Nash equilibrium in quantitative games

Thomas BRIHAYE Véronique BRUYÈRE Julie DE PRIL

University of Mons (Belgium)

UMONS

2^{nd} *Meeting GASICS 2009*
October 22, 2009

Outline

Outline

Features of our games

- games on finite graph with no deadlock
- n players
- turn based
- reachability objectives for all the players
- payoff: least number of edges to reach one's goal set, or $+\infty$
- aim: minimize one's payoff

$\rightsquigarrow$ multiplayer, non-zero-sum, quantitative games

Features of our games

- games on finite graph with no deadlock
- n players
- turn based
- reachability objectives for all the players
- payoff: least number of edges to reach one's goal set, or $+\infty$
- aim: minimize one's payoff

$\rightsquigarrow$ multiplayer, non-zero-sum, quantitative games

Features of our games

- games on finite graph with no deadlock
- n players
- turn based
- reachability objectives for all the players
- payoff: least number of edges to reach one's goal set, or $+\infty$
- aim: minimize one's payoff

$\rightsquigarrow$ multiplayer, non-zero-sum, quantitative games

Features of our games

- games on finite graph with no deadlock
- n players
- turn based
- reachability objectives for all the players
- payoff: least number of edges to reach one's goal set, or $+\infty$
- aim: minimize one's payoff

$\rightsquigarrow$ multiplayer, non-zero-sum, quantitative games

Features of our games

- games on finite graph with no deadlock
- *n* players
- turn based
- reachability objectives for all the players
- payoff: least number of edges to reach one's goal set, or $+\infty$
- aim: minimize one's payoff

$\rightsquigarrow$ multiplayer, non-zero-sum, quantitative games

Features of our games

- games on finite graph with no deadlock
- n players
- turn based
- reachability objectives for all the players
- payoff: least number of edges to reach one's goal set, or $+\infty$
- aim: minimize one's payoff

$\rightsquigarrow$ multiplayer, non-zero-sum, quantitative games

Features of our games

- games on finite graph with no deadlock
- n players
- turn based
- reachability objectives for all the players
- payoff: least number of edges to reach one's goal set, or $+\infty$
- aim: minimize one's payoff

$\rightsquigarrow$ multiplayer, non-zero-sum, quantitative games

Example

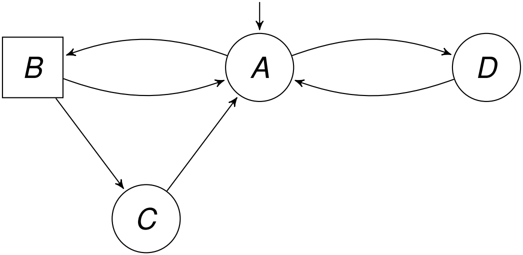

Player 1 ○ $Reach_1 = \{C\}$

Player 2 □ $Reach_2 = \{D\}$

Example

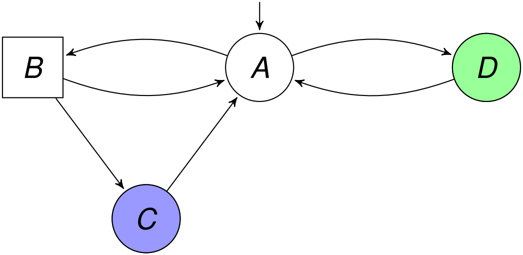

Player 1 ◯ $\textcolor{blue}{Reach_1} = \{C\}$
Player 2 ☐ $\textcolor{green}{Reach_2} = \{D\}$

Example

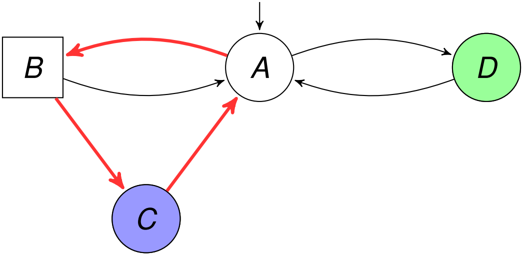

Player 1 ○ $Reach_1 = \{C\}$
Player 2 □ $Reach_2 = \{D\}$

Play: $(ABC)^\omega$
Payoffs: $(2, +\infty)$

Example

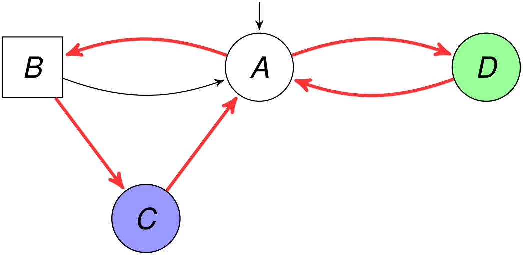

Player 1 ◯ $Reach_1 = \{C\}$
Player 2 □ $Reach_2 = \{D\}$

Play: $AD(ABC)^\omega$
Payoffs: $(4, 1)$

Nash equilibrium in such games

Idea: A strategy profile where no player has an incentive to deviate from the strategy chosen, given the choices of the other players.

Player 1 ○ $Reach_1 = \{C\}$
Player 2 □ $Reach_2 = \{D\}$

Nash equilibrium in such games

Idea: A strategy profile where no player has an incentive to deviate from the strategy chosen, given the choices of the other players.

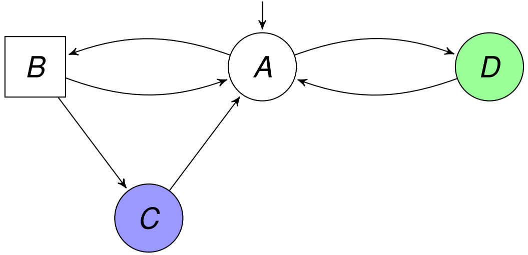

Player 1 ○ $Reach_1 = \{C\}$
Player 2 □ $Reach_2 = \{D\}$

Nash equilibrium in such games

Idea: A strategy profile where no player has an incentive to deviate from the strategy chosen, given the choices of the other players.

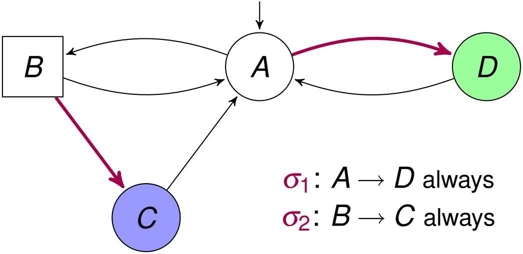

σ_1: $A \rightarrow D$ always
σ_2: $B \rightarrow C$ always

Player 1 ○ $Reach_1 = \{C\}$
Player 2 □ $Reach_2 = \{D\}$

Nash equilibrium in such games

Idea: A strategy profile where no player has an incentive to deviate from the strategy chosen, given the choices of the other players.

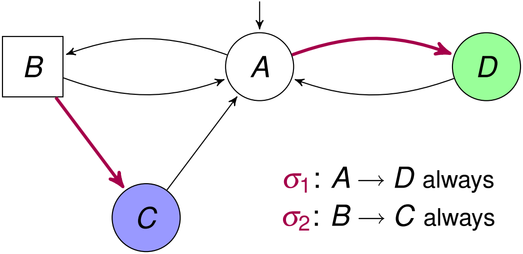

σ_1: $A \to D$ always
σ_2: $B \to C$ always

Player 1 ◯ $Reach_1 = \{C\}$
Player 2 ☐ $Reach_2 = \{D\}$

Play: $(AD)^\omega$
Payoffs: $(+\infty, 1)$
(σ_1, σ_2): NOT a Nash equilibrium

Nash equilibrium in such games

Idea: A strategy profile where no player has an incentive to deviate from the strategy chosen, given the choices of the other players.

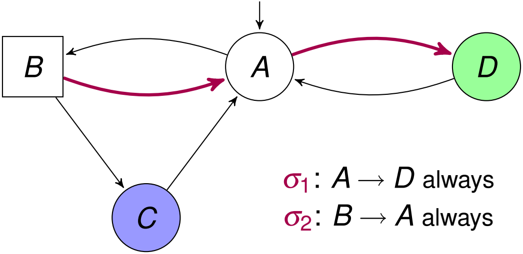

σ_1: $A \rightarrow D$ always
σ_2: $B \rightarrow A$ always

Player 1 ○ $Reach_1 = \{C\}$
Player 2 □ $Reach_2 = \{D\}$

Nash equilibrium in such games

Idea: A strategy profile where no player has an incentive to deviate from the strategy chosen, given the choices of the other players.

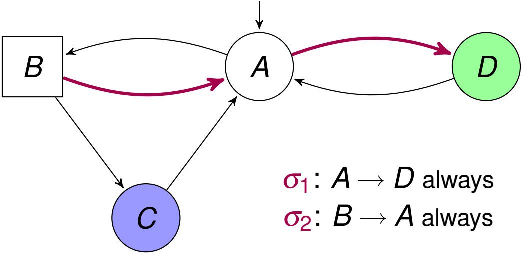

σ_1: $A \rightarrow D$ always
σ_2: $B \rightarrow A$ always

Player 1 ◯ $Reach_1 = \{C\}$
Player 2 ☐ $Reach_2 = \{D\}$

Play: $(AD)^\omega$
Payoffs: $(+\infty, 1)$
(σ_1, σ_2): Nash equilibrium

Nash equilibrium in such games

Idea: A strategy profile where no player has an incentive to deviate from the strategy chosen, given the choices of the other players.

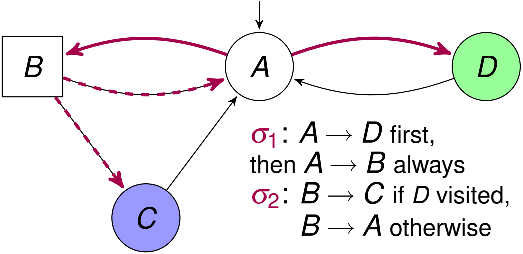

σ_1: $A \longrightarrow D$ first,
then $A \longrightarrow B$ always
σ_2: $B \longrightarrow C$ if D visited,
$B \longrightarrow A$ otherwise

Player 1 ○ $Reach_1 = \{C\}$
Player 2 □ $Reach_2 = \{D\}$

Nash equilibrium in such games

Idea: A strategy profile where no player has an incentive to deviate from the strategy chosen, given the choices of the other players.

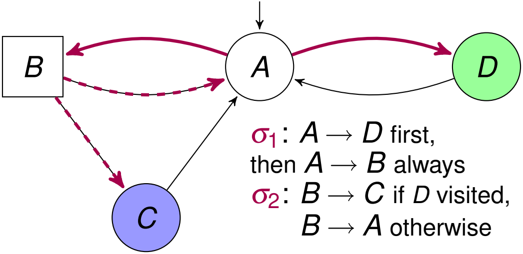

σ_1: $A \to D$ first,
then $A \to B$ always
σ_2: $B \to C$ if D visited,
$B \to A$ otherwise

Player 1 ○ $Reach_1 = \{C\}$
Player 2 □ $Reach_2 = \{D\}$

Play: $AD(ABC)^\omega$
Payoffs: $(4, 1)$
(σ_1, σ_2): Nash equilibrium

Qualitative NE $\neq$ Quantitative NE

Player 1 ○ $Reach_1 = \{C\}$
Player 2 □ $Reach_2 = \{B\}$

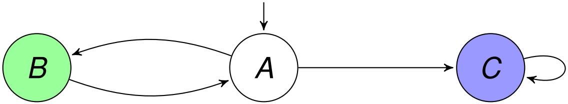

$\sigma_1 : A \rightarrow B$ first, then $A \rightarrow C$

Play: $ABAC^{\omega}$

Qualitative payoff profile: (Win, Win) $\rightsquigarrow$ qualitative Nash equilibrium
Quantitative payoff profile: $(3, 1)$ but not a quantitative Nash equilibrium

Qualitative NE $\neq$ Quantitative NE

Player 1 ◯ $\mathit{Reach}_1 = \{C\}$
Player 2 □ $\mathit{Reach}_2 = \{B\}$

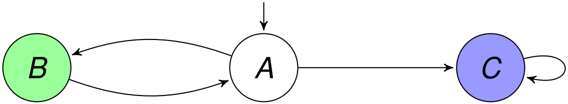

$\sigma_1 : A \to B$ first, then $A \to C$

Play: $ABAC^{\omega}$

Qualitative payoff profile: $(\mathit{Win}, \mathit{Win})$ $\rightsquigarrow$ qualitative Nash equilibrium
Quantitative payoff profile: $(3, 1)$ but not a quantitative Nash equilibrium

Qualitative NE $\neq$ Quantitative NE

Player 1 ◯ $\mathit{Reach}_1 = \{C\}$
Player 2 □ $\mathit{Reach}_2 = \{B\}$

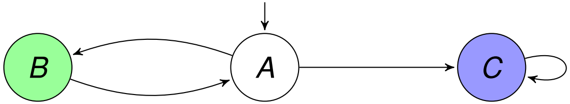

σ_1: $A \to B$ first, then $A \to C$

Play: $ABAC^\omega$

Qualitative payoff profile: $(\mathit{Win}, \mathit{Win})$ $\rightsquigarrow$ qualitative Nash equilibrium

Quantitative payoff profile: $(3, 1)$ but not a quantitative Nash equilibrium

Qualitative NE $\neq$ Quantitative NE

Player 1 ◯ $Reach_1 = \{C\}$
Player 2 □ $Reach_2 = \{B\}$

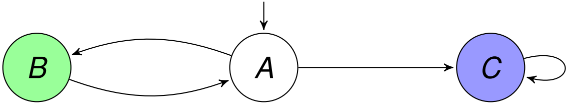

σ_1: $A \to B$ first, then $A \to C$

Play: $ABAC^\omega$

Qualitative payoff profile: (Win, Win) $\rightsquigarrow$ qualitative Nash equilibrium
Quantitative payoff profile: $(3, 1)$ but not a quantitative Nash equilibrium

Overview

	Zero-sum	Non-zero-sum
Qualitative	[Tho08],...	[CHJ06], [GU09],...
Quantitative	[HTW08], [BCHJ09],...	[BG09],...

[BG09] Why chess and back gammon can be solved in pure positional uniformly optimal strategies

[Tho08] Church's problem and a tour through automata theory

[HTW08] Optimal strategy synthesis in request response games

[BCHJ09] Better quality in synthesis through quantitative objectives

[CHJ06] Games with secure equilibria

[GU09] Solution concepts and algorithms for infinite multiplayer games

Motivation: e.g., solving games where everyone wants to minimize the time (or the energy consumption) to reach a given state set.

Overview

	Zero-sum	Non-zero-sum
Qualitative	[Tho08],...	[CHJ06], [GU09],...
Quantitative	[HTW08], [BCHJ09],...	[BG09],...

[BG09]	Why chess and back gammon can be solved in pure positional uniformly optimal strategies
[Tho08]	Church's problem and a tour through automata theory
[HTW08]	Optimal strategy synthesis in request response games
[BCHJ09]	Better quality in synthesis through quantitative objectives
[CHJ06]	Games with secure equilibria
[GU09]	Solution concepts and algorithms for infinite multiplayer games

Motivation: e.g., solving games where everyone wants to minimize the time (or the energy consumption) to reach a given state set.

Overview

	Zero-sum	Non-zero-sum
Qualitative	[Tho08],...	[CHJ06], [GU09],...
Quantitative	[HTW08], [BCHJ09],...	[BG09],...

[BG09]	Why chess and back gammon can be solved in pure positional uniformly optimal strategies
[Tho08]	Church's problem and a tour through automata theory
[HTW08]	Optimal strategy synthesis in request response games
[BCHJ09]	Better quality in synthesis through quantitative objectives
[CHJ06]	Games with secure equilibria
[GU09]	Solution concepts and algorithms for infinite multiplayer games

Motivation: e.g., solving games where everyone wants to minimize the time (or the energy consumption) to reach a given state set.

Outline

Theorem

There exists a Nash equilibrium in every multiplayer quantitative game with reachability objectives for all the players.

Player 1 ○ $Reach_1 = \{C\}$
Player 2 □ $Reach_2 = \{B\}$

No Nash equilibrium where both players reach their goal set...

But σ_1 leads to a Nash equilibrium, with payoff profile $(1, +\infty)$!

Existence of a Nash equilibrium

Theorem

There exists a Nash equilibrium in every multiplayer quantitative game with reachability objectives for all the players.

Player 1 ○ $Reach_1 = \{C\}$
Player 2 □ $Reach_2 = \{B\}$

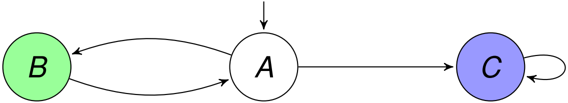

No Nash equilibrium where both players reach their goal set...

But σ_1 leads to a Nash equilibrium, with payoff profile $(1, +\infty)$!

Existence of a Nash equilibrium

Theorem

There exists a Nash equilibrium in every multiplayer quantitative game with reachability objectives for all the players.

Player 1 ○ $Reach_1 = \{C\}$
Player 2 □ $Reach_2 = \{B\}$

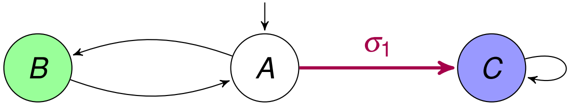

No Nash equilibrium where both players reach their goal set...

But σ_1 leads to a Nash equilibrium, with payoff profile $(1, +\infty)$!

Idea of the proof (1/4)

Consider the arena of an initialized game like an infinite tree.

Assume that the tree is finite. Then Kuhn's theorem provides a quantitative Nash equilibrium.

Backward induction

Idea of the proof (1/4)

Consider the arena of an initialized game like an infinite tree.

Assume that the tree is finite. Then Kuhn's theorem provides a quantitative Nash equilibrium.

Backward induction

Idea of the proof (1/4)

Consider the arena of an initialized game like an infinite tree.

Assume that the tree is finite. Then Kuhn's theorem provides a quantitative Nash equilibrium.

Player 1 ◯ $Reach_1 = \{C\}$
Player 2 □ $Reach_2 = \{D\}$

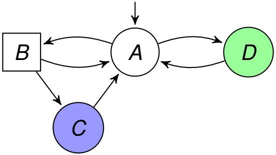

Backward induction

Idea of the proof (1/4)

Consider the arena of an initialized game like an infinite tree.

Assume that the tree is finite. Then Kuhn's theorem provides a quantitative Nash equilibrium.

Player 1 ○ $Reach_1 = \{C\}$
Player 2 □ $Reach_2 = \{D\}$

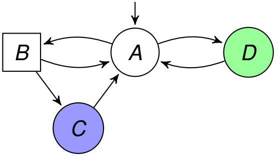

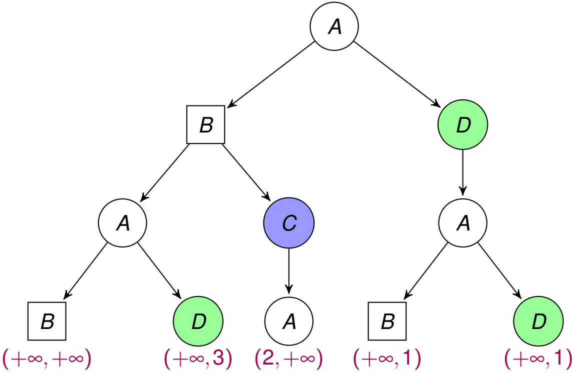

Idea of the proof (1/4)

Consider the arena of an initialized game like an infinite tree.

Assume that the tree is finite. Then Kuhn's theorem provides a quantitative Nash equilibrium.

Player 1 ○ $Reach_1 = \{C\}$
Player 2 □ $Reach_2 = \{D\}$

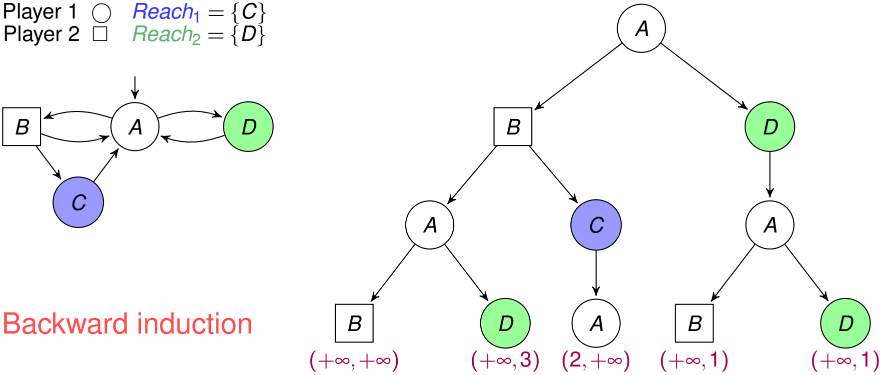

Backward induction

Idea of the proof (1/4)

Consider the arena of an initialized game like an infinite tree.

Assume that the tree is finite. Then Kuhn's theorem provides a quantitative Nash equilibrium.

Player 1 ○ $Reach_1 = \{C\}$
Player 2 □ $Reach_2 = \{D\}$

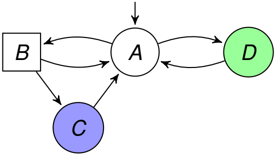

Backward induction

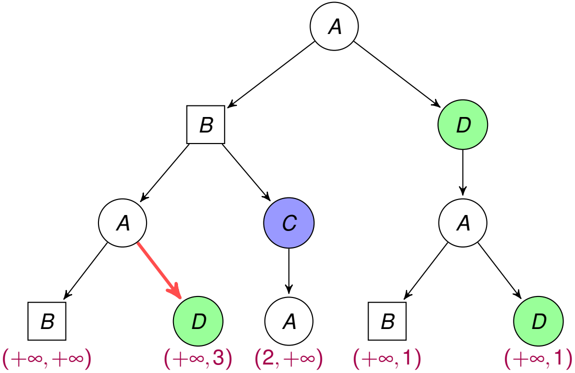

Idea of the proof (1/4)

Consider the arena of an initialized game like an infinite tree.

Assume that the tree is finite. Then Kuhn's theorem provides a quantitative Nash equilibrium.

Player 1 ○ $\mathit{Reach}_1 = \{C\}$
Player 2 □ $\mathit{Reach}_2 = \{D\}$

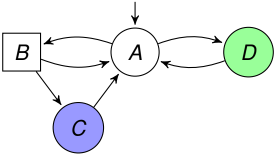

Backward induction

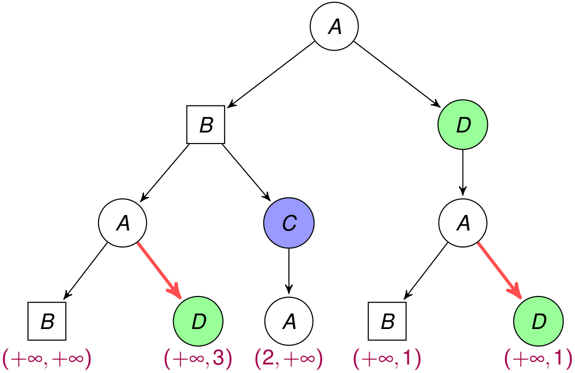

Idea of the proof (1/4)

Consider the arena of an initialized game like an infinite tree.

Assume that the tree is finite. Then Kuhn's theorem provides a quantitative Nash equilibrium.

Player 1 ○ $\mathit{Reach}_1 = \{C\}$
Player 2 □ $\mathit{Reach}_2 = \{D\}$

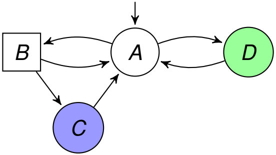

Backward induction

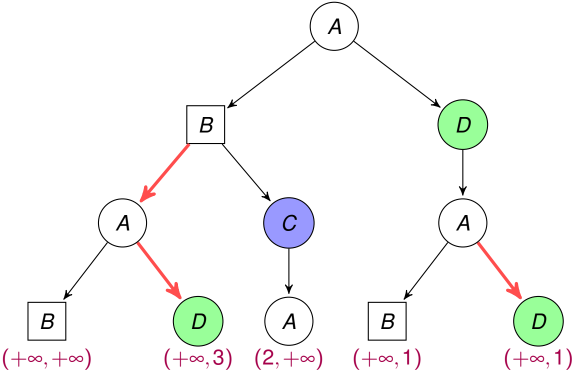

Idea of the proof (1/4)

Consider the arena of an initialized game like an infinite tree.

Assume that the tree is finite. Then Kuhn's theorem provides a quantitative Nash equilibrium.

Player 1 ○ $\mathit{Reach}_1 = \{C\}$
Player 2 □ $\mathit{Reach}_2 = \{D\}$

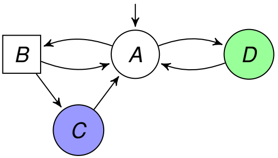

Backward induction

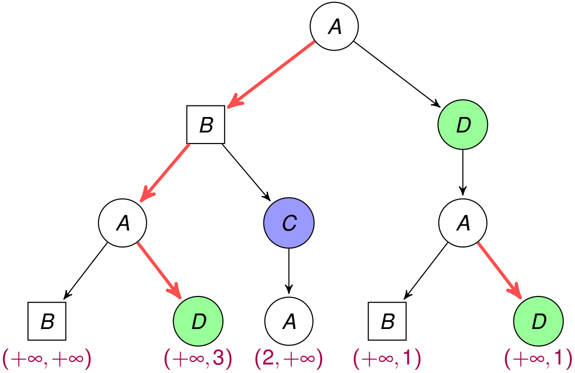

Thanks to Kuhn's theorem, if we prove the following result, we are done.

Proposition

Let T be the unravelling of the game and T_f be this tree limited to the depth $(n+1)*2*|V|$.
If there exists a Nash equilibrium in T_f, then there exists a Nash equilibrium in T.

Idea of the proof (3/4)

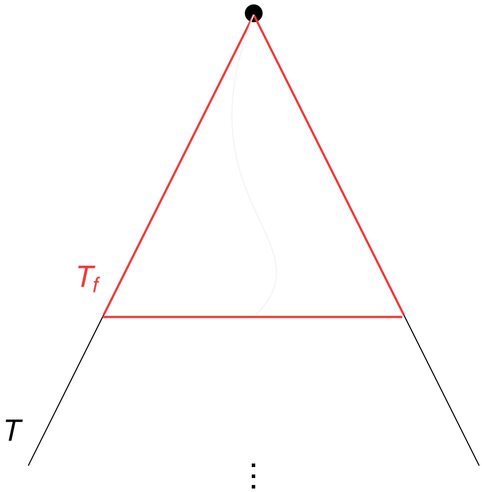

Idea of the proof (3/4)

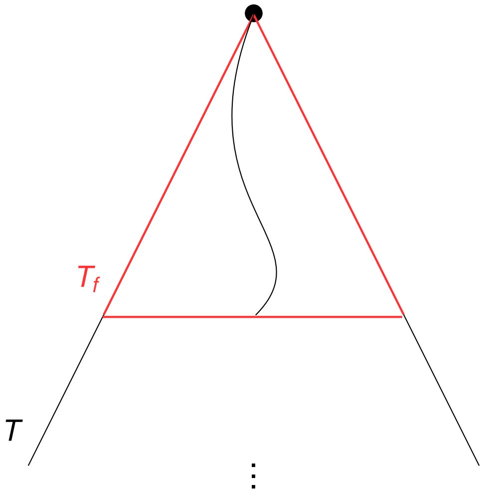

Idea of the proof (3/4)

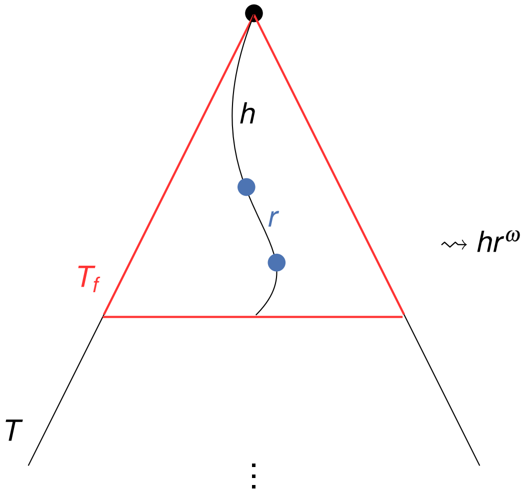

Idea of the proof (4/4)

Nash equilibrium in T:
$\longrightarrow$ everyone plays according to hr^ω
$\longrightarrow$ as soon as one player deviates, the other ones punish him by playing a strategy that prevents him from reaching his goal set (more quickly).

Remark: payoff profile not preserved...
But no hope to do better (example)!

Idea of the proof (4/4)

Nash equilibrium in T:
$\rightarrow$ everyone plays according to hr^{ω}
$\rightarrow$ as soon as one player deviates, the other ones punish him by playing a strategy that prevents him from reaching his goal set (more quickly).

Remark: payoff profile not preserved...
But no hope to do better (example)!

Nash equilibrium with bounded memory

Theorem

There exists a NE of payoff profile $(c_i)_i$ iff there exists a NE of payoff profile $(c_i)_i$ where the strategies have bounded memory.

Idea of the proof:

- transform the play of the NE into a play of the form hr^ω *(keeping the same payoff profile)* ;
- the Nash equilibrium with bounded memory will be:
  $\rightarrow$ everyone plays according to hr^ω
  $\rightarrow$ as soon as one player deviates, the other ones punish him by playing a bounded memory strategy that prevents him from reaching his goal set (more quickly).

Nash equilibrium with bounded memory

Theorem

There exists a NE of payoff profile $(c_i)_i$ iff there exists a NE of payoff profile $(c_i)_i$ where the strategies have bounded memory.

Idea of the proof:

- transform the play of the NE into a play of the form hr^ω *(keeping the same payoff profile)* ;
- the Nash equilibrium with bounded memory will be:
  $\rightarrow$ everyone plays according to hr^ω
  $\rightarrow$ as soon as one player deviates, the other ones punish him by playing a bounded memory strategy that prevents him from reaching his goal set (more quickly).

Secure equilibrium in such games

Idea: A strategy profile

- that is a Nash equilibrium; and
- such that if a player gets the same payoff when deviating, he can not increase the other players' payoffs and increase strictly one player's payoff.

σ_1: $A \to D$ always
σ_2: $B \to A$ always

Player 1 ○ $Reach_1 = \{C\}$
Player 2 □ $Reach_2 = \{D\}$

Play: $(AD)^\omega$
Payoffs: $(+\infty, 1)$

(σ_1, σ_2): Nash equilibrium
(σ_1, σ_2): NOT a secure equilibrium

Secure equilibrium in such games

Idea: A strategy profile

- that is a Nash equilibrium; and
- such that if a player gets the same payoff when deviating, he can not increase the other players' payoffs and increase strictly one player's payoff.

Player 1 ○ $Reach_1 = \{C\}$
Player 2 □ $Reach_2 = \{D\}$

Play: $(AD)^\omega$
Payoffs: $(+\infty, 1)$

(σ_1, σ_2): Nash equilibrium
(σ_1, σ_2): NOT a secure equilibrium

Secure equilibrium in such games

Idea: A strategy profile

- that is a Nash equilibrium; and
- such that if a player gets the same payoff when deviating, he can not increase the other players' payoffs and increase strictly one player's payoff.

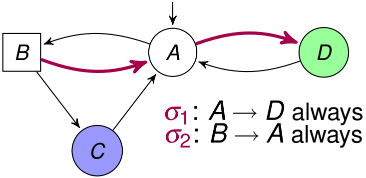

Player 1 ○ $Reach_1 = \{C\}$
Player 2 □ $Reach_2 = \{D\}$

Play: $(AD)^\omega$
Payoffs: $(+\infty, 1)$

(σ_1, σ_2): Nash equilibrium
(σ_1, σ_2): NOT a secure equilibrium

Secure equilibrium in such games

Idea: A strategy profile

- that is a Nash equilibrium; and
- such that if a player gets the same payoff when deviating, he can not increase the other players' payoffs and increase strictly one player's payoff.

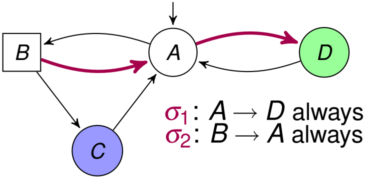

Player 1 ◯ $Reach_1 = \{C\}$
Player 2 □ $Reach_2 = \{D\}$

(σ_1, σ_2): Nash equilibrium

(σ_1, σ_2): NOT a secure equilibrium

Play: $(AD)^\omega$
Payoffs: $(+\infty, 1)$

Secure equilibrium in such games

Idea: A strategy profile

- that is a Nash equilibrium; and
- such that if a player gets the same payoff when deviating, he can not increase the other players' payoffs and increase strictly one player's payoff.

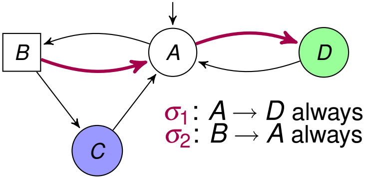

Player 1 ○ $Reach_1 = \{C\}$
Player 2 □ $Reach_2 = \{D\}$

Play: $(AD)^\omega$
Payoffs: $(+\infty, 1)$

(σ_1, σ_2): Nash equilibrium
(σ_1, σ_2): NOT a secure equilibrium

Existence of a secure equilibrium

Theorem

There exists a secure equilibrium in every multiplayer quantitative game with reachability objectives for all the players.

Proof: Same ideas as for the existence of a Nash equilibrium.

Existence of a secure equilibrium

Theorem

There exists a secure equilibrium in every multiplayer quantitative game with reachability objectives for all the players.

Proof: Same ideas as for the existence of a Nash equilibrium.

Outline

Future work

- Positive costs on edges: OK
  (payoff = sum of costs to reach one's goal set)
- Complexity
- Subgame perfect equilibrium
- Büchi objectives
- Request-response games
- Memory size
- Mixed strategies
- ...

Future work

- Positive costs on edges: OK
  (payoff = sum of costs to reach one's goal set)
- Complexity
- Subgame perfect equilibrium
- Büchi objectives
- Request-response games
- Memory size
- Mixed strategies
- ...

Future work

- Positive costs on edges: OK
  (payoff = sum of costs to reach one's goal set)
- Complexity
- Subgame perfect equilibrium
- Büchi objectives
- Request-response games
- Memory size
- Mixed strategies
- ...

Future work

- Positive costs on edges: OK
  (payoff = sum of costs to reach one's goal set)
- Complexity
- Subgame perfect equilibrium
- Büchi objectives
- Request-response games
- Memory size
- Mixed strategies
- ...

Future work

- Positive costs on edges: OK
  (payoff = sum of costs to reach one's goal set)
- Complexity
- Subgame perfect equilibrium
- Büchi objectives
- Request-response games
- Memory size
- Mixed strategies
- ...

Future work

- Positive costs on edges: OK
  (payoff = sum of costs to reach one's goal set)
- Complexity
- Subgame perfect equilibrium
- Büchi objectives
- Request-response games
- Memory size
- Mixed strategies
- ...

Future work

- Positive costs on edges: OK
  (payoff = sum of costs to reach one's goal set)
- Complexity
- Subgame perfect equilibrium
- Büchi objectives
- Request-response games
- Memory size
- Mixed strategies
- ...

Future work

- Positive costs on edges: OK
  (payoff = sum of costs to reach one's goal set)
- Complexity
- Subgame perfect equilibrium
- Büchi objectives
- Request-response games
- Memory size
- Mixed strategies
- ...

Thank you for your attention!