Probabilistic Aspects of Computer Science: TD5
Concurrent Stochastic Games

Chargé de TD: Adrien Koutsos

October 10, 2017

Exercise 1: Normal-Form Games

A *normal-form game* $\mathcal{G} = (D_{\mathsf{Max}}, D_{\mathsf{Min}}, f)$ is defined by:

- D_{Max} (respectively D_{Min}), the finite set of *pure* strategies of Player Max (respectively Min);
- f, a mapping from $D_{\mathsf{Max}} \times D_{\mathsf{Min}}$ to $\mathbb{R}$.

The set of *mixed* strategies for player $X \in \{\mathsf{Max}, \mathsf{Min}\}$, A_X is $Dist(D_X)$ the set of distributions over D_X. One extends the mapping f over strategies in the usual way:

$$f(\sigma, \tau) \stackrel{\text{def}}{=} \sum_{(d,d') \in D_{\mathsf{Max}} \times D_{\mathsf{Min}}} \sigma(d)\tau(d')f(d,d')$$

The semantic of a normal form game can be described as a one-shot concurrent game. Any player randomly selects an action (i.e. a pure strategy) following the distribution of the mixed strategy. These choices are independent. Then the reward is evaluated for the pair of selected actions. As randomness is introduced by the strategies Max (resp. Min) aim at maximizing (resp. minimizing) the expected value.

Max \ Min	Head	Tail
Head	1	0
Tail	0	1

Figure 1: The two-coins game

In the two-coins game, any player selects a possible biased coin. Then the two coins are tossed. Player Max wins one euro if the results are identical. The matrix of this game is presented in Figure 1.

Question 1. Show that no pure strategy is optimal in this game.

Question 2. Find the optimal strategies in this game and show that there are unique.

Question 3. Show that optimal values for both players in a normal-form game can be expressed as the solution of a linear program.

Question 4. Show that these linear programs are dual. Deduce that normal-form games are determined and that there exist optimal strategies.

A one-sided countable normal-form game extends normal form game by allowing a player to have a countable set of pure strategies. W.l.o.g we assume that Max has a countable set of pure strategies.

Question 5. Show that the one-sided countable normal-form game is determined and that the player with finite set of pure strategies has an optimal strategy.

Question 6. Show that extend normal-form games where both players have countable sets of pure strategies are not necessarily determined.

Exercise 2: Reachability Games

A *reachability stochastic game* $\mathcal{G} = (S, \{X_s\}_{s\in S}, \{Y_s\}_{s\in S}, p, r, s^+)$ is defined by:

- S, a non empty finite set of states including a target state s^+;
- For all $s \in S$, X_s and Y_s are non empty finite sets of actions;
- p, the transition probability that associates with all $s, s' \in S$, $x \in X_s$, $y \in Y_s$, the transition probability $p(s'|s, x, y) \geq 0$ with $\sum_{s'\in S} p(s'|s, x, y) = 1$ and $p(s^+|s^+, x, y) = 1$.

Player Max (resp. Min) aims at maximizing (resp. minimizing) the probability to reach s^+.

Question 1. We study the reachability game of Figure 2. Show that the game is determined but that player Max has no optimal strategy (*Hint: define* $p_n^{\sigma,\tau}$ *as the probability to stay in* s^0 *for the first* n *steps and* $(1 - p_n^{\sigma,\tau})q_n^{\sigma,\tau}$ *as the probability to reach* s^+ *from* s_0 *during the first* n *steps w.r.t. strategies* σ, τ.).

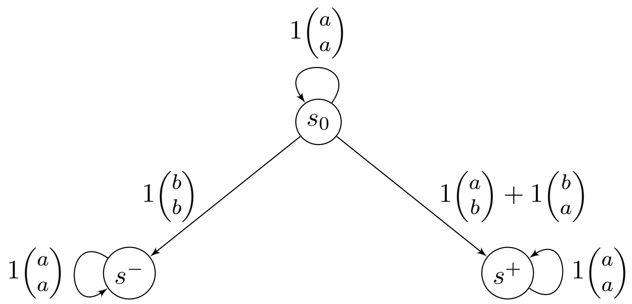

Figure 2: An example of reachability stochastic game

For any s and any vector $\mathbf{v}$ in $[0, 1]^S$ with $\mathbf{v}[s^+] = 1$, we define the normal form game $\Delta_s(\mathbf{v})$ by:

$$f(x, y) = \sum_{s'\in S} p(s'|s, x, y)\mathbf{v}[s']$$

Let us introduce the operator $\mathcal{L}_\mathcal{G}$ defined by: for all $s \neq s^+$, $\mathcal{L}_\mathcal{G}(\mathbf{v})[s] = val_{\Delta_s(\mathbf{v})}$ (observe that this implies $\mathcal{L}_\mathcal{G}(\mathbf{v})[s^+] = 1$).

Question 2. Prove that $\mathcal{L}_\mathcal{G}$ has a minimal fixed point that will be denoted $\mathbf{m}_\mathcal{G}$.

Question 3. Prove that $\mathbf{m}_\mathcal{G} \geq val_\mathcal{G}^\uparrow$ using a memoryless strategy of player Min.

Question 4. Prove that $\mathbf{m}_\mathcal{G} \leq val_\mathcal{G}^\downarrow$. Conclude that reachability games are determined and that player Min has an optimal memoryless strategy.

Question 5. Compute the value of the game of Figure 3. Compare with the value of a turn-based stochastic reachability game.

Exercise 3: Discounted Games

A *discounted stochastic game* $\mathcal{G} = (S, \{X_s\}_{s\in S}, \{Y_s\}_{s\in S}, p, r, \lambda)$ is defined by:

- S, a non empty finite set of states.
- For all $s \in S$, X_s and Y_s are non empty finite sets of actions;
- p, the transition probability that associates with all $s, s' \in S$, $x \in X_s$, $y \in Y_s$, the transition probability $p(s'|s, x, y) \geq 0$ with $\sum_{s'\in S} p(s'|s, x, y) = 1$.
- r, the reward function that associates to all $s \in S$, $x \in X_s$, $y \in Y_s$ the reward $r(s, x, y) \in \mathbb{R}$.

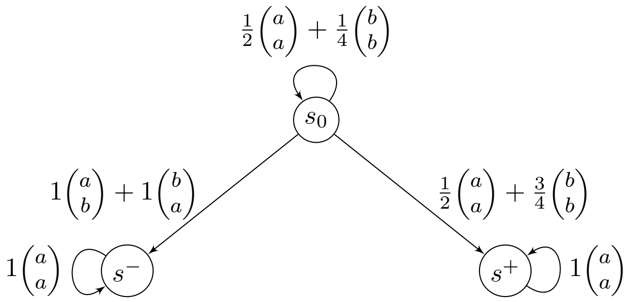

Figure 3: Another example of reachability stochastic game

- λ, the discounted factor with $0 < \lambda < 1$.

Player Max (resp. Min) aims at maximizing (resp. minimizing) the expected discounted reward:

$$\sum_{i \in \mathbb{N}} \lambda^n \, \mathbb{E}_h(r(S_n(h), X_n(h), Y_n(h)))$$

For any state s and vector $\mathbf{v}$ in $\mathbb{R}^{|S|}$ we define the normal form game $\Delta_s(\mathbf{v})$ by:

$$f(x, y) = r(s, x, y) + \lambda \sum_{s' \in S} p(s'|s, x, y)\mathbf{v}[s']$$

Let us introduce the operator $\mathcal{L}_\mathcal{G}$ defined by for all s, $\mathcal{L}_\mathcal{G}(\mathbf{v})[s] = val_{\Delta_s(\mathbf{v})}$.

Question 1. Prove that $\mathcal{L}_\mathcal{G}$ has a unique fixed point that will be denoted $\mathbf{m}_\mathcal{G}$ (*Hint: study the Lipschitz continuity of* $\mathcal{L}_\mathcal{G}$).

Question 2. For all s, let σ_s, τ_s be the optimal strategies for, respectively, Max and Min in $\Delta_s(\mathbf{m}_\mathcal{G})$. Let σ^∞ be the memoryless strategy such that $\sigma^\infty(s) = \sigma_s$. Similarly let τ^∞ be the memoryless strategy such that $\tau(s) = \tau_s$. Show that $\mathcal{G}$ is determined with optimal memoryless strategies $\sigma^\infty, \tau^\infty$ (*Hint: show that* $\mathbf{m}_\mathcal{G} \leq val^\downarrow_\mathcal{G}$ *and* $\mathbf{m}_\mathcal{G} \geq val^\uparrow_\mathcal{G}$).