Cassting

Mixed Nash Equilibria in Concurrent Games

Daniel STAN

Supervised by Nicolas Markey and Patricia Bouyer

Second Cassting meeting

SEVENTH FRAMEWORK PROGRAMME

Definition: concurrent game

Definition

$$\mathcal{G} = \left\langle \mathrm{States}, \mathrm{Agt}, \mathrm{Act}, \mathrm{Mov}, \mathrm{Tab}, (\mathrm{Allow}_A)_{A \in \mathrm{Agt}}, (\phi_A)_{A \in \mathrm{Agt}} \right\rangle$$

with

- $|\mathrm{States}| < +\infty$, $|\mathrm{Agt}| < +\infty$ *and* $|\mathrm{Act}| < +\infty$
- $\mathrm{Tab} : \mathrm{States} \times \mathrm{Act}^{\mathrm{Agt}} \longrightarrow \mathrm{States}$
- $\forall A \in \mathrm{Agt} \quad \mathrm{Allow}_A : \mathrm{States} \longrightarrow 2^{\mathrm{Act}}$
- $\forall A \in \mathrm{Agt} \quad \phi_A : \mathrm{States}^{\omega} \longrightarrow \mathbb{R}$

Strategy and outcome

Definition

Let $h \in \mathrm{States}^+$ *be a history,*

$$\sigma_A(h) \in \mathrm{Dist}(\mathrm{Allow}_A(\mathrm{last}(h)))$$

We denote by $\mathbb{S}_A$ *the set of (randomized) strategies of player* A*, and* $\mathbb{S}$ *the set of strategy profiles*

Definition (Semantics)

If $h \in \mathrm{States}^+$ *is a history, and* $\sigma \in \mathbb{S}$*, the action chosen by each agent* A *is a random variable* $a_A \sim \sigma_A(h)$*.*
Next state is then $\mathrm{Tab}(\mathrm{last}(h), (a_A)_{A \in \mathrm{Agt}})$

Definition (Probability measure)

The game generates an infinite random run for every strategy profile σ *and initial history* h*. For every* $r \in \mathrm{States}^\omega$*, we denote by* $\mathbb{P}^\sigma(r \mid h)$ *the probability to get generate the run* r *from* h*.*
For every $h' \in \mathrm{States}^+$*, we note* $\mathbb{P}^\sigma(h' \mid h) = \mathbb{P}^\sigma(h' \mathrm{States}^\omega \mid h)$

Utility functions and expectation

We only consider terminal reward utilities.

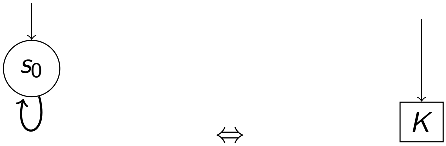

$$\forall h \in \mathrm{States}^* \quad \phi(h \cdot s_0^\omega) = K \in \mathbb{R}^{\mathrm{Agt}}$$

Definition (Conditionnal expectation)

For $\phi : \mathrm{States}^\omega \longrightarrow \mathbb{R}^{\mathrm{Act}}$ *with finite support, let*

$$\mathbb{E}^\sigma(\phi \mid h) = \sum_{x \in \mathrm{Img}(\phi)} x \cdot \mathbb{P}^\sigma\left(\phi^{-1}(x) \mid h\right)$$

Deviation and Nash equilibria

Definition

A Nash Equilibrium (NE) is a couple (σ, h) *where*

- $\sigma \in \mathbb{S}$
- $h \in \mathrm{States}^+$ *is an initial history*
- *No player can improve his utility by changing her own strategy (deviation),*

$$\forall A \in \mathrm{Agt} \quad \forall \sigma'_A \in \mathbb{S}_A \quad \mathbb{E}^{\sigma[A/\sigma'_A]}(\phi_A \mid h) \leq \mathbb{E}^{\sigma}(\phi_A \mid h)$$

When $\sigma(h)$ *is degenerated for all h, the equilibrium is said to be pure. Otherwise, this is a mixted equilibrium.*

Concurrent Games and Equilibria

Examples and tools

Existence of an equilibrium in terminal-reward games

Rock-paper-scissors

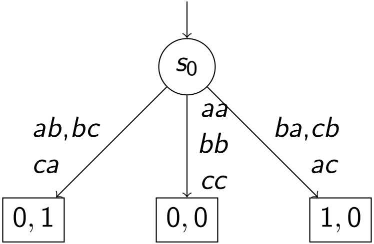

$$\sigma(s_0) = aa \Rightarrow \mathbb{E}^{\sigma}(\phi \mid s_0) = (0, 0)$$

Agent 1 can deviate:

$$\sigma_1'(s_0) = b \Rightarrow \mathbb{E}^{\sigma[1/\sigma_1']}(\phi \mid s_0) = (1, 0)$$

Rock-paper-scissors

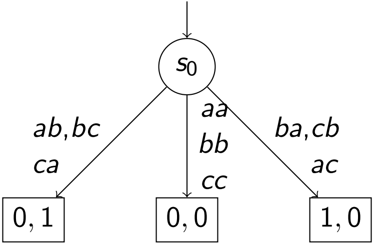

$$\sigma(s_0) \sim \mathcal{U}(\{a, b, c\})^2 \Rightarrow \mathbb{E}^{\sigma}(\phi \mid s_0) = \left(\frac{1}{3}, \frac{1}{3}\right)$$

is an equilibrium:

$$\forall i \quad \forall \sigma' \quad \mathbb{E}^{\sigma[i/\sigma']}(\phi \mid s_0) = \left(\frac{1}{3}, \frac{1}{3}\right)$$

One stage game with two players and two actions

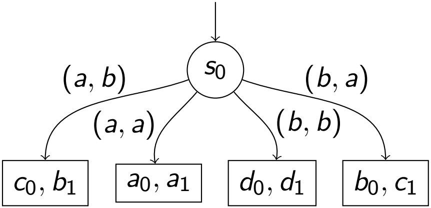

1 \ 0	a	b
a	a_0, a_1	b_0, c_1
b	c_0, b_1	d_0, d_1

$$\forall i \in \{0,1\} \quad \forall \sigma' \in \mathbb{S} \quad \mathbb{E}^{\sigma[i/\sigma']}(\phi_i \mid h) \leq \mathbb{E}^{\sigma}(\phi_i \mid h)$$

One stage game with two players and two actions

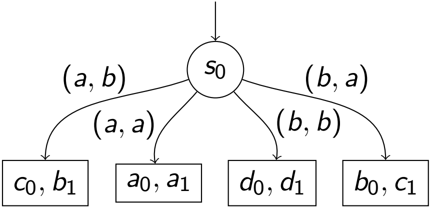

1 \ 0	a	b
a	a_0, a_1	b_0, c_1
b	c_0, b_1	d_0, d_1

$$\forall i\ [(d_i - c_i) + (a_i - b_i)] \cdot \sigma_{1-i}(a \mid s_0) = d_i - c_i$$

One stage game with two players and two actions

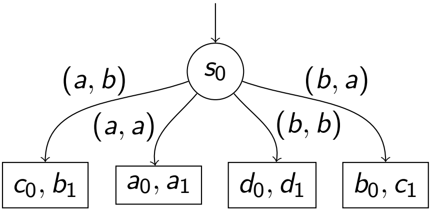

1 \ 0	a	b
a	a_0, a_1	b_0, c_1
b	c_0, b_1	d_0, d_1

$$\forall i \begin{cases} \sigma_i(a \mid s_0) < 1 \Rightarrow [(d_i - c_i) + (a_i - b_i)] \cdot \sigma_{1-i}(a \mid s_0) \leq d_i - c_i \\ \sigma_i(a \mid s_0) > 0 \Rightarrow [(d_i - c_i) + (a_i - b_i)] \cdot \sigma_{1-i}(a \mid s_0) \geq d_i - c_i \end{cases}$$

Infinite runs and abstractions

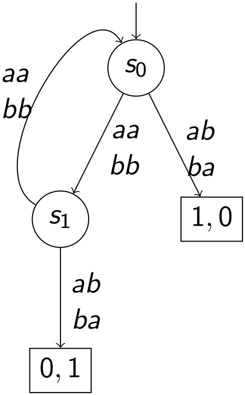

Infinite runs and abstractions

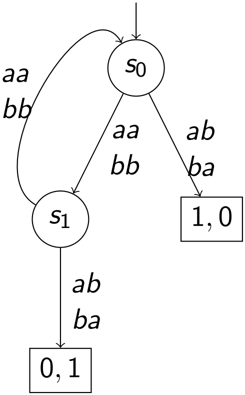

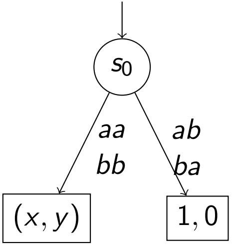

$$(x, y) = \mathbb{E}^{\sigma}(\phi \mid s_0 s_1)$$

Concurrent Games and Equilibria

Examples and tools

Existence of an equilibrium in terminal-reward games

Problem

Problem

Given a terminal-reward concurrent game $\mathcal{G}$, an initial state s_0 and $\varphi : \mathbb{R}^{\mathrm{Act}} \to \mathbb{R}$ a linear map,
does there exist a $\sigma \in \mathbb{S}$ such that (σ, s_0) is a Nash Equilibrium and $\varphi(\mathbb{E}^{\sigma}(\phi \mid s_0)) \geq 0$?

Problem

Problem

Given a terminal-reward concurrent game $\mathcal{G}$, an initial state s_0 and $\varphi : \mathbb{R}^{\mathrm{Act}} \to \mathbb{R}$ a linear map,
does there exist a $\sigma \in \mathbb{S}$ such that (σ, s_0) is a Nash Equilibrium and $\varphi(\mathbb{E}^{\sigma}(\phi \mid s_0)) \geq 0$?

Without constraints,

- It does exist an equilibrium in any finite stage game (Nash's theorem cf [Nas50])
- If we allow non-positive terminal-rewards, there exist concurrent games without equilibrium
- This still holds if we add Büchi conditions for the rewards
- If we restrict to non-negative rewards

Problem

Problem

Given a terminal-reward concurrent game $\mathcal{G}$, an initial state s_0 and $\varphi : \mathbb{R}^{\mathrm{Act}} \to \mathbb{R}$ a linear map,
does there exist a $\sigma \in \mathbb{S}$ such that (σ, s_0) is a Nash Equilibrium and $\varphi(\mathbb{E}^{\sigma}(\phi \mid s_0)) \geq 0$?

Without constraints,

- It does exist an equilibrium in any finite stage game (Nash's theorem cf [Nas50])
- If we allow non-positive terminal-rewards, there exist concurrent games without equilibrium
- This still holds if we add Büchi conditions for the rewards
- If we restrict to non-negative rewards. . . ?

Undecidability of the problem in turn-based games

Theorem ([UW11])

The constrained existence problem is undecidable for 14*-player turn-based games with terminal rewards.*
Turn-based: $\forall s \in \mathrm{States} \quad |\{A \mid |\mathrm{Allow}_A(s)| > 1\}| \leq 1$

- Reduction to a 2-counter machine
- c_i as expected reward $\frac{1}{2^{c_i}}$
- (Safety) constraint on one player
- Complex proof (14 agents required)
- Concurrency not exploited

2-counter machine

Definition

(Q, q_0, Δ) *where*

- Q *finite set of states*
- $q_0 \in Q$ *initial state*
- $\Delta \subseteq Q \times \Gamma \times Q$ *transition table with*
  $\Gamma = \{inc(j), dec(j), zero(j) \mid j \in \{1, 2\}\}$

We note $q\Delta = \{(\gamma, q') \mid (q, \gamma, q') \in \Delta\}$

Counting modules

- Goal: encode (c_1, c_2) as expected reward $\frac{1}{2^{c_1} 3^{c_2}}$
- Zero test ?
- We build a game with a countable number of equilibria
- $\mathrm{Agt} = \{0, 1, 2\}$, 1 and 2 are antagonistic
- Safety condition: "Agent 0 should get an expected reward of 1"

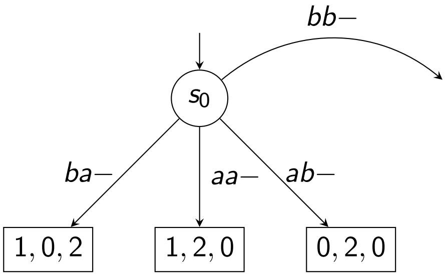

Counting modules

- Goal: encode (c_1, c_2) as expected reward $\frac{1}{2^{c_1}3^{c_2}}$
- Zero test ?
- We build a game with a countable number of equilibria
- $\mathrm{Agt} = \{0, 1, 2\}$, 1 and 2 are antagonistic
- Safety condition: "Agent 0 should get an expected reward of 1"

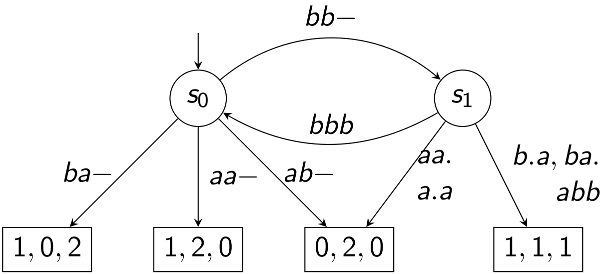

Counting modules (iteration)

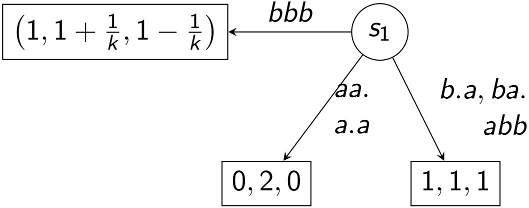

1 \ 0	a	b
a	$0, 2$	$1, 1$
b	$1, 1$	$1, 1+\frac{1}{k}$

Figure: Projection on players $0, 1$ (assuming 2 plays b)

Counting modules (iteration)

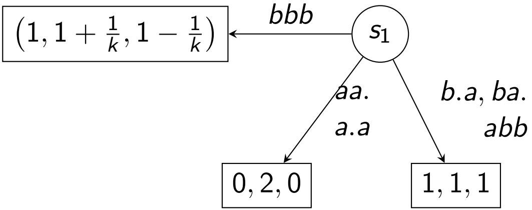

1 \ 0	a	b
a	$0, 2$	$1, 1$
b	$1, 1$	$1, 1+\frac{1}{k}$

Figure: Projection on players 0, 1 (assuming 2 plays *b*)

$$\Rightarrow \sigma_0(a) \leq \frac{1+\frac{1}{k}-1}{1+\frac{1}{k}-1+2-1} = \frac{1}{k+1} \Rightarrow \mathbb{E}^{\sigma}(\phi_1) \leq 1 + \frac{1}{k+1}$$

Countable set of equilibria

Theorem

If we restrict to safe equilibria, then:
the set of expected rewards of the possible Nash equilibria is exactly the set $I_r = \left\{ \left(1, 1 + \frac{1}{n+1}, 1 - \frac{1}{n+1}\right) \;\middle|\; n \in \overline{\mathbb{N}} \right\}$.

Countable set of equilibria

Theorem

If we restrict to safe equilibria, then:
the set of expected rewards of the possible Nash equilibria is exactly the set $I_r = \left\{ \left(1, 1 + \frac{1}{n+1}, 1 - \frac{1}{n+1}\right) \;\middle|\; n \in \overline{\mathbb{N}} \right\}$.

With the same method, we build a game whose set of equilibria is:

$$I_r = \left\{ \left(1, 1 + \frac{1}{2^n}, 1 - \frac{1}{2^n}\right) \;\middle|\; n \in \overline{\mathbb{N}} \right\}$$

Tweaking the modules

We duplicate agents 1 and 2, so we got 5 agents: $0, A_1^0, A_2^0, A_1^1, A_2^1$.
Last two players have linear combinations of the previous payoffs.

Tweaking the modules

We duplicate agents 1 and 2, so we got 5 agents: $0, A_1^0, A_2^0, A_1^1, A_2^1$.
Last two players have linear combinations of the previous payoffs.
For instruction $\mathrm{dec}(1)$, we have:

$$I_r = \left\{ \left(1, 1 + \frac{1}{n+1}, 1 - \frac{1}{n+1}, 1 - \frac{1}{n+1}, 1 + \frac{1}{n+1}\right) \;\middle|\; n \in \overline{\mathbb{N}} \right\}$$

.

One step reduction

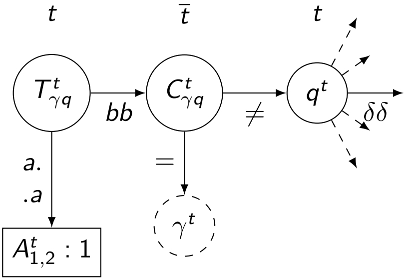

Where

- γ^t is one of the *tweaked* module
- states annotated with t are controlled by A^t_1 and A^t_2
- $\delta = \gamma' q' \in q\Delta$ is the next transition

One step reduction

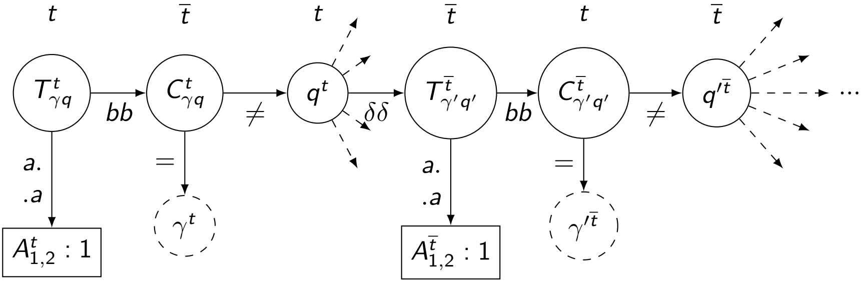

Where

- γ^t is one of the *tweaked* module
- states annotated with t are controlled by A^t_1 and A^t_2
- $\delta = \gamma' q' \in q\Delta$ is the next transition

One step reduction (sketch for dec(1))

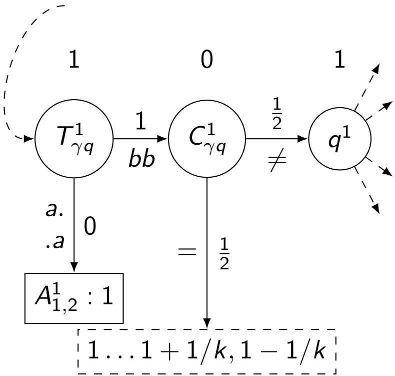

One step reduction (sketch for $\mathrm{dec}(1)$)

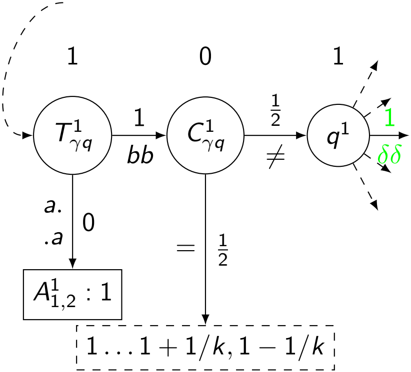

One step reduction (sketch for $\mathrm{dec}(1)$)

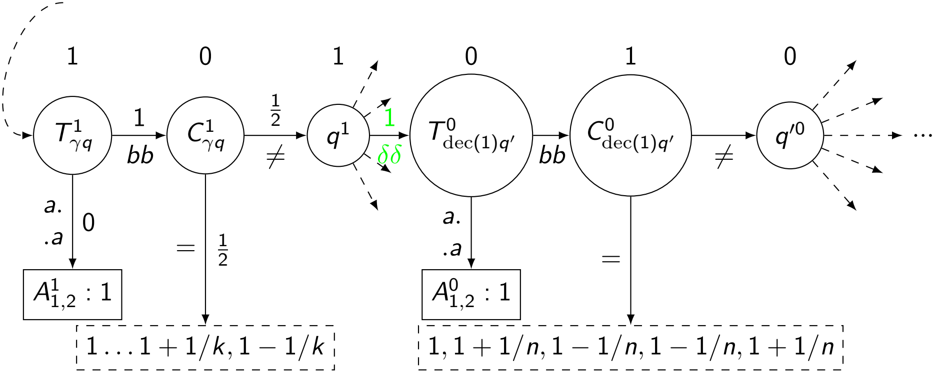

One step reduction (sketch for $\mathrm{dec}(1)$)

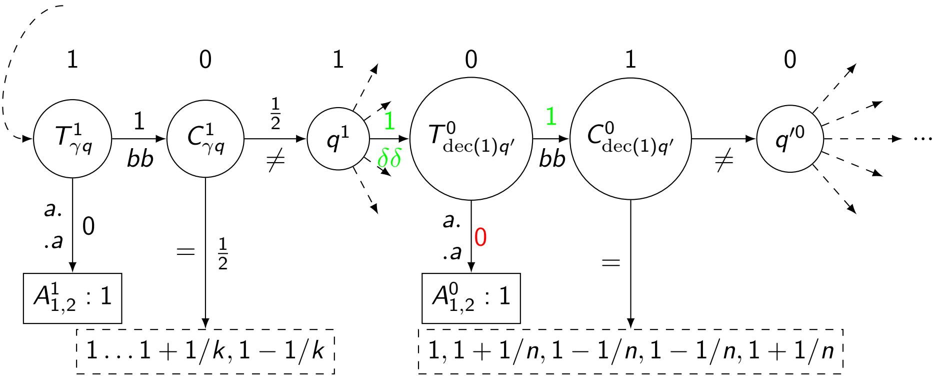

One step reduction (sketch for $\mathrm{dec}(1)$)

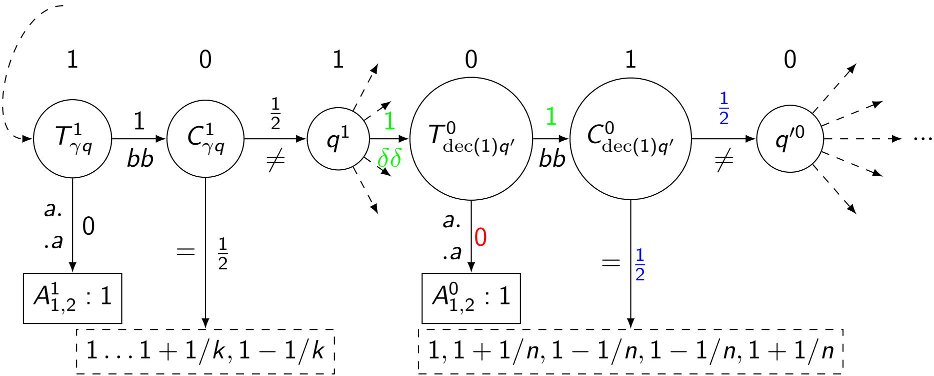

One step reduction (sketch for $\mathrm{dec}(1)$)

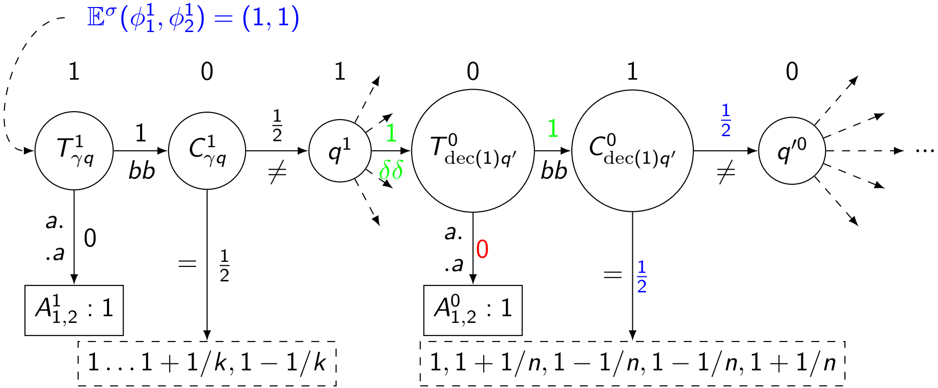

One step reduction (sketch for $\mathrm{dec}(1)$)

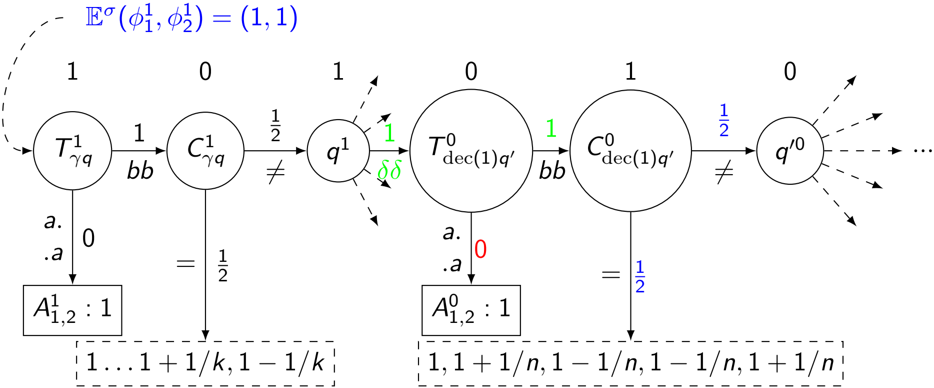

We should have:

$$1 = \frac{1}{2}\left(1 + \frac{1}{k}\right) + \frac{1}{4}\left(1 - \frac{1}{n}\right) + \frac{1}{4} \Rightarrow n = \frac{k}{2}$$

Theorem

The constrained existence problem is undecidable for 5-player concurrent games.

Theorem

The constrained existence problem is undecidable for 5*-player concurrent games.*

Still hope for 2-player games !

Conclusion

- Study of simple cases
- Build of simple modules with non-trivial equilibria
- Reduction to 5-player games using concurrency
- Lower bound on the number of players ?

Outlooks

- Reduction to 3 players ?
- Regularity with no-constraint ? (closure)
- Restriction to simpler strategy profiles (computable, automaton ?)
- Restriction to weaker winning conditions
- Adapt the tools for pure Nash Equilibria (eg Suspect Game [Bre12])

Thank you for your attention

Bibliography I

Romain Brenguier.
Équilibres de Nash dans les Jeux Concurrents – Application aux Jeux Temporisés.
Thèse de doctorat, Laboratoire Spécification et Vérification, ENS Cachan, France, November 2012.

John F. Nash.
Equilibrium points in n-person games.
Proceedings of the National Academy of Sciences, 36(1):48–49, 1950.

Bibliography II

Michael Ummels and Dominik Wojtczak.
The complexity of Nash equilibria in limit-average games.
In Joost-Pieter Katoen and Barbara König, editors, *Proceedings of the 22nd International Conference on Concurrency Theory (CONCUR'11)*, volume 6901 of *Lecture Notes in Computer Science*, pages 482–496, Aachen, Germany, September 2011. Springer.