On the Complexity of Branching Games with Regular Conditions

Marcin Przybyłko[1]
(joint work with Michał Skrzypczak[2])

University of New Caledonia[1]
University of Warsaw[1,2]

Highlights 2016

Branching games

(a.k.a *tree games*) proposed by M.Mio as a framework for game semantics for the probabilistic μ-calculus

achieved by Mio in the form of
stochastic meta-parity games

extend games on graphs by adding new type of positions – branching positions – that split game into several independent and concurrent sub-games

... so the outcome is a tree,
but strategies still work on words.

Outline

Branching games:

1. the definitions,
2. the determinacy,
3. and the computation of games' values.

Games

Branching game $G = \langle \mathtt{B}, \Phi \rangle$ is played by two players (Eve, Adam) and consists of

- a *branching board* $\mathtt{B}$,
- and an *objective* Φ.

... that can be seen as a *rulebook* ...

Board and objective

Branching games – board

the *board* – (unfolding of a) binary graph, consisting of Adam's, Eve's, Nature's and branching vertices.

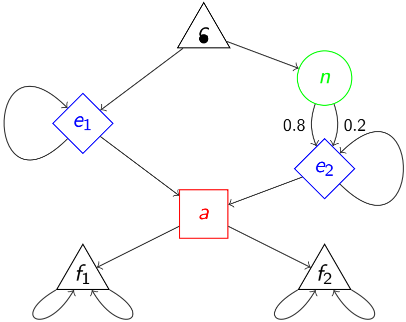

The black token marks the initial vertex.

Branching games, the game-play

Board

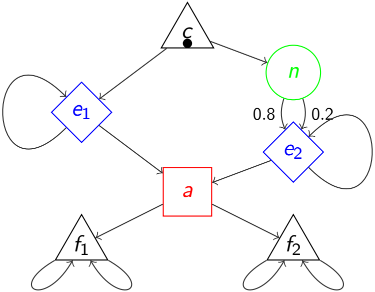

Outcome

c

Branching games, the game-play

Board

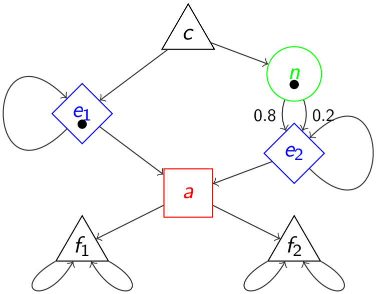

Outcome

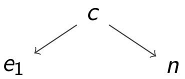

Branching games, the game-play

Board

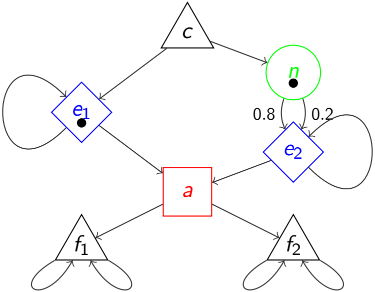

Outcome

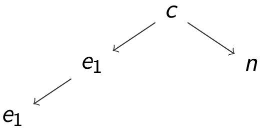

Branching games, the game-play

Board

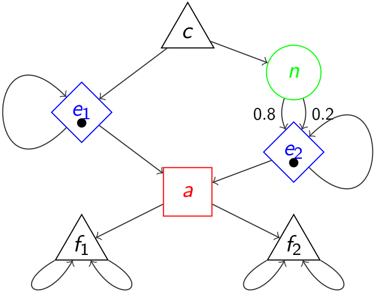

Outcome

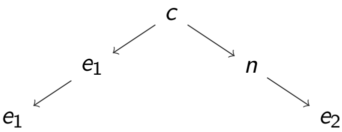

Branching games, the game-play

Board

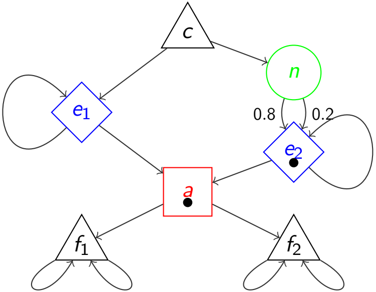

Outcome

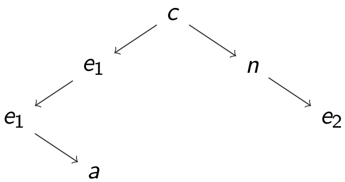

Branching games, the game-play

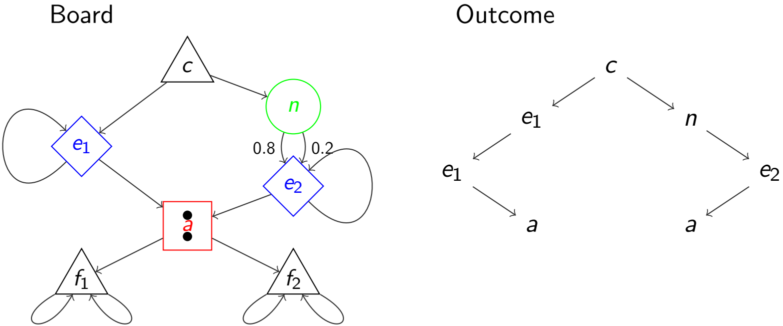

Branching games, the game-play

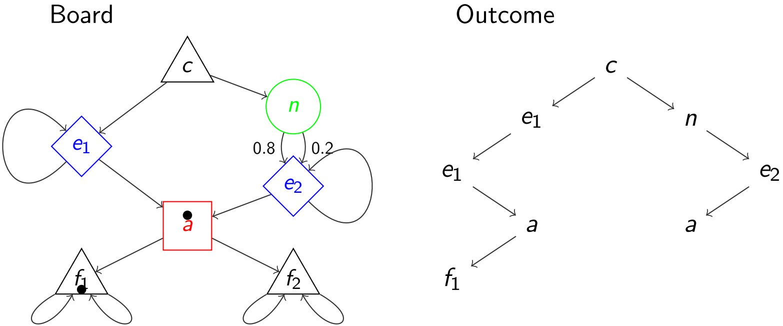

Branching games, the game-play

Board

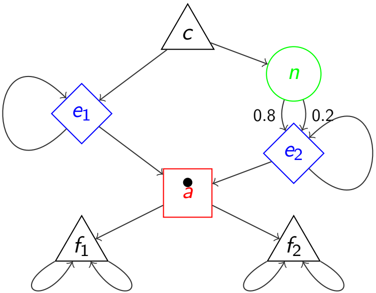

Outcome

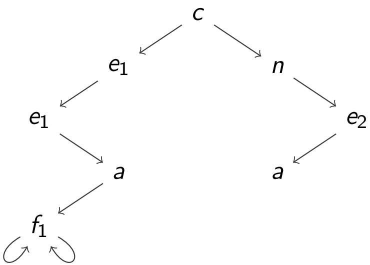

Branching games, the game-play

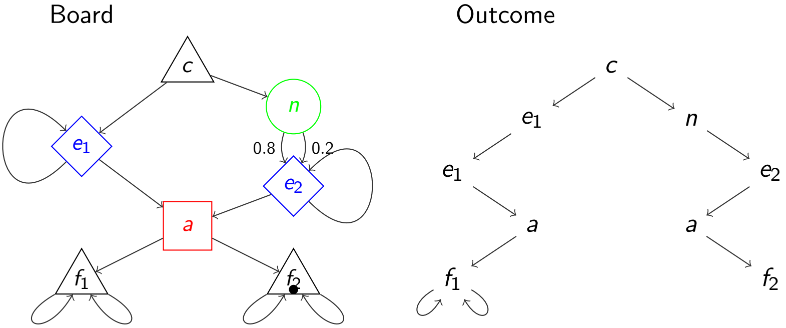

Branching games, the game-play

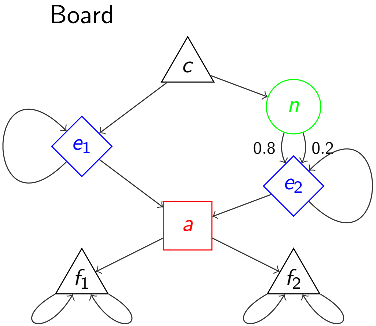

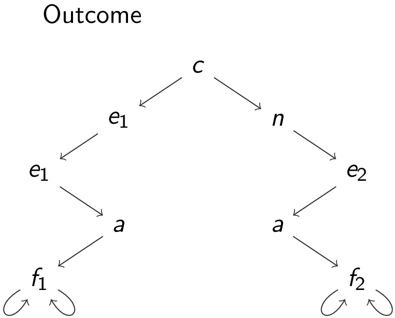

Branching games, the game-play

Board

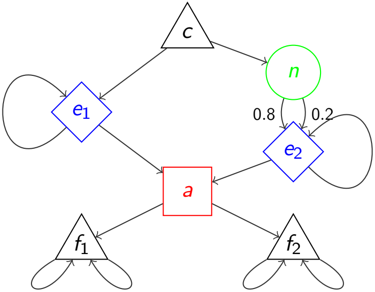

Outcome

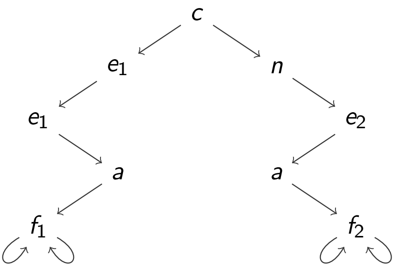

outcome – an infinite labelled tree

Branching games – objective

objective – a function $\Phi : Tr_\Gamma \to \mathbb{R}^+$ evaluated on the outcome

Φ is called the pay-off function
Eve wants to maximise the pay-off,
Adam to minimise it.

outcome – a measure $\mu^{\sigma,\pi}$ on set Tr_Γ defined by the actions (*strategies* σ, π) of the players.

Tr_Γ – a set of trees labelled with some finite set Γ.

Terminology

We say that a game $G = \langle \mathtt{B}, \Phi \rangle$

- is *finitary* if the set of vertices is finite and the probability values are rational.
- is $\mathcal{P}$-*branching* if the board $\mathtt{B}$ uses a subset $\mathcal{P} \subseteq \{A, E, \mathcal{N}, \mathcal{B}\}$ of position types.

  e.g. $\{A, E, \mathcal{N}\}$-game is a simple game on graphs

- has regular winning objective $\Phi : Tr_\Gamma \to \mathbb{R}^+$ if $\Phi(t) = \begin{cases} 1 & \text{if } t \in L \\ 0 & \text{otherwise} \end{cases}$ for some regular language of trees L.

Game values

Strategies

Let $G = \langle \mathrm{B}, \Phi \rangle$ be a branching game.

- Eve's *pure* strategy $\sigma \in \Sigma_{\mathrm{B}}^{E}$ is any function $\sigma : \{L, R\}^* \to \{L, R\}$.
- Eve's *behavioural* strategy $\sigma_B \in \Sigma_{\mathrm{B}}^{BE}$ is any function $\sigma_B : \{L, R\}^* \to \mu(\{L, R\})$.

  $\mu(S)$ denotes a Borel probabilistic measure on the set S
- Eve's *mixed* strategy $\sigma_M \in \Sigma_{\mathrm{B}}^{ME}$ is any function $\sigma_M \in \mu(\Sigma_{\mathrm{B}}^{E})$.
- Adam's sets of strategies are $\Sigma_{\mathrm{B}}^{A}, \Sigma_{\mathrm{B}}^{BA}, \Sigma_{\mathrm{B}}^{MA}$

Partial game values

The play value val_{G} of two strategies σ_m, π_m :

$$val_{\mathrm{G}}(\sigma_m, \pi_m) \stackrel{\text{def}}{=} \int \Phi d\mu^{\sigma_m, \pi_m}$$

in fact, $val_{\mathrm{G}}(\sigma_m, \pi_m) \stackrel{\text{def}}{=} \mu^{\sigma_m, \pi_m}(\Phi^{-1}(1))$

The X *value of* G *for Eve* (resp. Adam) is defined as

$$val_{\mathrm{G}}^{XE} \stackrel{\text{def}}{=} \sup_{\sigma \in \Sigma_{\mathrm{B}}^{XE}} \inf_{\pi \in \Sigma_{\mathrm{B}}^{A}} val_{\mathrm{G}}(\sigma, \pi),$$

$$val_{\mathrm{G}}^{XA} \stackrel{\text{def}}{=} \inf_{\pi \in \Sigma_{\mathrm{B}}^{XA}} \sup_{\sigma \in \Sigma_{\mathrm{B}}^{E}} val_{\mathrm{G}}(\sigma, \pi).$$

where $X \in \{\varepsilon, B, M\}$ – 6 values in total

Partial game values

The play value val_{G} of two strategies σ_m, π_m :

$$val_{\mathrm{G}}(\sigma_m, \pi_m) \stackrel{\text{def}}{=} \int \Phi d\mu^{\sigma_m,\pi_m}$$

in fact, $val_{\mathrm{G}}(\sigma_m, \pi_m) \stackrel{\text{def}}{=} \mu^{\sigma_m,\pi_m}(\Phi^{-1}(1))$

The X *value of* G *for Eve* (resp. Adam) is defined as

$$val_{\mathrm{G}}^{XE} \stackrel{\text{def}}{=} \sup_{\sigma \in \Sigma_{\mathrm{B}}^{XE}} \inf_{\pi \in \Sigma_{\mathrm{B}}^{A}} val_{\mathrm{G}}(\sigma, \pi),$$

$$val_{\mathrm{G}}^{XA} \stackrel{\text{def}}{=} \inf_{\pi \in \Sigma_{\mathrm{B}}^{XA}} \sup_{\sigma \in \Sigma_{\mathrm{B}}^{E}} val_{\mathrm{G}}(\sigma, \pi).$$

where $X \in \{\varepsilon, B, M\}$ – 6 values in total

... so what do we know about the values?

Determinacy

Determinacy

We say that game G is *determined (under pure strategies)* if

$$val^E_{\mathrm{G}} = val^A_{\mathrm{G}},$$

i.e. if

$$\sup_{\sigma \in \Sigma^E_{\mathrm{B}}} \inf_{\pi \in \Sigma^A_{\mathrm{B}}} val_{\mathrm{G}}(\sigma, \pi) = \inf_{\pi \in \Sigma^A_{\mathrm{B}}} \sup_{\sigma \in \Sigma^E_{\mathrm{B}}} val_{\mathrm{G}}(\sigma, \pi).$$

e.g. parity games – one of the players has a winning strategy

We say that G is *determined under mixed strategies* if

$$val^{ME}_{\mathrm{G}} = val^{MA}_{\mathrm{G}}.$$

e.g. Blackwell games are determined under mixed strategies

Determinacy

We say that game G is *determined (under pure strategies)* if

$$val^E_{\mathrm{G}} = val^A_{\mathrm{G}},$$

i.e. if

$$\sup_{\sigma \in \Sigma^E_{\mathrm{B}}} \inf_{\pi \in \Sigma^A_{\mathrm{B}}} val_{\mathrm{G}}(\sigma, \pi) = \inf_{\pi \in \Sigma^A_{\mathrm{B}}} \sup_{\sigma \in \Sigma^E_{\mathrm{B}}} val_{\mathrm{G}}(\sigma, \pi).$$

e.g. parity games – one of the players has a winning strategy

We say that G is *determined under mixed strategies* if

$$val^{ME}_{\mathrm{G}} = val^{MA}_{\mathrm{G}}.$$

e.g. Blackwell games are determined under mixed strategies

mixed determinacy is implied by pure determinacy

Branching games and determinacy

Branching games with regular objectives

- do not have to be determined;
- not even under mixed strategies;
- nevertheless, there are natural classes of objectives that force determinacy.

Mixed determinacy

Theorem
There is a $\{E, A, \mathcal{B}\}$*-branching game with a regular winning set being a difference of two open sets that is not determined under mixed strategies.*

... but...

Theorem
If $\mathrm{G} = \langle \mathtt{B}, L \rangle$ *is a* $\{E, A, \mathcal{N}, \mathcal{B}\}$*-branching game and* L *is an arbitrary closed subset of* $\mathrm{plays}(\mathtt{B})$ *then* G *is determined under mixed strategies.*

Mixed determinacy

Theorem
There is a $\{E, A, \mathcal{B}\}$*-branching game with a regular winning set being a difference of two open sets that is not determined under mixed strategies.*

... but...

Theorem
If $\mathrm{G} = \langle \mathtt{B}, L \rangle$ *is a* $\{E, A, \mathcal{N}, \mathcal{B}\}$*-branching game and* L *is an arbitrary closed subset of* $\mathrm{plays}(\mathtt{B})$ *then* G *is determined under mixed strategies.*

... so what about computing values?

Computing values

Computing partial values

For a branching game with regular objectives

- non-probabilistic instances can be solved,

  reduction to satisfiability of *MSO* formula

- probabilistic features yield undecidability.

  can encode probabilistic automata

The value problem

Problem (The value V of a regular $\mathcal{P}$-branching game)

- **Input** *A finitary $\mathcal{P}$-branching game G with the winning condition given by a non-deterministic tree automaton.*
- **Output** *Does $V > \frac{1}{2}$?*

$V \in \{val^A, val^{BA}, val^{MA}, val^{ME}, val^{BE}, val^E\}$
$\mathcal{P} \subseteq \{A, E, \mathcal{N}, \mathcal{B}\}$

Non-stochastic games

Theorem

- *The value* val^E *problem of a regular* $\{A, E, \mathcal{B}\}$*-branching game is 2-EXP-complete,*
- *the value* val^A *problem of a regular* $\{A, E, \mathcal{B}\}$*-branching game is EXP-complete.*

upper bounds by reduction to MSO;
lack of symmetry caused by non-determinism
and nondeterministic automaton on the input

Stochastic games

Theorem
The value V problem of a regular $\{P, \mathcal{N}, \mathcal{B}\}$-branching game is undecidable.

$V \in \{val^{A}, val^{BA}, val^{MA}, val^{ME}, val^{BE}, val^{E}\}$
$P \in \{A, E\}$

Proof idea : reduce the *non-emptiness of probabilistic automata* to the value problem.

Theorem
The value V problem of a regular $\{E, A, \mathcal{B}\}$-branching game is undecidable.

$V \in \{val^{BA}, val^{MA}, val^{ME}, val^{BE}\}$

Proof idea: two players can simulate stochastic positions

Conclusions and future work

Conclusions:

- Branching games with regular objectives are not determined under mixed strategies
- but restricting the class of objectives can force determinacy.
- Non-stochastic computational problems can be solved in doubly-exponential time,
- stochastic versions of those problems are undecidable.

Possible directions of future work:

- Characterise the class of regular objectives that admit determinacy.
- Characterise the class of regular objectives that admit mixed determinacy.
- Find a class of objectives for which partial values are computable.

interesting special case – how to compute measure of a regular language $\mu(L)$

Matching pennies – different partial values

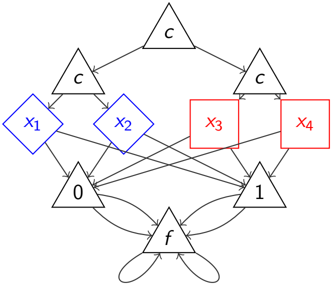

$$L \stackrel{\text{def}}{=} \{t \in \mathrm{plays}(\mathtt{B}) \mid x_1(t) = x_2(t) = x_3(t) = x_4(t) \;\lor\; x_3(t) \neq x_4(t)\}$$

Partial values:

$$val_{\mathrm{G}}^{A} = 1 \geq val_{\mathrm{G}}^{BA} = \frac{3}{4} \geq val_{\mathrm{G}}^{MA} = \frac{1}{2} = val_{\mathrm{G}}^{ME} \geq val_{\mathrm{G}}^{BE} = \frac{1}{4} \geq val_{\mathrm{G}}^{E} = 0.$$

Undecidability proof, idea

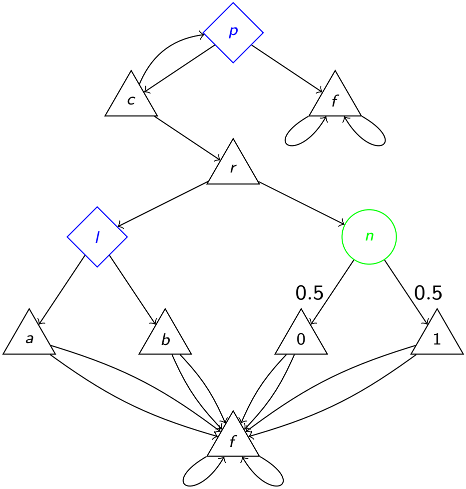

We reduce *non-emptiness of a probabilistic automaton*.
For a tree $t \in \mathrm{plays}(\mathtt{B})$

- the $\{a + b\}^* f^\omega$ word w encodes the input of the automaton;
- the $\{0 + 1\}^* f^\omega$ word r encodes a run on w.

The objective of the game accepts a tree only if the run r is accepting on the word w.

Pure determinacy

Theorem

If $\mathrm{G} = \langle \mathtt{B}, L \rangle$ *is a* $\{E, A, \mathcal{B}\}$*-branching game and* L *is a language defined by game automata then* G *is determined under pure strategies.*

effective reduction to parity games

can be extended to stochastic version

Mixed indeterminacy – a branching example

Board:

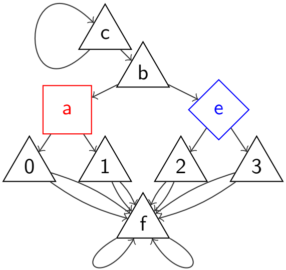

Language:

$L \stackrel{\text{def}}{=} \{t \in \mathrm{plays}(\mathtt{B}) \mid$
$e(t) < \infty$ and not $(e(t) < a(t) < \infty)\}$.

where $e(t)$ $[a(t)]$ is the depth of the first occurrence of 3 [1]

The games simulates "pick a number" game.

Mixed determinacy for closed objectives.

Theorem

If $\mathrm{G} = \langle \mathtt{B}, L \rangle$ *is a* $\{E, A, \mathcal{N}, \mathcal{B}\}$*-branching game and* L *is an arbitrary closed subset of* $\mathrm{plays}(\mathtt{B})$ *then* G *is determined under mixed strategies.*

Mixed determinacy for closed objectives.

Theorem

If $\mathrm{G} = \langle \mathtt{B}, L \rangle$ *is a* $\{E, A, \mathcal{N}, \mathcal{B}\}$*-branching game and* L *is an arbitrary closed subset of* $\mathrm{plays}(\mathtt{B})$ *then* G *is determined under mixed strategies.*

proof by an application of *Glicksberg's minimax theorem*

Useful theorem

Theorem (Glicksberg's minimax theorem)

Let A, B be compact metrisable spaces and $f : A \times B \to \mathbb{R}$ be an upper semi-continuous function. Then the following holds

$$\sup_{\mu} \inf_{\nu} \int_{A,B} f(a,b)\ \mu(\mathrm{d}a)\,\nu(\mathrm{d}b) = \inf_{\nu} \sup_{\mu} \int_{A,B} f(a,b)\ \mu(\mathrm{d}a)\,\nu(\mathrm{d}b),$$

where μ, ν range over the Borel probability measures on the sets A, B respectively.

Useful theorem

Theorem (Glicksberg's minimax theorem)

Let A, B be compact metrisable spaces and $f : A \times B \to \mathbb{R}$ be an upper semi-continuous function. Then the following holds

$$\sup_{\mu} \inf_{\nu} \int_{A,B} f(a,b)\ \mu(\mathrm{d}a)\,\nu(\mathrm{d}b) = \inf_{\nu} \sup_{\mu} \int_{A,B} f(a,b)\ \mu(\mathrm{d}a)\,\nu(\mathrm{d}b),$$

where μ, ν range over the Borel probability measures on the sets A, B respectively.

Sketch:

- Both $A = \Sigma_{\mathrm{B}}^{E}$, and $B = \Sigma_{\mathrm{B}}^{A}$ are compact;
- function $eval_G : (\sigma, \pi) \mapsto t^{\sigma,\pi}$ is continuous;
- characteristic function Φ of a closed set is semi-continuous;
- μ and ν range over mixed strategies.