On algorithmic aspects of infinite games played on weighted graphs

Véronique Bruyère
UMONS Belgium

MPS 2016

This talk

- Focus on algorithmic game theory for the synthesis of correct computer systems

This talk

- Focus on algorithmic game theory for the synthesis of correct computer systems
- Some classical results and recent UMONS results

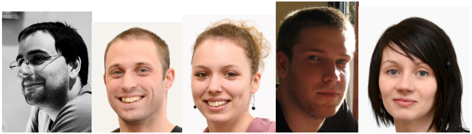

This talk

- Focus on algorithmic game theory for the synthesis of correct computer systems
- Some classical results and recent UMONS results

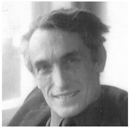


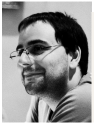


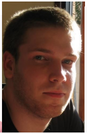


Algorithmic game theory

Reactive system

- System embedded into an uncontrollable environment
- It must satisfy some property against any behavior of the environment
- How to automatically design a correct controller for the system?

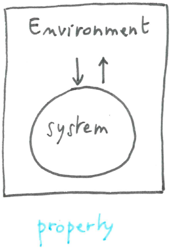

Algorithmic game theory

Reactive system

- System embedded into an uncontrollable environment
- It must satisfy some property against any behavior of the environment
- How to automatically design a correct controller for the system?

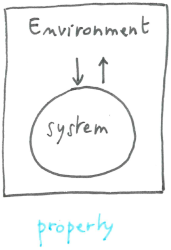

Example

- System: airplane
  and environment: weather
- The airplane must land safely in any weather conditions
- How to design a correct autopilot?

Algorithmic game theory

Reactive system

- System embedded into an uncontrollable environment
- It must satisfy some property against any behavior of the environment
- How to automatically design a correct controller for the system?

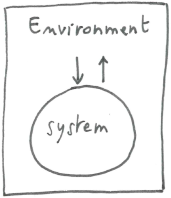

property

Modelization

- Two-player zero-sum game played on a finite directed graph
- Property $=$ objective for the system
- Synthesis of a controller $=$ construction of a winning strategy

game played on a graph

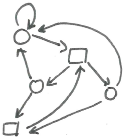

objective

Basic model

Basic model

Definition

Two-player zero-sum game $G = (V, V_1, V_2, E, v_0)$:

- (V, E) finite directed graph (with no deadlock)
- (V_1, V_2) partition of V with V_p controlled by player $p \in \{1, 2\}$
- initial vertex v_0

Paths

- Play: infinite path from v_0
  $\rho = \rho_0\rho_1 \ldots \in V^\omega$ in G
- History: prefix h of a play
- Unravelling of G: infinite tree of all paths from v_0

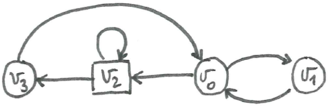

Player 1 ◯, Player 2 □

Basic model

Objective: set $\Omega \subseteq V^\omega$ of plays

Zero-sum game:

- objective Ω for player 1
- opposite objective $V^\omega \setminus \Omega$ for player 2

Basic model

Objective: set $\Omega \subseteq V^\omega$ of plays

Zero-sum game:

- objective Ω for player 1
- opposite objective $V^\omega \setminus \Omega$ for player 2

Definition

Given a set $U \subseteq V$, classical qualitative objectives are:

- Reachability objective: visit a vertex of U at least once
- Büchi objective: visit a vertex of U infinitely often

Basic model

Objective: set $\Omega \subseteq V^\omega$ of plays

Zero-sum game:

- objective Ω for player 1
- opposite objective $V^\omega \setminus \Omega$ for player 2

Definition

Given a set $U \subseteq V$, classical qualitative objectives are:

- Reachability objective: visit a vertex of U at least once
- Büchi objective: visit a vertex of U infinitely often

Given a coloring $c : V \to \{0, 1, \ldots, C\}$

- Parity objective: the maximum color seen infinitely often is even

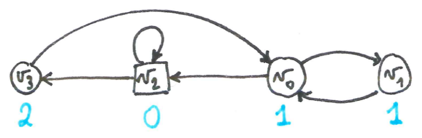

Strategies

Strategy for player p:
function $\sigma : V^* V_p \to V$ such that
$\sigma(hv) = v'$ with $(v, v') \in E$

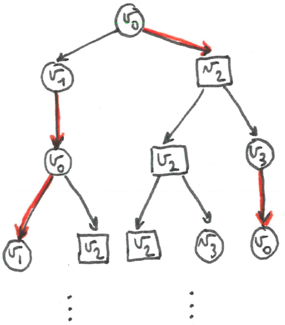

Strategies

Strategy for player p:
function $\sigma : V^* V_p \to V$ such that
$\sigma(hv) = v'$ with $(v, v') \in E$

Memoryless: when $\sigma(hv) = \sigma(v)$

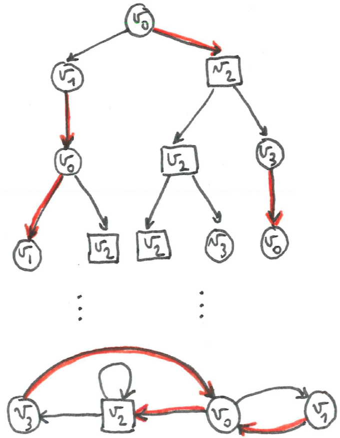

Strategies

Strategy for player p:
function $\sigma : V^* V_p \to V$ such that
$\sigma(hv) = v'$ with $(v, v') \in E$

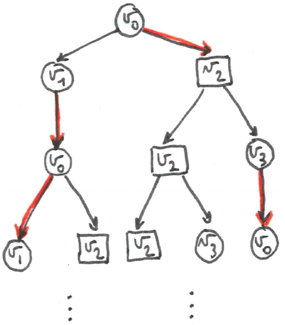

Memoryless: when $\sigma(hv) = \sigma(v)$

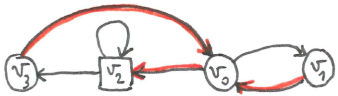

Finite-memory: when σ is recorded by a finite automaton

Strategies

Winning strategy for player p: ensure his objective against any strategy of the other player

A game is determined from initial vertex v_0 when

- either player 1 is winning for Ω from v_0
- or player 2 is winning for $V^\omega \setminus \Omega$ from v_0

Strategies

Winning strategy for player p: ensure his objective against any strategy of the other player

A game is determined from initial vertex v_0 when

- either player 1 is winning for Ω from v_0
- or player 2 is winning for $V^\omega \setminus \Omega$ from v_0

Example

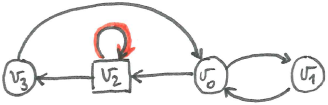

- Objective of player 1: visit v_3
- Player 1 is (trivially) winning from v_3
- Player 2 is winning from v_0, v_1, and v_2
  Memoryless strategy: looping on v_2

Strategies

Winning strategy for player p: ensure his objective against any strategy of the other player

A game is determined from initial vertex v_0 when

- either player 1 is winning for Ω from v_0
- or player 2 is winning for $V^\omega \setminus \Omega$ from v_0

Another example

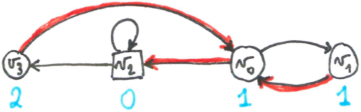

Parity game: Player 1 is winning from every vertex with a memoryless strategy

- Either player 2 eventually stays at v_2
  $\rightarrow$ max color seen infinitely often $= 0$
- Or he infinitely often visits v_3
  $\rightarrow$ max color seen infinitely often $= 2$

Martin's theorem

Theorem [Mar75]

Every game with Borel objectives is determined

Martin's theorem

Theorem [Mar75]

Every game with Borel objectives is determined

- Need of the axiom of choice to exhibit a non-determined game
- No information about the winning strategies

Martin's theorem

Theorem [Mar75]

Every game with Borel objectives is determined

- Need of the axiom of choice to exhibit a non-determined game
- No information about the winning strategies

Corollary

Every game with ω-regular objectives is determined

Martin's theorem

Theorem [Mar75]

Every game with Borel objectives is determined

- Need of the axiom of choice to exhibit a non-determined game
- No information about the winning strategies

Corollary

Every game with ω-regular objectives is determined

Algorithmic questions

- Who is the winner from initial vertex v_0?
- Complexity class of this decision problem?
- Can we construct a winning strategy for the winner?
- What kind of winning strategy? Memoryless, finite-memory?

Algorithmic results for one-player games

Classical question in automata theory: Player 1 wins iff there exists a play satisfying the objective

- Reachability objective: emptiness of automata on finite words
- Büchi objective: emptiness of automata on infinite words

Algorithmic results for one-player games

Classical question in automata theory: Player 1 wins iff there exists a play satisfying the objective

- Reachability objective: emptiness of automata on finite words
- Büchi objective: emptiness of automata on infinite words

Finite-memory winning strategy iff the winning play is eventually periodic

Algorithmic results for one-player games

Classical question in automata theory: Player 1 wins iff there exists a play satisfying the objective

- Reachability objective: emptiness of automata on finite words
- Büchi objective: emptiness of automata on infinite words

Finite-memory winning strategy iff the winning play is eventually periodic

- reachable cycle in the graph
- reachable simple cycle for memoryless strategies

Algorithmic results for one-player games

Classical question in automata theory: Player 1 wins iff there exists a play satisfying the objective

- Reachability objective: emptiness of automata on finite words
- Büchi objective: emptiness of automata on infinite words

Finite-memory winning strategy iff the winning play is eventually periodic

- reachable cycle in the graph
- reachable simple cycle for memoryless strategies

Remark: Given a two-player game G, by fixing a strategy σ for player p, we get a one-player (infinite/finite) game G_σ

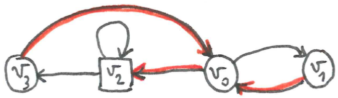

Algorithmic results for two-player games

[Bee80, Imm81, EJ91], see also [Zie98, GTW02]

	Reach	Büchi	Parity
Complexity	P-complete		NP $\cap$ co-NP
Player 1 strategy	memoryless		
Player 2 strategy	memoryless		

Algorithmic results for two-player games

[Bee80, Imm81, EJ91], see also [Zie98, GTW02]

	Reach	Büchi	Parity
Complexity	P-complete		NP $\cap$ co-NP
Player 1 strategy	memoryless		
Player 2 strategy	memoryless		

- Remember the previous examples
- More information on the proofs in the next slides

Algorithmic results for two-player games

[Bee80, Imm81, EJ91], see also [Zie98, GTW02]

	Reach	Büchi	Parity
Complexity	P-complete		NP $\cap$ co-NP
Player 1 strategy	memoryless		
Player 2 strategy	memoryless		

- Remember the previous examples
- More information on the proofs in the next slides

Major open problem: can we solve Parity games in P?

Basic model

Extension with weights on the edges

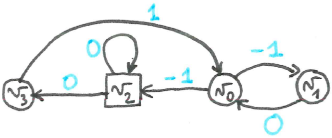

Definition

Two-player zero-sum game $G = (V, V_1, V_2, E, v_0, w)$ as before, with:

- $w : E \to \mathbb{Z}$ weight function

Basic model

Extension with weights on the edges

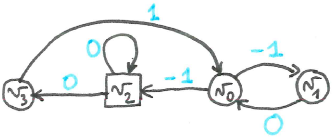

Definition

Two-player zero-sum game $G = (V, V_1, V_2, E, v_0, w)$ as before, with:

- $w : E \to \mathbb{Z}$ weight function

Classical payoff $f(\rho)$ of a play $\rho = \rho_0\rho_1\rho_2\ldots$:

- $\mathrm{Inf}(\rho) = \inf_{n\in\mathbb{N}} w(\rho_n, \rho_{n+1})$
- $\mathrm{LimInf}(\rho) = \liminf_{n\to\infty} w(\rho_n, \rho_{n+1})$
- Total-payoff $\underline{\mathrm{TP}}(\rho) = \liminf_{n\to\infty} \sum_{k=0}^{n-1} w(\rho_k, \rho_{k+1})$
- Mean-payoff $\underline{\mathrm{MP}}(\rho) = \liminf_{n\to\infty} \frac{1}{n} \sum_{k=0}^{n-1} w(\rho_k, \rho_{k+1})$

Basic model

Extension with weights on the edges

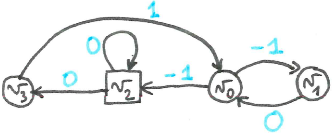

Definition

Two-player zero-sum game $G = (V, V_1, V_2, E, v_0, w)$ as before, with:

- $w : E \to \mathbb{Z}$ weight function

Classical payoff $f(\rho)$ of a play $\rho = \rho_0\rho_1\rho_2\dots$:

- $\mathrm{Inf}(\rho) = \inf_{n\in\mathbb{N}} w(\rho_n, \rho_{n+1})$
- $\mathrm{LimInf}(\rho) = \liminf_{n\to\infty} w(\rho_n, \rho_{n+1})$
- Total-payoff $\underline{\mathrm{TP}}(\rho) = \liminf_{n\to\infty} \sum_{k=0}^{n-1} w(\rho_k, \rho_{k+1})$
- Mean-payoff $\underline{\mathrm{MP}}(\rho) = \liminf_{n\to\infty} \frac{1}{n} \sum_{k=0}^{n-1} w(\rho_k, \rho_{k+1})$

Similar definitions with $\mathrm{Sup}(\rho), \mathrm{LimSup}(\rho), \overline{\mathrm{TP}}(\rho), \overline{\mathrm{MP}}(\rho)$

Quantitative objectives

Example

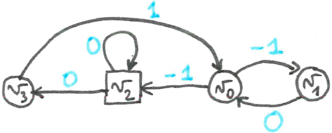

play $\rho = (v_0 v_1)^\omega$

- $\underline{\mathrm{TP}}(\rho) = \liminf_{n\to\infty} \sum_{k=0}^{n-1} w(\rho_k, \rho_{k+1})$

$(-1, -1, -2, -2, -3, -3, \dots, -n, -n, \dots) \to -\infty = \underline{\mathrm{TP}}(\rho) = \overline{\mathrm{TP}}(\rho)$

- $\underline{\mathrm{MP}}(\rho) = \liminf_{n\to\infty} \frac{1}{n} \sum_{k=0}^{n-1} w(\rho_k, \rho_{k+1})$

$(\frac{-1}{1}, \frac{-1}{2}, \frac{-2}{3}, \frac{-2}{4}, \frac{-3}{5}, \frac{-3}{6}, \dots, \frac{-n}{2n-1}, \frac{-n}{2n}, \dots) \to -\frac{1}{2} = \underline{\mathrm{MP}}(\rho) = \overline{\mathrm{MP}}(\rho)$

Quantitative objectives

Example

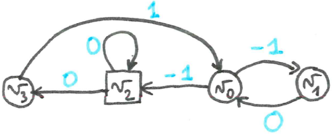

play $\rho = (v_0 v_1)^\omega$

- $\underline{\mathrm{TP}}(\rho) = \liminf_{n\to\infty} \sum_{k=0}^{n-1} w(\rho_k, \rho_{k+1})$

$$(-1, -1, -2, -2, -3, -3, \ldots, -n, -n, \ldots) \to -\infty = \underline{\mathrm{TP}}(\rho) = \overline{\mathrm{TP}}(\rho)$$

- $\underline{\mathrm{MP}}(\rho) = \liminf_{n\to\infty} \frac{1}{n} \sum_{k=0}^{n-1} w(\rho_k, \rho_{k+1})$

$$\left(\frac{-1}{1}, \frac{-1}{2}, \frac{-2}{3}, \frac{-2}{4}, \frac{-3}{5}, \frac{-3}{6}, \ldots, \frac{-n}{2n-1}, \frac{-n}{2n}, \ldots\right) \to -\frac{1}{2} = \underline{\mathrm{MP}}(\rho) = \overline{\mathrm{MP}}(\rho)$$

Lemma: If eventually periodic $\rho = hg^\omega$,
then $\underline{\mathrm{MP}}(\rho) = \overline{\mathrm{MP}}(\rho) =$ mean payoff of the cycle g

Quantitative objectives

Definition

Classical quantitative objectives are, given a threshold $\nu \in \mathbb{Q}$:

- Inf objective: ensure $\mathsf{Inf}(\rho) \geq \nu$
- Similarly for the other payoff functions LimInf, $\underline{\mathsf{TP}}$, ...

Quantitative objectives

Definition

Classical quantitative objectives are, given a threshold $\nu \in \mathbb{Q}$:

- Inf objective: ensure $\mathsf{Inf}(\rho) \geq \nu$
- Similarly for the other payoff functions LimInf, TP, ...

- Player 1 wants to maximize the payoff (interest in the highest ν)
- Player 2 wants to minimize the payoff

Quantitative objectives

Definition

Classical quantitative objectives are, given a threshold $\nu \in \mathbb{Q}$:

- Inf objective: ensure $\mathsf{Inf}(\rho) \geq \nu$
- Similarly for the other payoff functions LimInf, $\underline{\mathsf{TP}}$, ...

- Player 1 wants to maximize the payoff (interest in the highest ν)
- Player 2 wants to minimize the payoff

Corollary of Martin's Theorem

Games with such quantitative objectives are determined

- Inf, LimInf: ω-regular
- $\underline{\mathsf{TP}}$, $\underline{\mathsf{MP}}$: not ω-regular, but Borel

Algorithmic results for two-player games

[EM79, ZP96, BSV04]

	Reach	Büchi	Parity	$\underline{\mathrm{TP}}, \overline{\mathrm{TP}}$	$\underline{\mathrm{MP}}, \overline{\mathrm{MP}}$
Complexity	P-complete		NP $\cap$ co-NP		
Player 1 strategy	memoryless				
Player 2 strategy	memoryless				

Algorithmic results for two-player games

[EM79, ZP96, BSV04]

	Reach	Büchi	Parity	$\underline{\mathsf{TP}}, \overline{\mathsf{TP}}$	$\underline{\mathsf{MP}}, \overline{\mathsf{MP}}$
Complexity	P-complete		NP $\cap$ co-NP		
Player 1 strategy	memoryless				
Player 2 strategy	memoryless				

Reductions:

- Parity games $\to$ Mean-payoff games [Jur98]
- Total-payoff games $\leftrightarrow$ Mean-payoff games

Major open problem: can we solve Parity, Mean-payoff and Total-payoff games in P?

Algorithmic results

Theorem [GZ04]

Let G be a weighted game. If the payoff function f is fairly mixing, i.e.:

1. $f(\rho) \leq f(\rho') \Rightarrow f(h\rho) \leq f(h\rho')$
2. $\min\{f(\rho), f(h^\omega)\} \leq f(h\rho) \leq \max\{f(\rho), f(h^\omega)\}$
3. $\min\{f(h_0 h_2 h_4 \ldots), f(h_1 h_3 h_5 \ldots), \inf_i f(h_i^\omega)\}$
   $\leq f(h_0 h_1 h_2 h_3 \ldots) \leq \max\{f(h_0 h_2 h_4 \ldots), f(h_1 h_3 h_5 \ldots), \sup_i f(h_i^\omega)\}$

then both players have memoryless (optimal) winning strategies

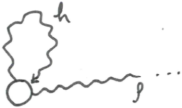

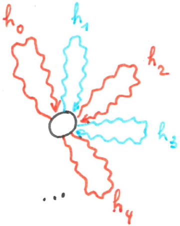

Algorithmic results

Theorem [GZ04]

Let G be a weighted game. If the payoff function f is fairly mixing, i.e.:

1. $f(\rho) \leq f(\rho') \Rightarrow f(h\rho) \leq f(h\rho')$
2. $\min\{f(\rho), f(h^\omega)\} \leq f(h\rho) \leq \max\{f(\rho), f(h^\omega)\}$
3. $\min\{f(h_0 h_2 h_4 \dots), f(h_1 h_3 h_5 \dots), \inf_i f(h_i^\omega)\}$
   $\leq f(h_0 h_1 h_2 h_3 \dots) \leq \max\{f(h_0 h_2 h_4 \dots), f(h_1 h_3 h_5 \dots), \sup_i f(h_i^\omega)\}$

then both players have memoryless (optimal) winning strategies

- Many applications: Parity, Inf, LimInf, Mean-payoff, Total-payoff, ...
- If the payoff function is prefix-independent, i.e. $f(\rho) = f(h\rho)$, then conditions 1. and 2. are satisfied
- Simple proof by induction on the number of edges

Algorithmic results

Parity games in NP $\cap$ co-NP

Algorithmic results

Parity games in NP $\cap$ co-NP

- in NP:
  - Guess a memoryless winning strategy σ player 1
  - In the one-player game G_σ, check in polynomial time whether there exists a reachable cycle with odd maximum color

Algorithmic results

Parity games in NP ∩ co-NP

- in NP:
  - Guess a memoryless winning strategy σ player 1
  - In the one-player game G_σ, check in polynomial time whether there exists a reachable cycle with odd maximum color
- in co-NP: symmetrically for player 2

Algorithmic results

Parity games in NP $\cap$ co-NP

- in NP:
  - Guess a memoryless winning strategy σ player 1
  - In the one-player game G_σ, check in polynomial time whether there exists a reachable cycle with odd maximum color
- in co-NP: symmetrically for player 2

Mean-payoff games in NP $\cap$ co-NP

Algorithmic results

Parity games in NP ∩ co-NP

- in NP:
  - Guess a memoryless winning strategy σ player 1
  - In the one-player game G_σ, check in polynomial time whether there exists a reachable cycle with odd maximum color
- in co-NP: symmetrically for player 2

Mean-payoff games in NP ∩ co-NP

- Same approach
- One can compute in polynomial time the minimum (resp. maximum) mean weight cycle in a weighted graph [Kar78]

Window games

Objective to satisfy inside a finite window sliding along the play [CHH09], [CDRR15, BHRR16]

Definition

Given a window size λ and a threshold ν

- WMP objective: ensure for all positions i of the window, $\frac{1}{\ell}\sum_{k=0}^{\ell-1} w(\rho_{i+k},\rho_{i+k+1}) \geq \nu$ for some $\ell \leq \lambda$

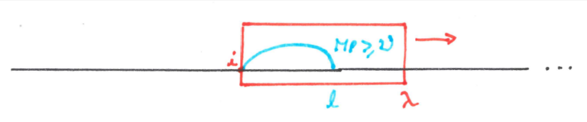

Window games

Objective to satisfy inside a finite window sliding along the play [CHH09], [CDRR15, BHRR16]

Definition

Given a window size λ and a threshold ν

- WMP objective: ensure for all positions i of the window, $\frac{1}{\ell}\sum_{k=0}^{\ell-1} w(\rho_{i+k}, \rho_{i+k+1}) \geq \nu$ for some $\ell \leq \lambda$

Given a window size λ and a coloring c

- WParity objective: ensure for all positions i of the window, $\max_{0\leq k\leq \ell} c(\rho_{i+k})$ even for some $\ell \leq \lambda$

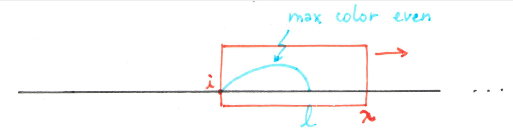

Window games

Motivations

- Strengthening of the Mean-payoff and Parity objectives
  Guarantee within a bounded time, not in the limit
- More computationally tractable
  (Open problem: Mean-payoff and Parity games are in P?)

Window games

Motivations

- Strengthening of the Mean-payoff and Parity objectives
  Guarantee within a bounded time, not in the limit
- More computationally tractable
  (Open problem: Mean-payoff and Parity games are in P?)

Remark: ω-regular objectives, thus determined games

Window games

Motivations

- Strengthening of the Mean-payoff and Parity objectives
  Guarantee within a bounded time, not in the limit
- More computationally tractable
  (Open problem: Mean-payoff and Parity games are in P?)

Remark: ω-regular objectives, thus determined games

Example

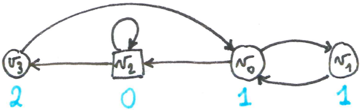

- Player 1 winning for Parity
- but loosing for WParity
  for all window sizes λ

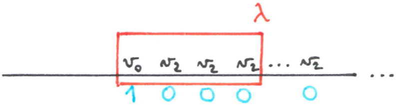

Window games

Results

	$\underline{\text{MP}}, \overline{\text{MP}}$	Parity	WMP	WParity
Complexity	NP $\cap$ co-NP		P-complete polynomial windows	P-complete
Player 1 strategy	memoryless		finite-memory	
Player 2 strategy	memoryless		finite-memory	

Model

Extension to k-dimensional weighted games with $k \geq 2$

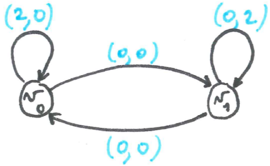

Definition

Two-player zero-sum game $G = (V, V_1, V_2, E, v_0, w)$ as before except:

- $w : E \to \mathbb{Z}^k$ weight function

Model

Extension to *k-dimensional* weighted games with $k \geq 2$

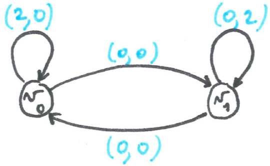

Definition

Two-player zero-sum game $G = (V, V_1, V_2, E, v_0, w)$ as before except:

- $w : E \to \mathbb{Z}^k$ weight function

Definition

Quantitative objective Ω:

- Intersection $\cap_{\ell=1}^{k} \Omega_\ell$ such that Ω_ℓ is an objective for dimension ℓ
- More generally, Boolean combination of such objectives Ω_ℓ, $1 \leq \ell \leq k$

Example

- One-player game
- $k = 2$, $\Omega = \Omega_1 \cap \Omega_2$
  with $\Omega_1 = \underline{\mathrm{MP}}(\rho) \geq 1$ for dimension 1
  and $\Omega_2 = \underline{\mathrm{MP}}(\rho) \geq 1$ for dimension 2

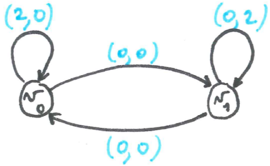

Example

- One-player game
- $k = 2$, $\Omega = \Omega_1 \cap \Omega_2$
  with $\Omega_1 = \underline{\mathrm{MP}}(\rho) \geq 1$ for dimension 1
  and $\Omega_2 = \underline{\mathrm{MP}}(\rho) \geq 1$ for dimension 2

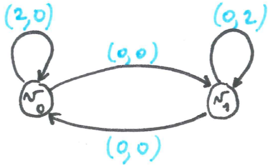

- Player 1 is losing with finite-memory strategies
  - Eventually periodic play $\rho = hg^\omega$
  - Mean-payoff of cycle g equal to

$$a \cdot (2,0) + b \cdot (0,0) + c \cdot (0,2) = (2 \cdot a, 2 \cdot c) \not\geq (1,1)$$

  with $a + b + c = 1$ and $b > 0$

Example

- One-player game
- $k = 2$, $\Omega = \Omega_1 \cap \Omega_2$
  with $\Omega_1 = \underline{\mathsf{MP}}(\rho) \geq 1$ for dimension 1
  and $\Omega_2 = \underline{\mathsf{MP}}(\rho) \geq 1$ for dimension 2

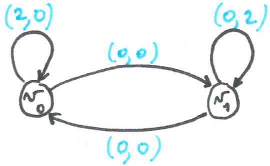

- Player 1 is losing with finite-memory strategies
  - Eventually periodic play $\rho = hg^\omega$
  - Mean-payoff of cycle g equal to

$$a \cdot (2,0) + b \cdot (0,0) + c \cdot (0,2) = (2 \cdot a, 2 \cdot c) \not\geq (1,1)$$

  with $a + b + c = 1$ and $b > 0$
- Player 1 is winning with infinite-memory strategies
  - Alternate visits to v_0 and v_1
  - At alternation n, loop n times on v_0 and then loop n times on v_1

$$a_n \cdot (2,0) + \epsilon_n \cdot (0,0) + a_n \cdot (0,2)$$

  with $\epsilon_n \to 0$ and $a_n \to \frac{1}{2}$
  - Mean-payoff equal to $(1,1)$

Results with the same objective on all dimensions

[CDHR10],[CRR14, CDRR15]

	$\underline{\mathsf{TP}}$, $\overline{\mathsf{TP}}$	$\overline{\mathsf{MP}}$	$\underline{\mathsf{MP}}$
Complexity	Undecidable	NP $\cap$ co-NP	coNP-complete
Player 1 strategy	-	infinite-memory	
Player 2 strategy	-	memoryless	

Results with the same objective on all dimensions

[CDHR10],[CRR14, CDRR15]

	$\underline{\mathrm{TP}}$, $\overline{\mathrm{TP}}$	$\overline{\mathrm{MP}}$	$\underline{\mathrm{MP}}$
Complexity	Undecidable	NP $\cap$ co-NP	coNP-complete
Player 1 strategy	-	infinite-memory	
Player 2 strategy	-	memoryless	

[CDRR15, BHRR16]

	WMP	WParity
Complexity	Exptime-complete	
Player 1 strategy	finite-memory	
Player 2 strategy	finite-memory	

Results with heterogeneous objectives

Only preliminary results

Theorem

- [Vel15]: Undecidability for Boolean combinations of $\underline{\mathrm{MP}}$ and $\overline{\mathrm{MP}}$ objectives
- [BHR15]: Decidability for Boolean combinations of WMP, Inf, Sup, LimInf, LimSup objectives, and finite-memory winning strategies for both players

Model

Summary

- 2-player zero-sum games, one player against the other
- Qualitative/quantitative uni/multidimensional objectives

Model

Summary

- 2-player zero-sum games, one player against the other
- Qualitative/quantitative uni/multidimensional objectives

Multiplayer non zero-sum games

- Several players with their own objective
- Non necessarily antagonistic objectives

Model

Summary

- 2-player zero-sum games, one player against the other
- Qualitative/quantitative uni/multidimensional objectives

Multiplayer non zero-sum games

- Several players with their own objective
- Non necessarily antagonistic objectives

Definition

n-player non zero-sum game $G = (V, (V_p)_{p \in \Pi}, E, v_0, \bar{w})$:

- Set Π of n players, $n \geq 1$
- $(V_p)_{p \in \Pi}$ partition of V with V_p controlled by player $p \in \Pi$
- optional: $\bar{w} = (w_p)_{p \in \Pi} : E \to \mathbb{Z}^n$ such that
  - w_p is the weight function of player p
  - leading to his payoff function f_p

Model

Definition

Objective Ω_p for each player $p \in \Pi$

- qualitative
- quantitative (depending on f_p)

Strategy profile $\bar{\sigma} = (\sigma_p)_{p \in \Pi}$ and outcome $\rho = \langle \bar{\sigma} \rangle_{v_0}$ from initial vertex v_0.

Model

Definition

Objective Ω_p for each player $p \in \Pi$

- qualitative
- quantitative (depending on f_p)

Strategy profile $\bar{\sigma} = (\sigma_p)_{p \in \Pi}$ and outcome $\rho = \langle \bar{\sigma} \rangle_{v_0}$ from initial vertex v_0.

Example

- 3-player game (player 1 ◯, player 2 □, player 3 ◊)
- Each player wants to visit one of his vertices infinitely often
- Outcome $(v_0 v_1 v_2 v_1)^\omega$: good solution for each player that needs cooperation

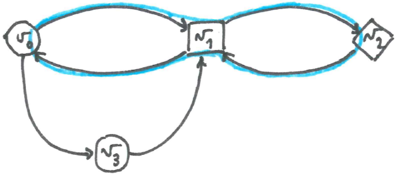

Nash equilibria (NE)

Classical notion such that each player is

- rational (he wants to maximize his payoff)
- selfish (he is only concerned with his own payoff)

Nash equilibria (NE)

Classical notion such that each player is

- rational (he wants to maximize his payoff)
- selfish (he is only concerned with his own payoff)

Informally, $\bar{\sigma}$ is an NE if no player has an incentive to deviate from his strategy, if the other players stick to their own strategies

Definition [Nas50]

The strategy profile $\bar{\sigma}$ is a Nash equilibrium from v_0 if, for each player $p \in \Pi$, for each strategy σ'_p of p,

$$f_p(\langle \sigma'_p, \bar{\sigma}_{-p} \rangle_{v_0}) \quad \leq \quad f_p(\langle \bar{\sigma} \rangle_{v_0}).$$

Notation: $\bar{\sigma}_{-p} = (\sigma_i)_{i \in \Pi \setminus \{p\}}$.

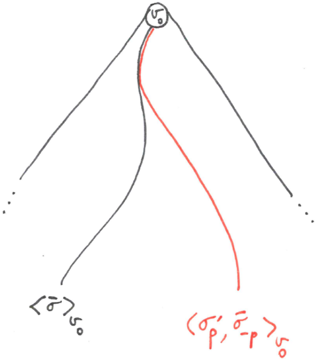

Nash equilibria

Example of NE: outcome $v_0 v_2 v_4^\omega$ with payoff $(3, 2)$

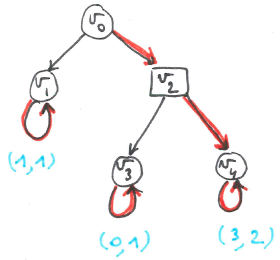

No incentive to deviate:

- If player 1 deviates to v_1, he will get 1 instead of 3
- If player 2 deviates to v_3, he will get 1 instead of 2

Nash equilibria

Example of NE: outcome $v_0 v_2 v_4^\omega$ with payoff $(3, 2)$

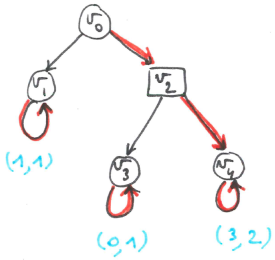

No incentive to deviate:

- If player 1 deviates to v_1, he will get 1 instead of 3
- If player 2 deviates to v_3, he will get 1 instead of 2

Algorithmic questions

- Does there exist an NE from initial vertex v_0?
- Can we construct it?
- With what kind of strategies? Memoryless, finite-memory?

Some results

Theorem

Qualitative objectives

- [GU08]: Existence of an NE in case of Borel objectives

Some results

Theorem

Qualitative objectives

- [GU08]: Existence of an NE in case of Borel objectives

Quantitative objectives

- [Kuh53]: Construction of an NE for games played on a finite tree
- [FL83, Har85]: Existence of an NE if the payoff function f_p of each player p is bounded and continuous

Some results

Theorem

Qualitative objectives

- [GU08]: Existence of an NE in case of Borel objectives

Quantitative objectives

- [Kuh53]: Construction of an NE for games played on a finite tree
- [FL83, Har85]: Existence of an NE if the payoff function f_p of each player p is bounded and continuous

Proof of [Kuh53]: Backward induction from the leaves to the root

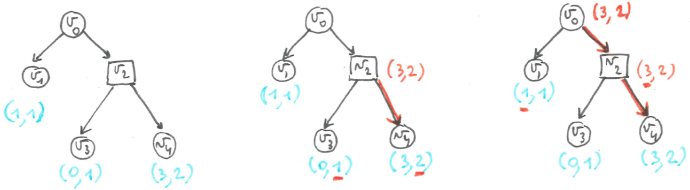

Some results

Theorem

Qualitative objectives

- [GU08]: Existence of an NE in case of Borel objectives

Quantitative objectives

- [Kuh53]: Construction of an NE for games played on a finite tree
- [FL83, Har85]: Existence of an NE if the payoff function f_p of each player p is bounded and continuous

Proof of [FL83, Har85]:

- On the unravelling of the game G truncated at depth d, construction of an NE $\bar{\sigma}^d$ by [Kuh53]
- A subsequence of $(\bar{\sigma}^d)_{d \in \mathbb{N}}$ converges to a strategy profile $\bar{\sigma}^*$ that is proved to be an NE

Some results

Theorem

Qualitative objectives

- [GU08]: Existence of an NE in case of Borel objectives

Quantitative objectives

- [Kuh53]: Construction of an NE for games played on a finite tree
- [FL83, Har85]: Existence of an NE if the payoff function f_p of each player p is bounded and continuous

Remark: The LimInf, LimSup and MP payoff functions are not continuous

For example, $\rho^n = v_0^n v_1^\omega \to \rho = v_0^\omega$ and

$\lim_{n\to\infty} \text{LimInf}(\rho^n) = 1 \neq 0 = \text{LimInf}(\rho)$

Some results

Theorem [BDS13]

Constuction of a finite-memory NE in the game G if for all p

- the payoff function f_p satisfies: $f_p(\rho) \leq f_p(\rho') \Rightarrow f_p(h\rho) \leq f_p(h\rho')$
- the two-player zero-sum game G^p
  - where player 1 is p and player 2 is the coalition of the other players
  - is determined with memoryless optimal winning strategies for both players

Some results

Theorem [BDS13]

Constuction of a finite-memory NE in the game G if for all p

- the payoff function f_p satisfies: $f_p(\rho) \leq f_p(\rho') \Rightarrow f_p(h\rho) \leq f_p(h\rho')$
- the two-player zero-sum game G^p
  - where player 1 is p and player 2 is the coalition of the other players
  - is determined with memoryless optimal winning strategies for both players

Proof:

- Let σ_p (resp. σ_{-p}) be an optimal winning strategy of player p (resp. the coalition) in G^p
- Construct the profile $\bar{\sigma}$:
  - play as σ_p for each player p
    (p plays selfishly and optimally for his own objective)
  - and as soon as some player p deviates, punish p by playing σ_{-p}
    (the coalition plays against p's objective)

Other kinds of equilibria

Subgame perfect equilibrium (SPE) [Sel65]

- takes into account the sequential nature of games played on graphs
- i.e., is an NE from the initial vertex v_0, but also after every history h of the game

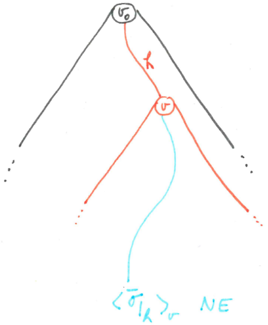

Other kinds of equilibria

Subgame perfect equilibrium (SPE) [Sel65]

- takes into account the sequential nature of games played on graphs
- i.e., is an NE from the initial vertex v_0, but also after every history h of the game

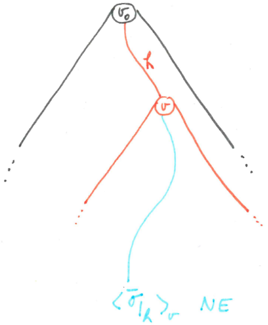

Theorem

- Previous results [Kuh53] and [FL83, Har85] provide NE and more generally SPE

Other kinds of equilibria

Subgame perfect equilibrium (SPE) [Sel65]

- takes into account the sequential nature of games played on graphs
- i.e., is an NE from the initial vertex v_0, but also after every history h of the game

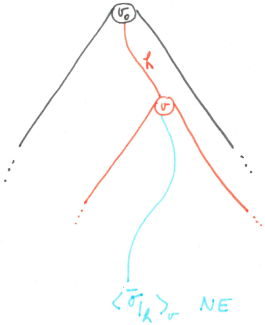

Theorem

- Previous results [Kuh53] and [FL83, Har85] provide NE and more generally SPE
- [BBMR15]: Construction of a finite-memory SPE for quantitative reachability games

Other kinds of equilibria

Secure equilibrium (SE) [CHJ06]

- each player wants to maximize his payoff, as a first objective
- and then minimize the payoff of the other players, as a second objective

Other kinds of equilibria

Secure equilibrium (SE) [CHJ06]

- each player wants to maximize his payoff, as a first objective
- and then minimize the payoff of the other players, as a second objective

Theorem

- [CHJ06]: Existence of an SE for 2-player games with Borel qualitative objectives
- [DFK+14]: Existence of an SE
  - if the payoff function f_p of each player p is bounded and continuous
  - or if each f_p is Borel measurable and have finite range
- [BMR14]: Previous result [BDS13] extended to SE for 2-player games

Summary

- 2-player zero-sum games, one player against the other
- Qualitative/quantitative uni/multidimensional objective
- Extension to multiplayer non zero-sum games
- Different notions of equilibria (NE, SPE, SE)

Summary

- 2-player zero-sum games, one player against the other
- Qualitative/quantitative uni/multidimensional objective
- Extension to multiplayer non zero-sum games
- Different notions of equilibria (NE, SPE, SE)

Need for further research on

- 2-player zero-sum games with heterogeneous objectives for the synthesis of correct reactive systems
- Adequate notions of equilibrium on multiplayer non zero-sum games for the synthesis of correct complex systems

Summary

- 2-player zero-sum games, one player against the other
- Qualitative/quantitative uni/multidimensional objective
- Extension to multiplayer non zero-sum games
- Different notions of equilibria (NE, SPE, SE)

Need for further research on

- 2-player zero-sum games with heterogeneous objectives for the synthesis of correct reactive systems
- Adequate notions of equilibrium on multiplayer non zero-sum games for the synthesis of correct complex systems

Other extensions

- Concurrent games
- Stochastic games
- Imperfect information

Summary

- 2-player zero-sum games, one player against the other
- Qualitative/quantitative uni/multidimensional objective
- Extension to multiplayer non zero-sum games
- Different notions of equilibria (NE, SPE, SE)

Need for further research on

- 2-player zero-sum games with heterogeneous objectives for the synthesis of correct reactive systems
- Adequate notions of equilibrium on multiplayer non zero-sum games for the synthesis of correct complex systems

Other extensions

- Concurrent games
- Stochastic games
- Imperfect information

Thank you!

- Thomas Brihaye, Véronique Bruyère, Noémie Meunier, and Jean-François Raskin, *Weak subgame perfect equilibria and their application to quantitative reachability*, CSL'15, LIPIcs, vol. 41, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015, pp. 504–518.
- Thomas Brihaye, Julie De Pril, and Sven Schewe, *Multiplayer cost games with simple nash equilibria*, LFCS'13, LNCS, vol. 7734, Springer, 2013, pp. 59–73.
- Catriel Beeri, *On the membership problem for functional and multivalued dependencies in relational databases*, ACM Trans. Database Syst. **5** (1980), no. 3, 241–259.
- Véronique Bruyère, Quentin Hautem, and Jean-François Raskin, *On the complexity of heterogeneous multidimensional quantitative games*, CoRR **abs/1511.08334** (2015).

- Véronique Bruyère, Quentin Hautem, Mickael Randour, and Jean-François Raskin, *Games with window parity objectives*, Tech. report, University of Mons, 2016.

- Véronique Bruyère, Noémie Meunier, and Jean-François Raskin, *Secure equilibria in weighted games*, CSL-LICS'14, ACM, 2014, pp. 26:1–26:26.

- Henrik Björklund, Sven Sandberg, and Sergei G. Vorobyov, *Memoryless determinacy of parity and mean payoff games: a simple proof*, Theor. Comput. Sci. **310** (2004), no. 1-3, 365–378.

- Krishnendu Chatterjee, Laurent Doyen, Thomas A. Henzinger, and Jean-François Raskin, *Generalized mean-payoff and energy games*, FSTTCS'10, LIPIcs, vol. 8, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2010, pp. 505–516.

- Krishnendu Chatterjee, Laurent Doyen, Mickael Randour, and Jean-François Raskin, *Looking at mean-payoff and total-payoff through windows*, Inf. Comput. **242** (2015), 25–52.

Krishnendu Chatterjee, Thomas A. Henzinger, and Florian Horn, *Finitary winning in omega-regular games*, ACM Trans. Comput. Log. **11** (2009), no. 1.

Krishnendu Chatterjee, Thomas A. Henzinger, and Marcin Jurdzinski, *Games with secure equilibria*, Theor. Comput. Sci. **365** (2006), no. 1-2, 67–82.

Krishnendu Chatterjee, Mickael Randour, and Jean-François Raskin, *Strategy synthesis for multi-dimensional quantitative objectives*, Acta Inf. **51** (2014), no. 3-4, 129–163.

Julie De Pril, János Flesch, Jeroen Kuipers, Gijs Schoenmakers, and Koos Vrieze, *Existence of secure equilibrium in multi-player games with perfect information*, MFCS'14, LNCS, vol. 8635, Springer, 2014, pp. 213–225.

E. Allen Emerson and Charanjit S. Jutla, *Tree automata, mu-calculus and determinacy (extended abstract)*, FOCS'91, IEEE Computer Society, 1991, pp. 368–377.

A. Ehrenfeucht and J. Mycielski, *Positional strategies for mean payoff games*, International Journal of Game Theory **8** (1979), no. 2, 109?113.

D. Fudenberg and D. Levine, *Subgame-perfect equilibria of finite- and infinite-horizon games*, Journal of Economic Theory **31 (2)** (1983), 251–268.

Erich Grädel, Wolfgang Thomas, and Thomas Wilke (eds.), *Automata, logics, and infinite games: A guide to current research [outcome of a dagstuhl seminar, february 2001]*, LNCS, vol. 2500, Springer, 2002.

Erich Grädel and Michael Ummels, *Solution Concepts and Algorithms for Infinite Multiplayer Games*, New Perspectives on Games and Interaction (Krzysztof Apt and Robert van Rooij, eds.), Texts in Logic and Games, vol. 4, Amsterdam University Press, 2008, pp. 151–178.

Hugo Gimbert and Wieslaw Zielonka, *When can you play positionally?*, MFCS'04, LNCS, vol. 3153, Springer, 2004, pp. 686–697.

- C. J. Harris, *Existence and characterization of perfect equilibrium in games of perfect information*, Econometrica **53 (3)** (1985), 613–628.
- Neil Immerman, *Number of quantifiers is better than number of tape cells*, J. Comput. Syst. Sci. **22** (1981), no. 3, 384–406.
- Marcin Jurdzinski, *Deciding the winner in parity games is in UP \cap co-up*, Inf. Process. Lett. **68** (1998), no. 3, 119–124.
- Richard M Karp, *A characterization of the minimum cycle mean in a digraph*, Discrete Mathematics **23** (1978), 309–311.
- H.W. Kuhn, *Extensive games and the problem of information*, Classics in Game Theory (1953), 46–68.
- Donald A. Martin, *Borel determinacy*, Annals of Mathematics **102** (1975), no. 2, 363–371.
- J. F. Nash, *Equilibrium points in n-person games*, PNAS, vol. 36, National Academy of Sciences, 1950, pp. 48–49.

Reinhard Selten, *Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit*, Zeitschrift für die gesamte Staatswissenschaft **121** (1965), 301–324 and 667–689.

Yaron Velner, *Robust multidimensional mean-payoff games are undecidable*, FoSSaCS'15, LNCS, vol. 9034, Springer, 2015, pp. 312–327.

Wieslaw Zielonka, *Infinite games on finitely coloured graphs with applications to automata on infinite trees*, Theor. Comput. Sci. **200** (1998), no. 1-2, 135–183.

Uri Zwick and Mike Paterson, *The complexity of mean payoff games on graphs*, Theor. Comput. Sci. **158** (1996), no. 1&2, 343–359.