \models \Box F) = 1$

positive safety existence of ν s.t. $\mathbb{P}^{\nu}(\mathcal{M} \models \Box F) > 0$

almost-sure repeated reachability existence of ν s.t. $\mathbb{P}^{\nu}(\mathcal{M} \models \Box\Diamond F) = 1$

are all EXPTIME-complete.

fixpoint algorithms on a powerset construction
belief-based strategies suffice except for positive safety

Decidability results

Good news: decidable problems for PA remain decidable

almost-sure safety existence of ν s.t. $\mathbb{P}^\nu(\mathcal{M} \models \Box F) = 1$

positive safety existence of ν s.t. $\mathbb{P}^\nu(\mathcal{M} \models \Box F) > 0$

almost-sure repeated reachability existence of ν s.t. $\mathbb{P}^\nu(\mathcal{M} \models \Box\Diamond F) = 1$

are all EXPTIME-complete.

fixpoint algorithms on a powerset construction
belief-based strategies suffice except for positive safety

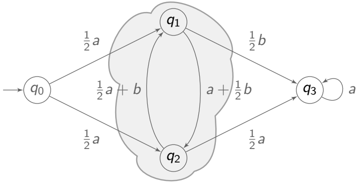

no belief-based strategy can achieve $\mathbb{P}^\nu(\mathcal{M} \models \Box\{q_0, q_1, q_2\}) > 0$
alternate a and b forever, guarantees a probability $\frac{1}{2}$

Decidability results

Good news: decidable problems for PA remain decidable

almost-sure safety existence of ν s.t. $\mathbb{P}^\nu(\mathcal{M} \models \Box F) = 1$

positive safety existence of ν s.t. $\mathbb{P}^\nu(\mathcal{M} \models \Box F) > 0$

almost-sure repeated reachability existence of ν s.t. $\mathbb{P}^\nu(\mathcal{M} \models \Box\Diamond F) = 1$

are all EXPTIME-complete.

fixpoint algorithms on a powerset construction
belief-based strategies suffice except for positive safety

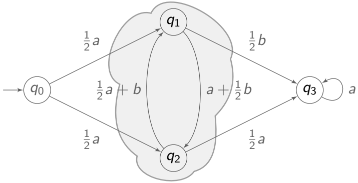

no belief-based strategy can achieve $\mathbb{P}^\nu(\mathcal{M} \models \Box\{q_0, q_1, q_2\}) > 0$
alternate a and b forever, guarantees a probability $\frac{1}{2}$

Open: decidability of non-null proportion with positive probability

$$\exists \nu,\ \mathbb{P}^\nu(\mathcal{M} \models \limsup_n \frac{\#\text{visits to } F \text{ in } n \text{ first steps}}{n} > 0) > 0?$$

Outline

Probabilistic automata

Partially observable MDP

Discussion

Life is hard...

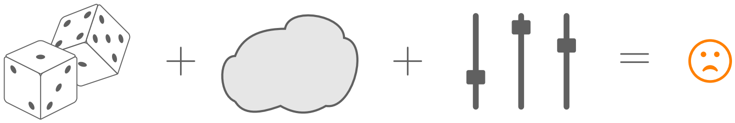

Life is hard...

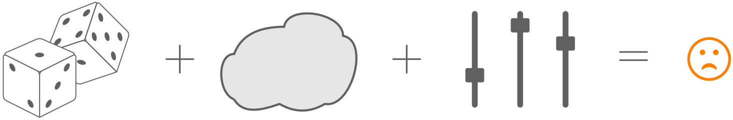

most optimization problems are undecidable

- notably quantitative questions
- but also some qualitative questions
- and undecidability is robust

... but there is still hope

- usual way arounds
  - decidable subclasses — Fijalkow *et al.*'12
  - restricted classes of strategies
  - approximations, although with no termination guarantees — Yu'06

... but there is still hope

- usual way arounds
  - decidable subclasses Fijalkow *et al.*'12
  - restricted classes of strategies
  - approximations, although with no termination guarantees Yu'06

- promising alternative: discretization
  - continuous distributions approximated by large discrete population

... but there is still hope

- usual way arounds
  - decidable subclasses — Fijalkow *et al.*'12
  - restricted classes of strategies
  - approximations, although with no termination guarantees — Yu'06

- promising alternative: discretization
  - continuous distributions approximated by large discrete population

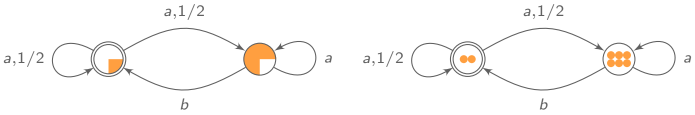

... but there is still hope

- usual way arounds
  - decidable subclasses — Fijalkow *et al.*'12
  - restricted classes of strategies
  - approximations, although with no termination guarantees — Yu'06

- promising alternative: discretization
  - continuous distributions approximated by large discrete population

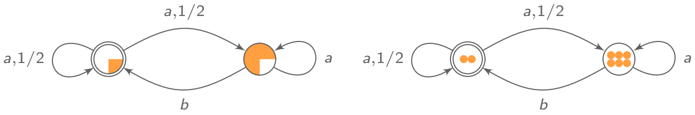

  - limit for large populations differs from continuous semantics
  - possible alternative semantics to PA/POMDP models, with more decidability results

... but there is still hope

- usual way arounds
  - decidable subclasses — Fijalkow *et al.*'12
  - restricted classes of strategies
  - approximations, although with no termination guarantees — Yu'06

- promising alternative: discretization
  - continuous distributions approximated by large discrete population

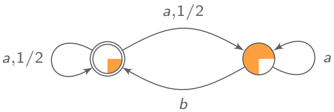

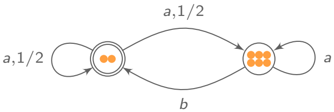

  - limit for large populations differs from continuous semantics
  - possible alternative semantics to PA/POMDP models, with more decidability results

Thank you!