Probabilities and partial observation

From probabilistic omega-automata to stochastic games of imperfect information

Nathalie Bertrand

Inria Rennes + Liverpool visitor

based on joint work with
Christel Baier & Marcus Größer
Blaise Genest & Hugo Gimbert

Motivation

Adversarial situations with

- probabilities: random choices, uncertainties, losses
- partial observation: distributed

Motivation

Adversarial situations with

- probabilities: random choices, uncertainties, losses
- partial observation: distributed

Monty Hall problem

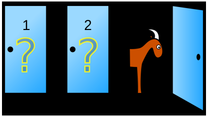

Motivation

Adversarial situations with

- probabilities: random choices, uncertainties, losses
- partial observation: distributed

Monty Hall problem

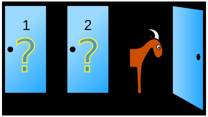

Ethernet protocol: random choice slot, collisions, maximum nb of trials

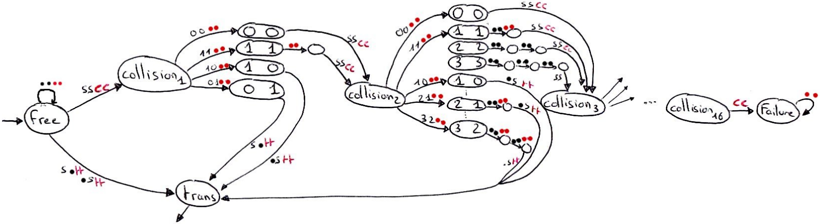

Partially observable MDP

Partially Observable MDP

A POMDP $(\mathcal{M}, \sim)$ consists of an MDP $\mathcal{M}$ equipped with an equivalence relation $\sim$ over states of $\mathcal{M}$.

Partially observable MDP

Partially Observable MDP

A POMDP $(\mathcal{M}, \sim)$ consists of an MDP $\mathcal{M}$ equipped with an equivalence relation $\sim$ over states of $\mathcal{M}$.

Problems for POMDPs

Given a POMDP $(\mathcal{M}, \sim)$, an LTL formula φ and $p \in [0, 1]$

Question is there a $\sim$ based scheduler $\mathcal{U}$ with $\mathbb{P}_{\mathcal{U}}(\varphi) \bowtie p$?

Partially observable MDP

Partially Observable MDP

A POMDP $(\mathcal{M}, \sim)$ consists of an MDP $\mathcal{M}$ equipped with an equivalence relation $\sim$ over states of $\mathcal{M}$.

Problems for POMDPs

Given a POMDP $(\mathcal{M}, \sim)$, an LTL formula φ and $p \in [0, 1]$

Question is there a $\sim$ based scheduler $\mathcal{U}$ with $\mathbb{P}_{\mathcal{U}}(\varphi) \bowtie p$?

State of the art (back in 2008)

- $\exists \mathcal{U} \sim$ based, $\mathbb{P}_{\mathcal{U}}(\Box F) > 0$ is EXPTIME-Comp. [de Alfaro 99]
- $\exists \mathcal{U} \sim$ based, $\mathbb{P}_{\mathcal{U}}(\Diamond F) \bowtie p$ is undecidable [Giro D'Argenio 07]
  ($p \in (0, 1)$ and $\bowtie \in \{\leq, \geq, =\}$)

Outline

1. Introduction
2. Probabilistic ω-automata
   - Examples and expressiveness
   - Emptiness problem
     - Positive semantics ($PBA_{>0}$)
     - Almost-sure semantics ($PBA_{=1}$)
3. Stochastic games of imperfect information
   - Framework
   - Qualitative determinacy
   - Memory requirements
   - Decidability
4. Conclusion

Probabilistic Büchi-automata

[Baier Größer 05]

Probabilistic Büchi Automata

NBA where nondeterminism is resolved by probabilities

$$\mathcal{L}_{>0}(\mathcal{A}) = \{w \in \Sigma^\omega \mid \mathbb{P}(\{\rho \in \mathsf{Runs}(w) \mid \rho \models \Box\Diamond F\}) > 0\}$$
$$\mathcal{L}_{=1}(\mathcal{A}) = \{w \in \Sigma^\omega \mid \mathbb{P}(\{\rho \in \mathsf{Runs}(w) \mid \rho \models \Box\Diamond F\}) = 1\}$$

Probabilistic Büchi-automata

[Baier Größer 05]

Probabilistic Büchi Automata

NBA where nondeterminism is resolved by probabilities

$$\mathcal{L}_{>0}(\mathcal{A}) = \{w \in \Sigma^\omega \mid \mathbb{P}(\{\rho \in \text{Runs}(w) \mid \rho \models \Box\Diamond F\}) > 0\}$$
$$\mathcal{L}_{=1}(\mathcal{A}) = \{w \in \Sigma^\omega \mid \mathbb{P}(\{\rho \in \text{Runs}(w) \mid \rho \models \Box\Diamond F\}) = 1\}$$

PBA are POMDP with trivial equivalence relation: $\forall p, q \; p \sim q$.

word for PBA $\equiv$ deterministic blind scheduler for MDP

Probabilistic Büchi-automata

[Baier Größer 05]

Probabilistic Büchi Automata

NBA where nondeterminism is resolved by probabilities

$$\mathcal{L}_{>0}(\mathcal{A}) = \{w \in \Sigma^\omega \mid \mathbb{P}(\{\rho \in \mathrm{Runs}(w) \mid \rho \models \Box\Diamond F\}) > 0\}$$
$$\mathcal{L}_{=1}(\mathcal{A}) = \{w \in \Sigma^\omega \mid \mathbb{P}(\{\rho \in \mathrm{Runs}(w) \mid \rho \models \Box\Diamond F\}) = 1\}$$

Example 1

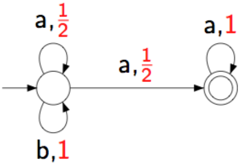

$$\mathcal{L}_{>0} = (a+b)^* a^\omega = \mathcal{L}_{\mathrm{NBA}}$$
$$\mathcal{L}_{=1} = b^* a^\omega$$

Probabilistic Büchi-automata

[Baier Größer 05]

Probabilistic Büchi Automata

NBA where nondeterminism is resolved by probabilities

$$\mathcal{L}_{>0}(\mathcal{A}) = \{w \in \Sigma^\omega \mid \mathbb{P}(\{\rho \in \mathrm{Runs}(w) \mid \rho \models \Box\Diamond F\}) > 0\}$$
$$\mathcal{L}_{=1}(\mathcal{A}) = \{w \in \Sigma^\omega \mid \mathbb{P}(\{\rho \in \mathrm{Runs}(w) \mid \rho \models \Box\Diamond F\}) = 1\}$$

Example 2

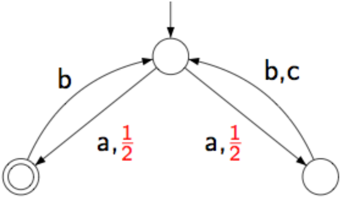

$$\mathcal{L}_{\mathrm{NBA}} = ((ac)^*(ab))^\omega$$
$$\mathcal{L}_{>0} = (ab + ac)^*(ab)^\omega$$
$$\mathcal{L}_{=1} = (ab)^\omega$$

Probabilistic Büchi-automata

[Baier Größer 05]

Probabilistic Büchi Automata

NBA where nondeterminism is resolved by probabilities

$$\mathcal{L}_{>0}(\mathcal{A}) = \{w \in \Sigma^\omega \mid \mathbb{P}(\{\rho \in \text{Runs}(w) \mid \rho \models \Box\Diamond F\}) > 0\}$$
$$\mathcal{L}_{=1}(\mathcal{A}) = \{w \in \Sigma^\omega \mid \mathbb{P}(\{\rho \in \text{Runs}(w) \mid \rho \models \Box\Diamond F\}) = 1\}$$

Example 3

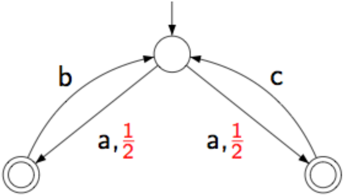

$$\mathcal{L}_{\text{NBA}} = (ab + ac)^\omega$$
$$\mathcal{L}_{>0} = \mathcal{L}_{=1} = \emptyset$$

Expressiveness: PBA vs NBA

Expressiveness

$\text{PBA}_{>0}$ are more expressive than NBA.
$\text{PBA}_{=1}$ and NBA are incomparable.

Expressiveness: PBA vs NBA

Expressiveness

$PBA_{>0}$ are more expressive than NBA.
$PBA_{=1}$ and NBA are incomparable.

- any NBA can be turned into an equivalent PBA
  $\rightarrow$ first turn NBA into an equivalent one deterministic in the limit
- example of a $PBA_{>0}$ whose language is not ω-regular

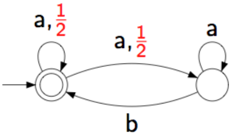

$$\mathcal{L} = \{a^{k_1} b a^{k_2} b \cdots \mid \prod_i (1 - \tfrac{1}{2}^{k_i}) > 0)\}$$

Expressiveness: PBA vs NBA

Expressiveness

$PBA_{>0}$ are more expressive than NBA.
$PBA_{=1}$ and NBA are incomparable.

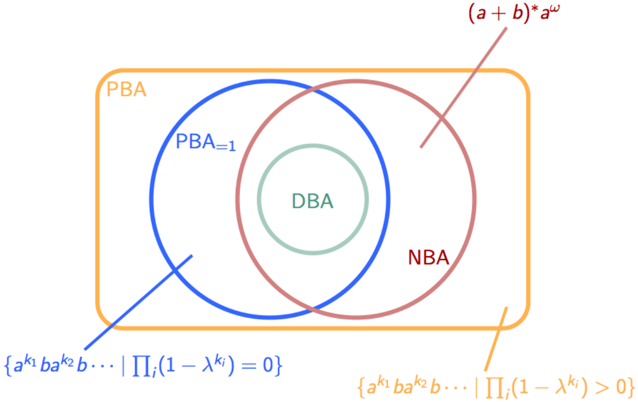

Probabilities matter

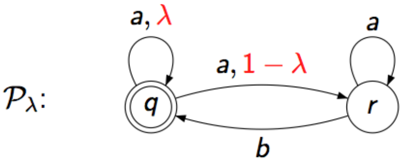

$$\mathcal{L}(\mathcal{P}_\lambda) = \{a^{k_1} b a^{k_2} b \cdots \mid \prod_i (1 - \lambda^{k_i}) > 0)\}$$

Probabilities matter

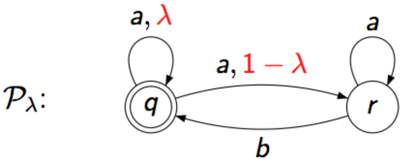

$$\mathcal{L}(\mathcal{P}_\lambda) = \{a^{k_1} b a^{k_2} b \cdots \mid \prod_i (1 - \lambda^{k_i}) > 0)\}$$

Lemma

For $0 < \lambda < \mu < 1$, $\mathcal{L}(\mathcal{P}_\lambda) \supsetneq \mathcal{L}(\mathcal{P}_\mu)$.

Emptiness for $PBA_{>0}$

[Baier B. Größer 08]

Theorem

The emptiness problem is undecidable for PBA.

Emptiness for $PBA_{>0}$

[Baier B. Größer 08]

Theorem

The emptiness problem is undecidable for PBA.

Proof sketch

Reduction of the modified emptiness problem for PFA

$$\mathcal{R} \text{ PFA with } \begin{cases} \forall w \; \mathbb{P}_{\mathcal{R}}(w) \leq \varepsilon & \text{or} \\ \exists w \; \mathbb{P}_{\mathcal{R}}(w) > 1 - \varepsilon \end{cases}$$

$\downarrow$

$\mathcal{P}_1$ and $\mathcal{P}_2$ PBA s.t.

$$\mathcal{L}^{>\varepsilon}(\mathcal{R}) = \emptyset \quad \Leftrightarrow \quad \mathcal{L}(\mathcal{P}_1) \cap \mathcal{L}(\mathcal{P}_2) = \emptyset$$

Consequences for POMDP

Undecidability results for POMDP

The following problems are undecidable

- Given $(\mathcal{M}, \sim)$ and F set of states of $\mathcal{M}$, is there a deterministic $\sim$ based $\mathcal{U}$ such that $\mathbb{P}_{\mathcal{U}}(\Box\Diamond F) > 0$.
- Given $(\mathcal{M}, \sim)$ and F set of states of $\mathcal{M}$, is there a deterministic $\sim$ based $\mathcal{U}$ such that $\mathbb{P}_{\mathcal{U}}(\Diamond\Box F) = 1$.

Consequences for POMDP

Undecidability results for POMDP

The following problems are undecidable

- Given $(\mathcal{M}, \sim)$ and F set of states of $\mathcal{M}$, is there a deterministic $\sim$ based $\mathcal{U}$ such that $\mathbb{P}_{\mathcal{U}}(\Box\Diamond F) > 0$.
- Given $(\mathcal{M}, \sim)$ and F set of states of $\mathcal{M}$, is there a deterministic $\sim$ based $\mathcal{U}$ such that $\mathbb{P}_{\mathcal{U}}(\Diamond\Box F) = 1$.

First undecidability results in qualitative verification of POMDP.

Emptiness for $PBA_{=1}$

Theorem

The emptiness problem is decidable for almost-sure PBA.

Emptiness for $PBA_{=1}$

Theorem

The emptiness problem is decidable for almost-sure PBA.

Proof sketch

1. emptiness problem and almost-sure reachability are interreducible for PBA

$$\exists w,\ \mathbb{P}^{w}_{\mathcal{P}}(\Box\Diamond F) = 1 \quad \equiv \quad \exists w,\ \mathbb{P}^{w}_{\mathcal{P}}(\Diamond M) = 1$$

Emptiness for $\text{PBA}_{=1}$

Theorem

The emptiness problem is decidable for almost-sure PBA.

Proof sketch

1. emptiness problem and almost-sure reachability are interreducible for PBA

$$\exists w,\ \mathbb{P}^{w}_{\mathcal{P}}(\Box\Diamond F) = 1 \quad \equiv \quad \exists w,\ \mathbb{P}^{w}_{\mathcal{P}}(\Diamond M) = 1$$

2. almost-sure reachability for POMDP is decidable

$$\exists \mathcal{U} \sim \text{based},\ \mathbb{P}^{\mathcal{U}}_{\mathcal{M}}(\Diamond M) = 1$$

Skip proof

Proof in more details: step 1

- $\exists w,\ \mathbb{P}^w_{\mathcal{P}}(\diamond M) = 1 \quad \ll \quad \exists w,\ \mathbb{P}^w_{\mathcal{P}}(\square\diamond F) = 1$
  Hint: $F = M$ and add self loops on F with probability one

Proof in more details: step 1

- $\exists w,\ \mathbb{P}^w_{\mathcal{P}}(\Diamond M) = 1 \quad \ll \quad \exists w,\ \mathbb{P}^w_{\mathcal{P}}(\Box\Diamond F) = 1$
  Hint: $F = M$ and add self loops on F with probability one
- $\exists w,\ \mathbb{P}^w_{\mathcal{P}}(\Box\Diamond F) = 1 \quad \ll \quad \exists w,\ \mathbb{P}^w_{\mathcal{P}}(\Diamond M) = 1$

Reduction:

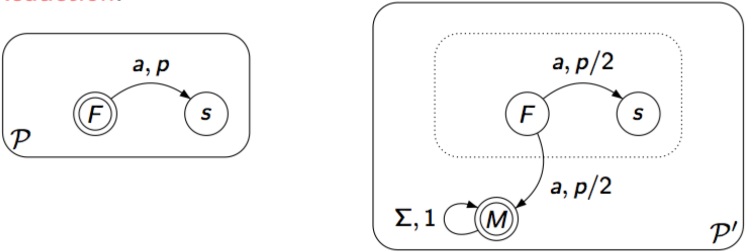

$$\forall w,\ \mathbb{P}^w_{\mathcal{P}}(\Box\Diamond F) = 1 \iff \mathbb{P}^w_{\mathcal{P}'}(\Diamond M) = 1$$

Proof in more details: step 2

Theorem

Given $(\mathcal{M}, \sim)$ and F set of states of $\mathcal{M}$, it is decidable whether there exists an observation-based $\mathcal{U}$ with $\mathbb{P}_{\mathcal{U}}(\Diamond M) = 1$.

Proof in more details: step 2

Theorem

Given $(\mathcal{M}, \sim)$ and F set of states of $\mathcal{M}$, it is decidable whether there exists an observation-based $\mathcal{U}$ with $\mathbb{P}_{\mathcal{U}}(\Diamond M) = 1$.

Idea: reduction to almost-sure reachability for MDP

From POMDP $(\mathcal{M}, \sim)$ build MDP $\mathcal{M}'$ by powerset construction with additional final state F'.

- if $\delta(r, a) \cap F = \emptyset$: traditional powerset construction
- if $\delta(r, a) \cap F \neq \emptyset$: go to F' with probability $1/2$, rest of probability mass uniformely distributed over non final successors

Outline

1. Introduction
2. Probabilistic ω-automata
   - Examples and expressiveness
   - Emptiness problem
     - Positive semantics ($PBA_{>0}$)
     - Almost-sure semantics ($PBA_{=1}$)
3. Stochastic games of imperfect information
   - Framework
   - Qualitative determinacy
   - Memory requirements
   - Decidability
4. Conclusion

Framework

[B. Genest Gimbert 09]

- two-player game
- partial observation on both sides
  $\longrightarrow$ signals received by the players
- probability on next state given current one and players' decisions
- qualitative objectives

Framework

[B. Genest Gimbert 09]

- two-player game
- partial observation on both sides
  $\rightarrow$ signals received by the players
- probability on next state given current one and players' decisions
- qualitative objectives

Strategy for Player i

Based on initial distribution and sequence of signals received so far, Player i chooses a distribution over actions.

Winning almost-surely and positively

Initial distribution δ, and strategy profile (σ, τ) induce a probability measure $\mathbb{P}^{\delta}_{\sigma,\tau}(\cdot)$ on maximal plays.

Winning almost-surely and positively

Initial distribution δ, and strategy profile (σ, τ) induce a probability measure $\mathbb{P}^{\delta}_{\sigma,\tau}(\cdot)$ on maximal plays.

φ: objective of the game for Player 1 (e.g. reachability, Büchi).

- From δ, Player 1 wins almost-surely if $\exists\sigma\forall\tau\ \mathbb{P}^{\delta}_{\sigma,\tau}(\varphi) = 1$.
- From δ, Player 1 wins positively if $\exists\sigma\forall\tau\ \mathbb{P}^{\delta}_{\sigma,\tau}(\varphi) > 0$.
- From δ, Player 2 wins almost-surely if $\exists\tau\forall\sigma\ \mathbb{P}^{\delta}_{\sigma,\tau}(\varphi) = 0$.
- From δ, Player 2 wins positively if $\exists\tau\forall\sigma\ \mathbb{P}^{\delta}_{\sigma,\tau}(\varphi) < 1$.

Winning almost-surely and positively

Initial distribution δ, and strategy profile (σ, τ) induce a probability measure $\mathbb{P}^{\delta}_{\sigma,\tau}(\cdot)$ on maximal plays.

φ: objective of the game for Player 1 (e.g. reachability, Büchi).

- From δ, Player 1 wins almost-surely if $\exists\sigma\forall\tau\ \mathbb{P}^{\delta}_{\sigma,\tau}(\varphi) = 1$.
- From δ, Player 1 wins positively if $\exists\sigma\forall\tau\ \mathbb{P}^{\delta}_{\sigma,\tau}(\varphi) > 0$.
- From δ, Player 2 wins almost-surely if $\exists\tau\forall\sigma\ \mathbb{P}^{\delta}_{\sigma,\tau}(\varphi) = 0$.
- From δ, Player 2 wins positively if $\exists\tau\forall\sigma\ \mathbb{P}^{\delta}_{\sigma,\tau}(\varphi) < 1$.

Player *i* a.-s. or pos. winning from δ only depends on support:

$$\mathbb{P}^{\delta}_{\sigma,\tau}(\varphi) = \sum_{s \in S} \delta(s) \cdot \mathbb{P}^{\mathbf{1}_s}_{\sigma,\tau}(\varphi)$$

Examples

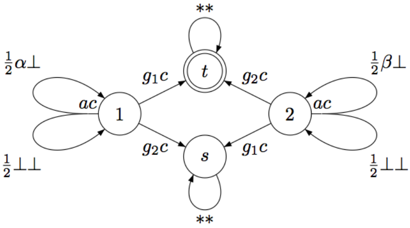

Initial support: $\{1, 2\}$.
Objective: reach t
Player 1 wins almost-surely.

Examples

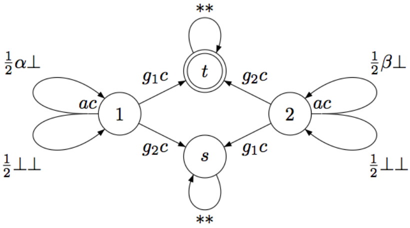

Initial support: $\{1, 2\}$.
Objective: reach t
Player 1 wins almost-surely.

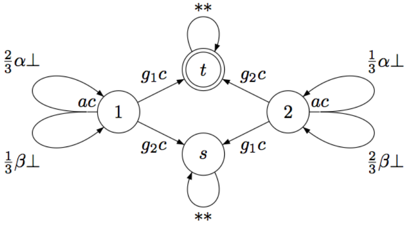

Initial support: $\{1, 2\}$.
Objective: reach t
Player 2 wins positively.

Consequence of undecidability of $PBA_{>0}$

Undecidability

Büchi games with positive probability (and co-Büchi games with probability one) are undecidable.

- 2-player Büchi games with positive probability generalize $PBA_{>0}$
- randomness for free in POMDP [Chatterjee Doyen Gimbert Henzinger 10]

Determinacy

Qualitative determinacy

In Büchi games, every initial distribution is

- almost-surely winning for Player 1, or
- positively winning for Player 2.

Determinacy

Qualitative determinacy

In Büchi games, every initial distribution is

- almost-surely winning for Player 1, or
- positively winning for Player 2.

Implies qualitative determinacy for reachability objectives as well.

NB: co-Büchi games are not qualitatively determined. ▸ Details

Determinacy: Proof sketch

Belief: possible states of the game according to signals received

Determinacy: Proof sketch

Belief: possible states of the game according to signals received

$\mathcal{L} \subseteq \mathcal{P}(S)$: set of supports where Player 2 does not win positively

Determinacy: Proof sketch

Belief: possible states of the game according to signals received

$\mathcal{L} \subseteq \mathcal{P}(S)$: set of supports where Player 2 does not win positively

- for every $L \in \mathcal{L}$, Player 1 has at least one safe action, i.e. can stay in $\mathcal{L}$
- σ strategy for Player 1 that, from $L \in \mathcal{L}$ selects (uniformly at random) a safe set of actions for L and plays uniform distribution

Determinacy: Proof sketch

Belief: possible states of the game according to signals received

$\mathcal{L} \subseteq \mathcal{P}(S)$: set of supports where Player 2 does not win positively

- for every $L \in \mathcal{L}$, Player 1 has at least one safe action, i.e. can stay in $\mathcal{L}$
- σ strategy for Player 1 that, from $L \in \mathcal{L}$ selects (uniformly at random) a safe set of actions for L and plays uniform distribution
  - σ ensures to stay in $\mathcal{L}$
  - σ is almost-surely winning

Determinacy: Proof sketch

Belief: possible states of the game according to signals received

$\mathcal{L} \subseteq \mathcal{P}(S)$: set of supports where Player 2 does not win positively

- for every $L \in \mathcal{L}$, Player 1 has at least one safe action, i.e. can stay in $\mathcal{L}$
- σ strategy for Player 1 that, from $L \in \mathcal{L}$ selects (uniformly at random) a safe set of actions for L and plays uniform distribution
  - σ ensures to stay in $\mathcal{L}$
  - σ is almost-surely winning
    - $M = \{s \in S \mid \exists \tau\ \mathbb{P}^{\mathbf{1}_s}_{\sigma,\tau}(\Box\neg F) = 1\}$

Determinacy: Proof sketch

Belief: possible states of the game according to signals received

$\mathcal{L} \subseteq \mathcal{P}(S)$: set of supports where Player 2 does not win positively

- for every $L \in \mathcal{L}$, Player 1 has at least one safe action, i.e. can stay in $\mathcal{L}$
- σ strategy for Player 1 that, from $L \in \mathcal{L}$ selects (uniformly at random) a safe set of actions for L and plays uniform distribution
  - σ ensures to stay in $\mathcal{L}$
  - σ is almost-surely winning
    - $M = \{s \in S \mid \exists \tau \; \mathbb{P}^{\mathbf{1}_s}_{\sigma,\tau}(\Box\neg F) = 1\}$
    - $\mathcal{L} \subseteq \mathcal{P}(S \setminus M)$

Determinacy: Proof sketch

Belief: possible states of the game according to signals received

$\mathcal{L} \subseteq \mathcal{P}(S)$: set of supports where Player 2 does not win positively

- for every $L \in \mathcal{L}$, Player 1 has at least one safe action, i.e. can stay in $\mathcal{L}$
- σ strategy for Player 1 that, from $L \in \mathcal{L}$ selects (uniformly at random) a safe set of actions for L and plays uniform distribution
  - σ ensures to stay in $\mathcal{L}$
  - σ is almost-surely winning
    - $M = \{s \in S \mid \exists \tau \ \mathbb{P}^{\mathbf{1}_s}_{\sigma,\tau}(\Box \neg F) = 1\}$
    - $\mathcal{L} \subseteq \mathcal{P}(S \setminus M)$
    - $\exists N \in \mathbb{N} \ \forall s \in S \setminus M \ \forall \tau \ \mathbb{P}^{\mathbf{1}_s}_{\sigma,\tau}(\Diamond F) > 1/N$

Determinacy: Proof sketch

Belief: possible states of the game according to signals received

$\mathcal{L} \subseteq \mathcal{P}(S)$: set of supports where Player 2 does not win positively

- for every $L \in \mathcal{L}$, Player 1 has at least one safe action, i.e. can stay in $\mathcal{L}$
- σ strategy for Player 1 that, from $L \in \mathcal{L}$ selects (uniformly at random) a safe set of actions for L and plays uniform distribution
  - σ ensures to stay in $\mathcal{L}$
  - σ is almost-surely winning
    - $M = \{s \in S \mid \exists \tau \ \mathbb{P}^{\mathbf{1}_s}_{\sigma,\tau}(\Box\neg F) = 1\}$
    - $\mathcal{L} \subseteq \mathcal{P}(S \setminus M)$
    - $\exists N \in \mathbb{N} \ \forall s \in S \setminus M \ \forall \tau \ \mathbb{P}^{\mathbf{1}_s}_{\sigma,\tau}(\Diamond F) > 1/N$

Player 1 wins almost-surely from $L \subseteq \mathcal{L}$, with belief-based strategy.

Memory requirements

Player 1 needs exponential memory to win almost-surely.
→ has to remember belief states

Player 2 needs doubly exponential memory to win positively.
→ has to remember possible beliefs of Player 1
(beliefs of beliefs)

Deciding stochastic games with signals

Decidability and complexity

Deciding whether Player 1 almost-surely wins a reachability or Büchi game is 2EXPTIME-complete.

Deciding stochastic games with signals

Decidability and complexity

Deciding whether Player 1 almost-surely wins a reachability or Büchi game is 2EXPTIME-complete.

- Player 1 better informed than Player 2: 2EXPTIME-complete
- Player 2 better informed than Player 1: EXPTIME-complete
- Player 1 perfectly informed: EXPTIME-complete

Concluding remarks

More results on probabilistic ω-automata in [Baier B. Größer 12]

Results on stochastic games with signals very dependent on the precise framework. E.g. for deterministic strategies or deterministic memory updates, the memory size for Player 1 may be a tower of exponential [Chatterjee Doyen]

Concluding remarks

More results on probabilistic ω-automata in [Baier B. Größer 12]

Results on stochastic games with signals very dependent on the precise framework. E.g. for deterministic strategies or deterministic memory updates, the memory size for Player 1 may be a tower of exponential [Chatterjee Doyen]

Baier, B. and Größer. Probabilistic ω-automata. Journal of the ACM, 2012.

B., Genest and Gimbert. Qualitative determinacy and decidability of stochastic games with signals. Proceedings of LICS, 2009.

Detour: co-Büchi games

Co-Büchi games are not qualitatively determined.

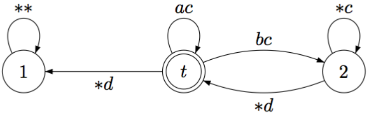

Initial state: t
Player 1 perfectly informed
Player 2 blind

Objective: avoid t from some point on

- Player 1 has no almost-surely winning strategy
- Player 2 has no positively winning strategy

Back to main