On Strong Determinacy of Countable Stochastic Games

S. Kiefer, Richard Mayr, M. Shirmohammadi, D. Wojtczak

University of Oxford, Edinburgh, Liverpool. UK

Highlights 2017

Main result

Countably infinite perfect information 2-player stochastic games.

Theorem

Almost-sure objectives are **strongly determined**, *i.e., there exist winning strategies.*

- *Either Player Max ($\Box$) has a strategy that enforces that the objective holds almost surely (i.e., with probability* 1*),*
- *Or Player Min ($\Diamond$) has a strategy that enforces that the objective holds with probability* < 1.

Unlike for other thresholds like $\geq 1/2$ where winning strategies do **not** exist in general.
(Not even if the game is finitely branching.)

2-Player Stochastic Games

Definition (Game)

$\mathcal{G} = (S, (S_\Box, S_\Diamond, S_\bigcirc), \longrightarrow, P)$.
Countable set of states S. Transition relation $\longrightarrow \subseteq S \times S$.
$S_\Box$ are states of player $\Box$ (Maximizer).
$S_\Diamond$ are states of player $\Diamond$ (Minimizer).
$S_\bigcirc$ are random states. $P : S_\bigcirc \to \mathcal{D}(S)$ assigns distribution over successors.

Games can be infinitely branching in general.

Strategies can use memory/randomization in general.
Objectives defined as usual: Borel σ-algebra $\mathcal{F}$ induced by cone sets and measure by Carathéodory's extension theorem.
E.g., Reachability, Büchi, co-Büchi, Parity, ω-regular, Borel.

Weak Determinacy

Definition (Weak Determinacy)

Given an objective $\mathcal{E}$ in a game $\mathcal{G}$, state s has **value** if

$$\sup_{\sigma\in\Sigma}\inf_{\pi\in\Pi}\mathcal{P}_{\mathcal{G},s,\sigma,\pi}(\mathcal{E}) = \inf_{\pi\in\Pi}\sup_{\sigma\in\Sigma}\mathcal{P}_{\mathcal{G},s,\sigma,\pi}(\mathcal{E}).$$

Let $\mathtt{val}_{\mathcal{G}}(s)$ denote the value, where it exists.
A game with a fixed objective is called **weakly determined** iff every state has value.

Theorem (follows immediately from [Maitra & Sudderth:1998])

Countable stochastic games with Borel objectives are weakly determined.

In particular, this implies the existence of ϵ-optimal strategies.

Threshold Objectives

Definition (Threshold objective)

Given objective $\mathcal{E}$ and $\rhd \in \{\geq, >\}$ and threshold $c \in [0,1]$, we define the *threshold objective* $(\mathcal{E}, \rhd c)$.

Definition (Strong Determinacy)

A threshold objective $(\mathcal{E}, \rhd c)$ is **strongly determined** iff either player $\Box$ or player $\Diamond$ has a **winning strategy**, i.e., either

- $\exists \sigma \forall \pi \; \mathcal{P}_{\mathcal{G},s,\sigma,\pi}(\mathcal{E}) \rhd c$, or
- $\exists \pi \forall \sigma \; \mathcal{P}_{\mathcal{G},s,\sigma,\pi}(\mathcal{E}) \not\rhd c$.

Does strong determinacy always hold?
If $\mathtt{val}_{\mathcal{G}}(s) > c$ or $\mathtt{val}_{\mathcal{G}}(s) < c$ then ϵ-optimal strategies are winning strategies.
But what if $\mathtt{val}_{\mathcal{G}}(s) = c$?

Counterexample to Strong Determinacy

Theorem (KMSW: LICS 2017, inspired by Kučera et al. 2011)

Consider Büchi objectives with thresholds $\triangleright c$ with $c \in (0, 1)$. Strong determinacy does **not hold**, *not even if the game is finitely branching. (Similarly for co-Büchi.)*

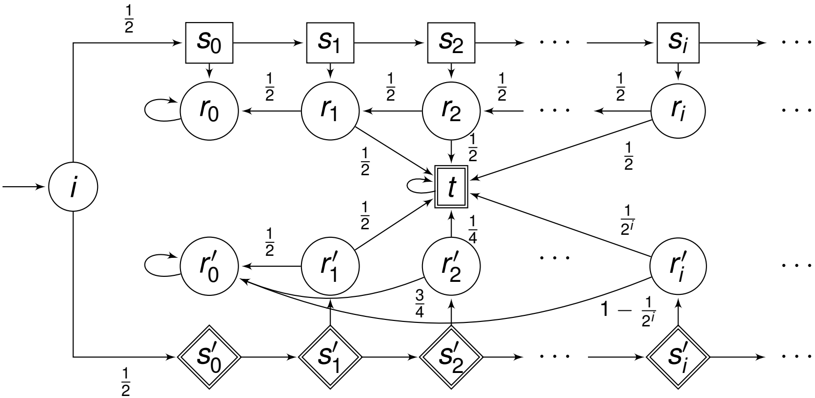

Almost-sure Objectives

Counterexamples only work for thresholds $c \in (0, 1)$.
What about the case of $c = 1$?
Almost-sure objective: Player $\Box$ wants the property to hold with probability 1.

Theorem (Gimbert & Horn: Soda 2010)

Almost-sure objectives in **finite** *stochastic games with* **tail** *objectives are strongly determined.*

Our main result: Drop these restrictions.

Theorem (KMSW: LICS 2017)

Almost-sure objectives in **countable** *stochastic games (even infinitely branching ones) are strongly determined.*

Note that this holds for general Borel objectives.

Almost-sure Objectives

Theorem (KMSW: LICS 2017)

Almost-sure objectives in **countable** *stochastic games (even infinitely branching ones) are strongly determined.*

Main ingredients of the proof:

1. Transformation of general objectives into (something similar to) tail objectives by encoding the history into the state.
   Easy, since we consider countable state spaces anyway.
   (Of course, this transformation does not preserve properties about the memory requirements of winning strategies.)
2. Transfinite induction. Construct a transfinite sequence of subgames, by removing parts of the game that player Max ($\Box$) cannot risk entering.

Almost-sure Objectives (cont.)

3. Weak determinacy of stochastic games [Maitra & Sudderth: 1998]. We can use ϵ-optimal strategies as building blocks to construct the winning strategies of players $\Box$ and $\Diamond$, respectively.
4. The concept of a "reset" strategy from [Gimbert & Horn: 2010]. Player $\Box$ might need to restart with a new ϵ-optimal strategy from the current state if the value of this state w.r.t. the previous strategy drops too low.
5. Lévy's zero-one law. This guarantees that the number of resets is almost-surely finite and the constructed strategy of player $\Box$ actually wins almost surely.

What about Memory?

On countably infinite arenas, just about every objective requires infinite memory (for at least one of the players).

Exceptions: Only for finitely branching games.

- Almost-sure reachability can be won MD.
- Almost-sure Büchi can be won MD.

Summary of Results

Infinitely branching games.

Objective	> 0	$> c$	$\geq c$	$= 1$
Reachability	✓(MD)	×	×	✔(¬FR)
Büchi	✔(¬FR)	×	×	✔(¬FR)
Borel	✔(¬FR)	×	×	✔(¬FR)

Finitely branching games.

Objective	> 0	$> c$	$\geq c$	$= 1$
Reachability	✓(MD)	✓(MD)	✓(¬FR)	✓**(MD)**
Büchi	✔(¬FR)	✖	✖	✔**(MD)**
Borel	✔(¬FR)	✖	✖	✔(¬FR)

Results for Safety and co-Büchi are implicit, since they are dual to Reachability and Büchi, respectively.