Secure Equilibria in Weighted Games

Véronique Bruyère (UMONS) Noémie Meunier (UMONS)
Jean-François Raskin (ULB)

April 2014

Topic

Reactive systems :

- modeled by two-player zero-sum games
- the system against the environment
- synthesis of a controller = construction of a winning strategy

Several extensions :

- from Boolean objectives to quantitative objectives
- from two-player zero-sum games to multi-player non zero-sum games
- construction of an equilibrium

Topic

Our work :

- two-player non zero-sum weighted games
- measures : sup, inf, lim sup, lim inf, mean-payoff, and discounted sum
- study of secure equilibria

Related work :

- Secure equilibria introduced in [CHJ06] for qualitative objectives.
- Different kinds of equilibria (Nash, secure, subgame secure perfect, and subgame perfect) for quantitative reachability objectives [BBD10, BBDG12, KLST12].
- General results for the existence of Nash equilibria in a large class of weighted games [BDS13].
- Existence and constraint existence of Nash equilibria in weighted games for the mean-payoff measure [UW11].
- . . .

Two player weighted games

Definition

$\mathcal{G} = (V, V_1, V_2, E, r, \text{Payoff})$:

- (V, E) finite graph, with V_1, V_2 a partition of V,
- $r = (r_1, r_2)$ weight function such that $r_i : E \to \mathbb{Q}$,
- $\text{Payoff} = (\text{Payoff}_1, \text{Payoff}_2)$ payoff function such that $\text{Payoff}_i : V^\omega \to \mathbb{R}$.

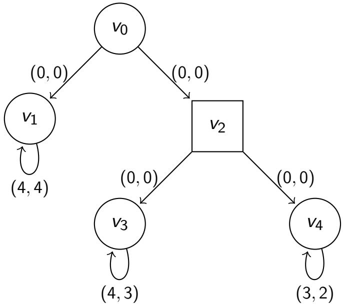

Strategy $\sigma : V^* V_i \to V$ for player i.
Finite-memory strategy. *Positional* strategy $\sigma : V_i \to V$.

Strategy profile (σ_1, σ_2) with *outcome* $\langle \sigma_1, \sigma_2 \rangle_{v_0}$ from initial vertex v_0.

Payoff functions

The payoff function Payoff is one of the payoffs in {Inf, Sup, LimInf, LimSup, InfMP, SupMP, Disc^λ for $\lambda \in]0,1[$}, where for $\rho \in V^\omega$:

- $\text{Inf}_i(\rho) = \inf_{n\in\mathbb{N}} r_i(\rho_n, \rho_{n+1})$,
- $\text{LimInf}_i(\rho) = \liminf_{n\to\infty} r_i(\rho_n, \rho_{n+1})$,
- $\text{InfMP}_i(\rho) = \liminf_{n\to\infty} \frac{1}{n} \sum_{k=0}^{n-1} r_i(\rho_k, \rho_{k+1})$,
- $\text{Disc}_i^\lambda(\rho) = (1-\lambda) \cdot \sum_{n=0}^{\infty} \lambda^n r_i(\rho_n, \rho_{n+1})$,

Similar definitions for $\text{Sup}_i(\rho), \text{LimSup}_i(\rho)$, and $\text{SupMP}_i(\rho)$.

Equilibria

Non zero-sum game.

Objective of each player : maximize his own payoff.

Definition

The strategy profile (σ_1, σ_2) is a *Nash equilibrium* from vertex v_0 if, for each strategy σ'_i, $i \in \{1, 2\}$,

$$\begin{aligned} \text{Payoff}_1(\langle \sigma'_1, \sigma_2 \rangle_{v_0}) &\leq \text{Payoff}_1(\langle \sigma_1, \sigma_2 \rangle_{v_0}), \\ \text{Payoff}_2(\langle \sigma_1, \sigma'_2 \rangle_{v_0}) &\leq \text{Payoff}_2(\langle \sigma_1, \sigma_2 \rangle_{v_0}). \end{aligned}$$

No player has an incentive to deviate from his strategy.

Equilibria

Objective of each player : first maximize his own payoff, and then minimize the payoff of the other player [CHJ06].

Definition

The strategy profile (σ_1, σ_2) is a *secure equilibrium* from vertex v_0 if, for each strategy σ'_i, $i \in \{1, 2\}$,

$$\begin{aligned} \text{Payoff}(\langle \sigma'_1, \sigma_2 \rangle_{v_0}) &\precsim_1 \text{Payoff}(\langle \sigma_1, \sigma_2 \rangle_{v_0}), \\ \text{Payoff}(\langle \sigma_1, \sigma'_2 \rangle_{v_0}) &\precsim_2 \text{Payoff}(\langle \sigma_1, \sigma_2 \rangle_{v_0}). \end{aligned}$$

No player has an incentive to deviate from his strategy w.r.t. his lexicographic order.

$$\begin{aligned} (x'_1, x'_2) \precsim_1 (x_1, x_2) &\iff (x'_1 < x_1) \vee (x'_1 = x_1 \wedge x'_2 \geq x_2), \\ (x'_1, x'_2) \precsim_2 (x_1, x_2) &\iff (x'_2 < x_2) \vee (x'_2 = x_2 \wedge x'_1 \geq x_1). \end{aligned}$$

Example of an InfMP weighted game

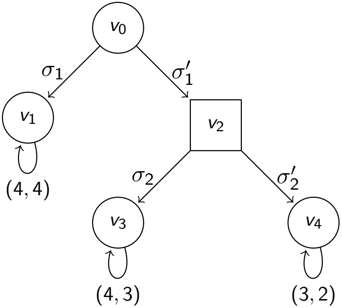

Example of an InfMP weighted game

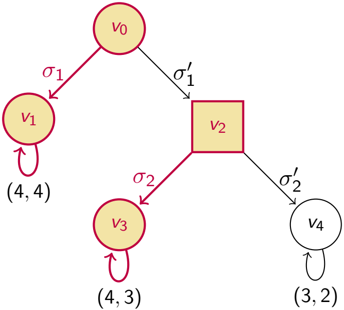

- (σ_1, σ_2) : Nash equilibrium with payoff $(4, 4)$, but not a secure equilibrium since $(4, 4) \prec_1 (4, 3)$

Example of an InfMP weighted game

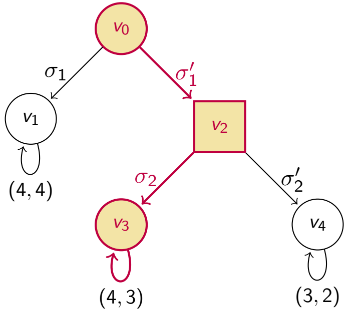

- (σ_1, σ_2) : Nash equilibrium with payoff $(4, 4)$, but not a secure equilibrium since $(4, 4) \prec_1 (4, 3)$
- (σ_1', σ_2) : secure equilibrium with payoff $(4, 3)$

Studied problems

Let $\mathcal{G}$ be a weighted game and v_0 be an initial vertex.

Problem

1. *Does there exist a (finite-memory) secure equilibrium in $\mathcal{G}$ from v_0 ?*
2. *What is the complexity of constructing a (finite-memory) secure equilibrium if one exists ?*
3. *Given two thresholds $\mu, \nu \in (\mathbb{Q} \cup \{\pm\infty\})^2$, is it decidable whether there exists a secure equilibrium (σ_1, σ_2) such that*

$$\mu \leq \mathrm{Payoff}(\langle \sigma_1, \sigma_2 \rangle_{v_0}) \leq \nu$$

i.e. $\mu_i \leq \mathrm{Payoff}_i(\langle \sigma_1, \sigma_2 \rangle_{v_0}) \leq \nu_i$ *for* $i \in \{1, 2\}$ *?*

Existence

General approach to construct a finite-memory secure equilibrium.

We need two tools :

- lexicographic payoff games
- prefix-linear payoff functions

Existence

General approach to construct a finite-memory secure equilibrium.

We need two tools :

- lexicographic payoff games
- prefix-linear payoff functions

Theorem

Let $\mathcal{G}$ be a weighted game and v_0 an initial vertex. If for each $i \in \{1, 2\}$,

- $\mathcal{G}^{\preccurlyeq_i}$ *is uniformly-determined,*
- Payoff_i *is prefix-linear,*

then there exists a finite-memory secure equilibrium in $\mathcal{G}$ from v_0 with memory size at most $|V| + 2$.

Lexicographic payoff games

With the non zero-sum game $\mathcal{G} = (V, V_1, V_2, E, r, \text{Payoff})$, we associate two zero-sum *lexicographic payoff games*

$$\mathcal{G}^{\preccurlyeq_1} = (V, V_1, V_2, E, r, \text{Payoff}, \preccurlyeq_1),$$
$$\mathcal{G}^{\preccurlyeq_2} = (V, V_1, V_2, E, r, \text{Payoff}, \preccurlyeq_2).$$

In $\mathcal{G}^{\preccurlyeq_1}$, player 1 wants to maximize $\text{Payoff}(\rho)$ w.r.t. $\preccurlyeq_1$, while player 2 wants to minimize it. Symmetrically for $\mathcal{G}^{\preccurlyeq_2}$.

Lexicographic payoff games

Let $\mathcal{G}^{\preccurlyeq_1}$ be a lexicographic payoff game.

Definition

- Let v be a vertex.

$$\overline{\mathsf{Val}}(v) = \inf_{\sigma_2 \in \Sigma_2} \sup_{\sigma_1 \in \Sigma_1} \mathsf{Payoff}(\langle \sigma_1, \sigma_2 \rangle_v),$$
$$\underline{\mathsf{Val}}(v) = \sup_{\sigma_1 \in \Sigma_1} \inf_{\sigma_2 \in \Sigma_2} \mathsf{Payoff}(\langle \sigma_1, \sigma_2 \rangle_v).$$

- $\mathcal{G}^{\preccurlyeq_1}$ is *determined* if, for every $v \in V$, we have $\overline{\mathsf{Val}}(v) = \underline{\mathsf{Val}}(v)$. Value $\mathsf{Val}(v)$ and optimal strategies.
- $\mathcal{G}^{\preccurlyeq_1}$ is *positionally-determined* if it is determined and has positional optimal strategies for both players and all vertices.
- $\mathcal{G}^{\preccurlyeq_1}$ is *uniformly-determined* if the positional optimal strategies can be chosen globally independently of the initial vertex v.

Example of an InfMP lexicographic payoff game

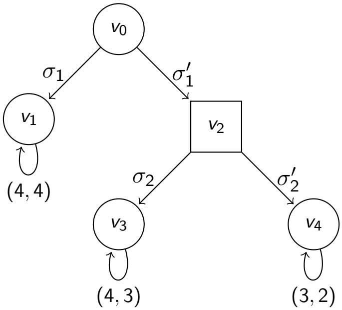

- lexicographic order $\preccurlyeq_1$ $\qquad (3,2) \preccurlyeq_1 (4,4) \preccurlyeq_1 (4,3)$

Example of an InfMP lexicographic payoff game

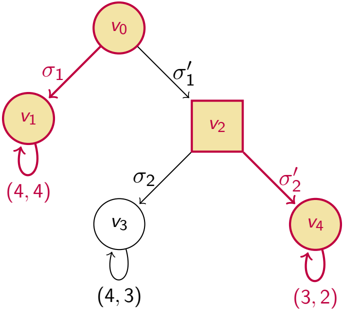

- lexicographic order $\preccurlyeq_1$ $\qquad (3, 2) \preccurlyeq_1 (4, 4) \preccurlyeq_1 (4, 3)$
- $\mathrm{Val}(v_0) = (4, 4), \mathrm{Val}(v_2) = (3, 2)$

Determinacy results

Theorem

Let $\mathcal{G}^{\preccurlyeq_1}$ *be a lexicographic payoff game. Then* $\mathcal{G}^{\preccurlyeq_1}$ *is uniformly-determined (resp. positionally-determined) for payoff functions* InfMP, SupMP, LimInf, LimSup *and* Disc^{λ} *(resp.* Inf *and* Sup*).*

Proof for

- InfMP, SupMP : inspired from [BCHJ09], however with several nontrivial adaptations ;
- LimInf, LimSup, Inf, Sup : based on adequate reductions to classical ω-regular games ;
- Disc^{λ} : based on one-dimensional discounted games.

Determinacy results (other interesting results)

Theorem

Let $\mathcal{G}^{\preccurlyeq_1}$ be a lexicographic payoff game.

1. *Let v be a vertex and $(\alpha, \beta) \in \mathbb{Q}^2$ be a pair of rationals. One can decide whether $(\alpha, \beta) \preccurlyeq_1 \mathsf{Val}(v)$.*
2. *One can compute the (rational) value $\mathsf{Val}(v)$ of each vertex v.*
3. *One can compute optimal strategies for both players in $\mathcal{G}^{\preccurlyeq_1}$.*

	1.	**2.**	**3.**
InfMP, SupMP	NP ∩ co-NP	*pseudo-polynomial*	*pseudo-polynomial*
LimInf, LimSup	P-*complete*	*polynomial*	*polynomial*
Inf, Sup	P-*complete*	*polynomial*	*polynomial*
Disc^{λ}	NP ∩ co-NP	*pseudo-polynomial*	*pseudo-polynomial*

Prefix-linear payoff function

Definition

Function Payoff_i is *prefix-linear* if, $\forall h \in V^+$, $\exists a \in \mathbb{R}$, $\exists b \in \mathbb{R}^+$, $\forall \rho \in V^\omega$

$$\text{Payoff}_i(h\rho) = a + b \cdot \text{Payoff}_i(\rho).$$

In particular, it is *prefix-independent* if $\text{Payoff}_i(h\rho) = \text{Payoff}_i(\rho)$.

Prefix-linear payoff function

Definition

Function Payoff_i is *prefix-linear* if, $\forall h \in V^+$, $\exists a \in \mathbb{R}$, $\exists b \in \mathbb{R}^+$, $\forall \rho \in V^\omega$

$$\text{Payoff}_i(h\rho) = a + b \cdot \text{Payoff}_i(\rho).$$

In particular, it is *prefix-independent* if $\text{Payoff}_i(h\rho) = \text{Payoff}_i(\rho)$.

- LimInf_i, LimSup_i, InfMP_i and SupMP_i : prefix-independent.
- Disc_i^λ : not prefix-independent, but prefix-linear.
- Inf_i, Sup_i : not prefix-linear.

Existence and construction

Theorem

Let $\mathcal{G}$ be a weighted game and v_0 an initial vertex. If for each $i \in \{1, 2\}$,

- Payoff_i *is prefix-linear,*
- $\mathcal{G}^{\preccurlyeq_i}$ *is uniformly-determined,*

then there exists a finite-memory secure equilibrium in $\mathcal{G}$ from v_0 with memory size at most $|V| + 2$.
Moreover, if the uniform optimal strategies can be computed in C time for each $\mathcal{G}^{\preccurlyeq_i}$, then this secure equilibrium can also be constructed in C time.

Existence and construction

Theorem

Let $\mathcal{G}$ be a weighted game and v_0 an initial vertex. If for each $i \in \{1, 2\}$,

- Payoff_i *is prefix-linear,*
- $\mathcal{G}^{\preccurlyeq_i}$ *is uniformly-determined,*

then there exists a finite-memory secure equilibrium in $\mathcal{G}$ from v_0 with memory size at most $|V| + 2$.
Moreover, if the uniform optimal strategies can be computed in C time for each $\mathcal{G}^{\preccurlyeq_i}$, then this secure equilibrium can also be constructed in C time.

Adaptation from [BDS13] : for Nash equilibria and multiple players.

Construction : each player i plays according to his optimal strategy in $\mathcal{G}^{\preccurlyeq_i}$. As soon as player j deviates, the other player shifts to $\mathcal{G}^{\preccurlyeq_j}$ and plays according to his optimal strategy in $\mathcal{G}^{\preccurlyeq_j}$.

Existence and construction

Recall that for InfMP, SupMP, LimInf, LimSup, and Disc^λ payoffs,

- each Payoff_i is prefix-linear,
- each $\mathcal{G}^{\preccurlyeq_i}$ is uniformly-determined.

Therefore we get :

Corollary

Existence of a finite-memory secure equilibrium in $\mathcal{G}$ from v_0.

	Memory size	*Construction time*
InfMP, SupMP	$\lvert V\rvert + 2$	*pseudo-polynomial*
LimInf, LimSup	$\lvert V\rvert + 2$	*polynomial*
Disc^λ	$\lvert V\rvert + 2$	*pseudo-polynomial*

Existence and construction

Recall that for InfMP, SupMP, LimInf, LimSup, and Disc^λ payoffs,

- each Payoff_i is prefix-linear,
- each $\mathcal{G}^{\preccurlyeq_i}$ is uniformly-determined.

Therefore we get :

Corollary

Existence of a finite-memory secure equilibrium in $\mathcal{G}$ from v_0.

	Memory size	*Construction time*
InfMP, SupMP	$\|V\| + 2$	*pseudo-polynomial*
LimInf, LimSup	$\|V\| + 2$	*polynomial*
Inf, Sup	$\|V\| \cdot \|E\|^2 + 3$	*polynomial*
Disc^λ	$\|V\| + 2$	*pseudo-polynomial*

Constrained existence

Let $\mu, \nu \in (\mathbb{Q} \cup \{\pm\infty\})^2$ be two thresholds.

General approach to decide the existence of a secure equilibrium such that $\mu \leq \text{Payoff}(\langle \sigma_1, \sigma_2 \rangle_{v_0}) \leq \nu$, using the two previous tools

and a solution to the next problem :

Problem (*)

Let $G = (V, E, v_0, r, \text{Val})$ *be a weighted graph with an initial vertex* v_0*, and a value function* $\text{Val} = (\text{Val}^1, \text{Val}^2)$*, with* $\text{Val}^i : V \to \mathbb{Q}$.
Is it decidable whether there exists an infinite path ρ *from* v_0 *such that*

- $\forall k \geq 0, \forall i \in \{1, 2\}, \text{Val}^i(\rho_k) \preccurlyeq_i \text{Payoff}(\rho_k \rho_{k+1} \ldots)$*, and*
- $\mu \leq \text{Payoff}(\rho) \leq \nu$ *?*

Constrained existence

Theorem

Let $\mathcal{G}$ be a weighted game, and $\mu, \nu \in (\mathbb{Q} \cup \{\pm\infty\})^2$ two thresholds. If

- *for each $i \in \{1, 2\}$, Payoff_i is prefix-linear, and $\mathcal{G}^{\preccurlyeq_i}$ is uniformly-determined with computable values,*
- *Problem (*) is decidable for the graph G constructed from $\mathcal{G}$ such that both functions Val^i, $i \in \{1, 2\}$, are defined from the values $\mathrm{Val}^i(v)$, $v \in V$, in $\mathcal{G}^{\preccurlyeq_i}$.*

One can decide whether there exists a secure equilibrium (σ_1, σ_2) in $\mathcal{G}$ from an initial vertex v_0 such that $\mu \leq \mathrm{Payoff}(\langle\sigma_1, \sigma_2\rangle_{v_0}) \leq \nu$.

Lemma

Let ρ be a play in $\mathcal{G}$ from v_0. Then ρ is the outcome of a secure equilibrium iff $\forall k \geq 0, \forall i \in \{1, 2\}, \mathrm{Val}^i(\rho_k) \preccurlyeq_i \mathrm{Payoff}(\rho_k \rho_{k+1} \ldots)$.

Solution to Problem (*)

Theorem

Problem (*) is decidable in polynomial time for InfMP, SupMP, LimInf *and* LimSup *payoffs.*

Corollary

Let $\mathcal{G}$ be a weighted game and v_0 be an initial vertex. One can decide whether there exists a secure equilibrium (σ_1, σ_2) such that $\mu \leq \mathrm{Payoff}(\langle \sigma_1, \sigma_2 \rangle_{v_0}) \leq \nu$.

	Complexity
InfMP, SupMP	NP $\cap$ co-NP
LimInf, LimSup	P

Solution to Problem (*)

Theorem

Problem (*) is decidable in polynomial time for InfMP, SupMP, LimInf *and* LimSup *payoffs.*

Corollary

Let $\mathcal{G}$ be a weighted game and v_0 be an initial vertex. One can decide whether there exists a secure equilibrium (σ_1, σ_2) such that $\mu \leq \text{Payoff}(\langle \sigma_1, \sigma_2 \rangle_{v_0}) \leq \nu$.

	Complexity
InfMP, SupMP	NP ∩ co-NP
LimInf, LimSup	P
Inf, Sup	P
Disc^{λ}	open

Conclusion

- General hypotheses for the existence and constrained existence of secure equilibria in two-player weighted games
  - as corollary, solutions for InfMP, SupMP, LimInf, LimSup, Inf, Sup and Disc^λ payoffs
  - except for the constrained existence for Disc^λ payoff
- Interesting results about the zero-sum games $\mathcal{G}^{\preccurlyeq_1}$, $\mathcal{G}^{\preccurlyeq_2}$

Future work

Solutions proposed under general hypotheses

- Generalization of these hypotheses ?
- Other applications of these hypotheses, e.g. for mixed payoffs (like InfMP for player 1, and LimSup for player 2) ?

Other questions

- Secure equilibria in multi-player weighted games ?
- If several secure equilibria exist, selection of an interesting one (Pareto optimal, with a bounded payoff, ...) ?
- Subgame perfect and secure subgame perfect equilibria in weighted games ?
- More than one payoff for each player ?

- Thomas Brihaye, Véronique Bruyère, and Julie De Pril, *Equilibria in quantitative reachability games*, CSR, LNCS, vol. 6072, Springer, 2010, pp. 72–83.
- Thomas Brihaye, Véronique Bruyère, Julie De Pril, and Hugo Gimbert, *On subgame perfection in quantitative reachability games*, Logical Methods in Computer Science **9** (2012).
- Roderick Bloem, Krishnendu Chatterjee, Thomas A. Henzinger, and Barbara Jobstmann, *Better quality in synthesis through quantitative objectives*, CAV, LNCS, vol. 5643, Springer, 2009, pp. 140–156.
- Thomas Brihaye, Julie De Pril, and Sven Schewe, *Multiplayer cost games with simple Nash equilibria*, LFCS, LNCS, vol. 7734, Springer, 2013, pp. 59–73.
- Udi Boker and Jan Otop, personal communication, 2014.

Krishnendu Chatterjee, Vojtech Forejt, and Dominik Wojtczak, *Multi-objective discounted reward verification in graphs and mdps*, LPAR, LNCS, vol. 8312, 2013, pp. 228–242.

K. Chatterjee, T. A. Henzinger, and M. Jurdzinski, *Games with secure equilibria*, Theor. Comput. Sci. **365** (2006), 67–82.

Miroslav Klimos, Kim G. Larsen, Filip Stefanak, and Jeppe Thaarup, *Nash equilibria in concurrent priced games*, LATA, LNCS, vol. 7183, Springer, 2012, pp. 363–376.

Michael Ummels and Dominik Wojtczak, *The complexity of Nash equilibria in limit-average games*, CONCUR, LNCS, vol. 6901, Springer, 2011, pp. 482–496.

Constrained existence

The case of Disc^λ payoff is left open. Indeed the next problem [CFW13] reduces to the constrained existence of a secure equilibrium for Disc^λ payoff.

Problem (**)

Given three rational numbers a, b and t, and a rational discount factor $\lambda \in]0, 1[$, does there exist an infinite sequence $w = w_0 w_1 \ldots \in \{a, b\}^\omega$ such that $\sum_{k=0}^{\infty} w_k \lambda^k$ is equal to t ?

Problem (**) is related to other hard open problems in diverse mathematical fields [BO14].