Positional equilibria in infinite perfect information games

Hugo Gimbert (LaBRI, CNRS), Wiesław Zielonka (LIAFA)

ACTS

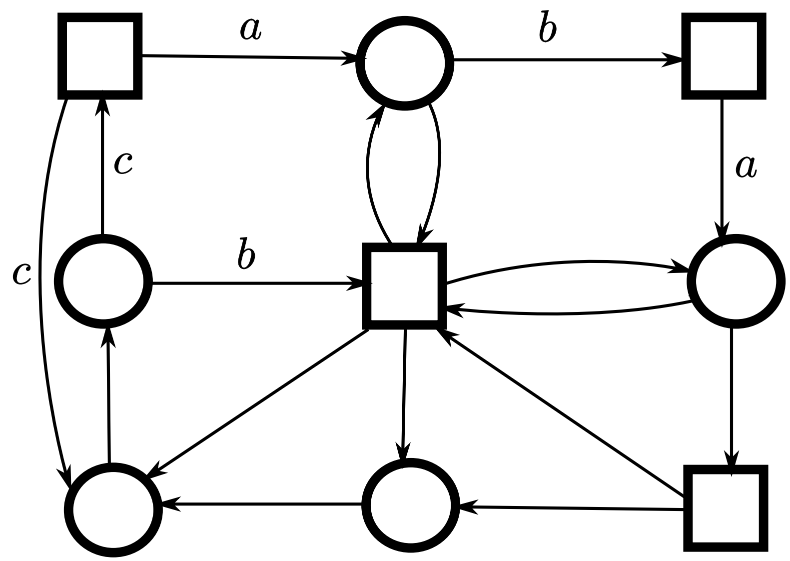
a
b
c
a
c
b

Transition systems

$$\mathcal{A} = (S, \Delta)$$

S— a finite set of states,
C— a set of colors,
$\Delta \subset S \times C \times S$ — a finite set of *actions*.

$$e = (s, c, t) \in \Delta$$

$s = \mathrm{source}(e)$ — the *source*,
$t = \mathrm{target}(e)$ — the *target*,
$c = \gamma_C(e)$ — the colour of e.

$$\Delta(s) = \{e \in \Delta \mid \mathrm{source}(e) = s\}$$

the set of actions *available at* s.

s a *sink state* if $\Delta(s) \neq \emptyset$.

Paths/Plays

$$p = e_1 e_2 e_3 \ldots \quad \forall i \geq 0, \mathrm{target}(e_i) = \mathrm{source}(e_{i+1}).$$

$$\gamma_C(p) = \gamma_C(e_1 e_2 \ldots) = \gamma_C(e_1)\gamma_C(e_2)\ldots.$$

Arenas

$$\mathcal{A} = (S, \Delta, \pi)$$

- (S, Δ) – a transition system without sink states,
- $\pi : S \to \{\mathrm{Min}, \mathrm{Max}\}$ is a mapping designating for each state s the player $\pi(s)$ controlling s.

Outcomes

An outcome of an infinite play p is the $\gamma_C(p)$.

The set of outcomes

$$\mathcal{O}(C) = \bigcup_{\substack{B \subset C \\ B \text{ finite nonempty}}} B^\omega.$$

Preference relation

A binary relation $\sqsupseteq$ over the set $\mathcal{O}(C)$ of outcomes

- reflexive, i.e. $u \sqsupseteq u$, for all $u \in \mathcal{O}(C)$,
- transitive, i.e. $u \sqsupseteq v$ and $v \sqsupseteq w$ imply $u \sqsupseteq w$, for $u, v, w \in \mathcal{O}(C)$ and
- total, either $u \sqsupseteq v$ or $v \sqsupseteq u$, for all $u, v \in \mathcal{O}(C)$.

A preference relation $=$ a total preorder relation over the set $\mathcal{O}(C)$ of outcomes.

Meaning

$$u \sqsupseteq v, \quad u, v \in \mathcal{O}(C).$$

u is no worse than v.

A player *strictly prefers* u to v, $u \sqsupset v$, if

$$u \sqsupseteq v \quad \text{but not} \quad v \sqsupseteq u.$$

If $u \sqsupseteq v$ and $v \sqsupseteq u$ — a player is indifferent between u and v.

$\sqsubseteq$ — the inverse of $\sqsupseteq$.

Two-person strictly antagonistic game

$$(\mathcal{A}, \sqsupseteq),$$

where $\mathcal{A}$ is an arena and $\sqsupseteq$ is a preference relation for Max the preference relation for player Min is $\sqsubseteq$.

The obvious aim of each player is to obtain the most advantageous outcome with respect to his preference relation.

Preferences versus payoff mappings

Payoff mapping

$$f : \mathcal{O}(C) \to \mathbb{R} \cup \{-\infty, +\infty\}$$

induces preference $\sqsupseteq_f$,

$$u \sqsupseteq_f v \quad \text{if} \quad f(u) \geq f(v).$$

Strategies and equilibria

$\mathcal{A} = (S, \Delta, \pi)$ – an arena.

$$S_{\mathrm{Max}} = \{s \in S \mid \pi(s) = \mathrm{Max}\}$$

states controlled by player Max

$$S_{\mathrm{Min}} = S \setminus S_{\mathrm{Max}}$$

states controlled by player Min.

A *strategy for player* $\mu \in \{\mathrm{Max}, \mathrm{Min}\}$ is a mapping

$$\sigma_\mu : \{p \in \mathscr{P}(\mathcal{A}) \mid \mathrm{target}(p) \in S_\mu\} \to \Delta,$$

such that $\sigma_\mu(p) \in \Delta(s)$, where $s = \mathrm{target}(p)$.

Plays consistent with a strategy

$p = e_0 e_1 e_2 \ldots$ is *consistent* with player μ's strategy σ_μ if, for each factorization $p = p'p''$, such that

- p'' is nonempty
- and $\text{target}(p') = \text{source}(p'')$ is controlled by player μ,

$\sigma_\mu(p')$ is the first action in p''.

Positional strategies

A *positional* (or memoryless) strategy for player μ

$$\sigma_\mu : S_\mu \to \Delta$$

such that, for all $s \in S_\mu$,

$$\sigma_\mu(s) \in \Delta(s)$$

A *strategy profile* is a pair (σ, τ) of strategies.

$$p_{\mathcal{A},s}(\sigma, \tau)$$

unique play with source s consistent with σ and τ.

Equilibria

A Nash equilibrium $(\sigma^{\#}, \tau^{\#})$ if for all states $s \in S$ and all strategies σ and τ,

$$\gamma_C(p_{\mathcal{A},s}(\sigma^{\#}, \tau)) \sqsupseteq \gamma_C(p_{\mathcal{A},s}(\sigma^{\#}, \tau^{\#})) \sqsupseteq \gamma_C(p_{\mathcal{A},s}(\sigma, \tau^{\#})) \ .$$

An equilibrium $(\sigma^{\#}, \tau^{\#})$ is said to be *positional* if the strategies $\sigma^{\#}$ and $\tau^{\#}$ are positional.

Mean-payoff games

$$C = \mathbb{R} \times \mathbb{R}_+$$

$$(r_1, t_1)(r_2, t_2)(r_3, t_3)\ldots \sqsupseteq (r'_1, t'_1)(r'_2, t'_2)(r'_3, t'_3)\ldots$$

if

$$\lim_n \frac{r_1 t_1 + r_2 t_2 + \cdots + r_n t_n}{t_1 + t_2 + \cdots + t_n} \geq \lim_n \frac{r'_1 t'_1 + r'_2 t'_2 + \cdots + r'_n t'_n}{t'_1 + t'_2 + \cdots + t'_n}$$

But then

$$1000, 1000, ..., 1000, 0^\omega \approx 0^\omega$$

Overtaking

$$(r_1, t_1)(r_2, t_2)(r_3, t_3) \ldots \sqsupseteq (r'_1, t'_1)(r'_2, t'_2)(r'_3, t'_3) \ldots$$

if

$$\exists N, \forall n > N, \quad \frac{r_1 t_1 + r_2 t_2 + \cdots + r_n t_n}{t_1 + t_2 + \cdots + t_n} \geq \frac{r'_1 t'_1 + r'_2 t'_2 + \cdots + r'_n t'_n}{t'_1 + t'_2 + \cdots + t'_n}$$

Weighted limits

$$C = \mathbb{R}, \quad \alpha \in [0,1]$$

$$f_\alpha(r_1 r_2 r_3 \ldots) = \alpha \cdot \limsup_i r_i + (1-\alpha) \cdot \liminf_i r_i$$

Extended preference relation and $\succeq$- equilibria

The *extended preference relation* $\succeq$ is defined as follows:

$$\text{for } x, y \in \mathcal{O}(C), \quad x \succeq y \quad \text{if} \quad \forall u \in C^*, ux \sqsupseteq uy.$$

Obviously, if $x \succeq y$ then $x \sqsupseteq y$.

$\succeq$ is transitive and reflexive, but maybe not total.

$\succeq$-equilibria

A strategy profile $(\sigma^{\#}, \tau^{\#})$ is a $\succeq$-*equilibrium* if for all strategies σ, τ

$$\gamma_C(p_{\mathcal{A},s}(\sigma^{\#}, \tau)) \succeq \gamma_C(p_{\mathcal{A},s}(\sigma^{\#}, \tau^{\#})) \succeq \gamma_C(p_{\mathcal{A},s}(\sigma, \tau^{\#})) \enspace .$$

Adherence operator

$$[\![\ ]\!] : 2^{C^*} \to 2^{C^\omega}$$

For $L \subseteq C^*$,

$$[\![L]\!] = \{u \in C^\omega \mid \mathrm{Pref}(u) \subset \mathrm{Pref}(L)\} \ .$$

Exercise

$[\![a^*b]\!] = \{a^\omega\}$

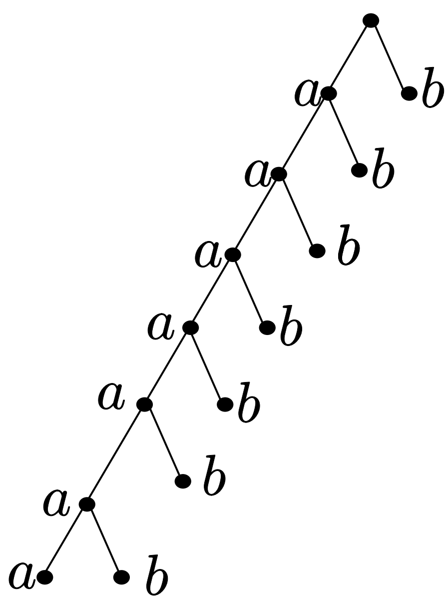

Why adherence?

$\mathcal{A} = (S, \Delta)$ an arena. Then

$$L_s^\omega(\mathcal{A}) = [\![L_s(\mathcal{A})]\!]$$

If $L \in \mathrm{Rec}(C^*)$ then $L_s^\omega(\mathcal{A}) = [\![L]\!]$ for some arena $\mathcal{A}$.

Properties of the adherence

Lemma 1. *Let* $L, M \subseteq C^*$ *be finitely generated and* $u \in C^*$*. Then*

$$[\![uL]\!] = u\,[\![L]\!]\,, \tag{1}$$

$$[\![\mathrm{Pref}([\![M]\!])]\!] = [\![M]\!]\,, \tag{2}$$

$$[\![L \cup M]\!] = [\![L]\!] \cup [\![M]\!]\,, \tag{3}$$

$$[\![LM]\!] = [\![L]\!] \cup L\,[\![M]\!]\,, \tag{4}$$

$$[\![L^*]\!] = (L \setminus \epsilon)^\omega \cup L^*\,[\![L]\!] \quad . \tag{5}$$

Conditions for positional equilibria

Let $u \in \mathcal{O}(C)$ and $X \subset \mathcal{O}(C)$.

Notation.

$$u \succeq X$$

if, for all $x \in X$, $u \succeq x$.

Ultimately periodic infinite words

Let $u, w \in C^*$ and $v \in C^+$.

An infinite word of the form

$$uv^\omega,$$

is called *ultimately periodic*.

Simple periodic languages

Let $u, w \in C^*$ and $v \in C^+$.

$$uv^*$$

Note

$$[\![uv^*]\!] = \{uv^\omega\}$$

Union selection.

$\succeq$ satisfies *union selection* condition if, for all ultimately periodic words $u_1 u_2^\omega$ and $v_1 v_2^\omega$, either

$$u_1 u_2^\omega \succeq v_1 v_2^\omega$$

or

$$v_1 v_2^\omega \succeq u_1 u_2^\omega$$

We can rewrite this condition as

$$\exists x \in \{u_1 u_2^\omega, v_1 v_2^\omega\}, \quad x \succeq [\![u_1 u_2^* \cup v_1 v_2^*]\!] \,.$$

Product selection.

We say that $\succeq$ satisfies *product selection* condition for player Max if, for all $u, v, w, z \in C^*$ such that $|v| > 0$ and $|w| > 0$,

$$\exists x \in \{uv^\omega, uwz^\omega\}, \quad x \succeq [\![uv^*wz^*]\!] .$$

Note that

$$\{uv^\omega, uwz^\omega\} \subset uv^\omega \cup uv^*wz^\omega = [\![uv^*wz^*]\!] .$$

Star selection.

$\succeq$ satisfies *star selection* condition for player Max if for each nonempty language $L \in \mathrm{Rec}(C^+)$

$$\exists x \in [\![L]\!] \cup \{u^\omega \mid u \in L\}, \quad x \succeq [\![L^*]\!] \,.$$

Note

$$[\![L]\!] \cup \{u^\omega \mid u \in L\} \subset [\![L]\!] \cup L^\omega = [\![L^*]\!]$$

Remark

If $\succeq$ satisfies all three selection conditions then for each $L \in \mathrm{Rec}(C)$, if $[\![L]\!] \neq \emptyset$ then

$$\exists uv^\omega \in [\![L]\!] \quad \text{such that} \quad uv^\omega \succeq [\![L]\!]$$

One player Max games

$\succeq$ satisfies all three selection conditions
if and only if
one-player Max games have optimal positional strategies for player Max.

Dual conditions.

$$\succeq \quad \leftrightarrow \quad \preceq .$$

Main result

Theorem 2. *Let* $\sqsupseteq$ *be a preference relation over* $\mathcal{O}(C)$ *and let* $\succeq$ *be the corresponding extended preference relation. The following conditions are equivalent:*

(1) There exist positional equilibria for all games $(\mathcal{A}, \sqsupseteq)$ *over finite arenas.*

(2) There exist positional $\succeq$*-equilibria for all games* $(\mathcal{A}, \sqsupseteq)$ *over finite arenas.*

(3) $\succeq$ *satisfies union selection, product selection and star selection conditions for player* Max *and player* Min.

(4) For all one-player games $(\mathcal{A}, \sqsupseteq)$ *the player controlling the arena* $\mathcal{A}$ *has an optimal positional strategy.*

(5) For all one-player games $(\mathcal{A}, \sqsupseteq)$ *the player controlling the arena* $\mathcal{A}$ *has a* $\succeq$*-optimal positional strategy.*