Nash Equilibria in Symmetric Games with Partial Observation

Patricia Bouyer [1] Nicolas Markey [1] Steen Vester [2]

[1]CNRS, LSV & ENS Cachan

[2]Technical University of Denmark

April 5, 2014

Introduction

We introduce a game model where

- There can be many players, each one having its own objectives
- The game is the same from the point of view of every player
- Players have partial information

As a starting point, we study the problem of finding pure Nash equilibria as well as pure symmetric Nash equilibria in this model for qualitative objectives.

Motivation

A set of **mobile phones** competing for **bandwidth**

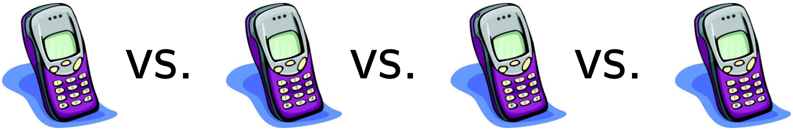

Example

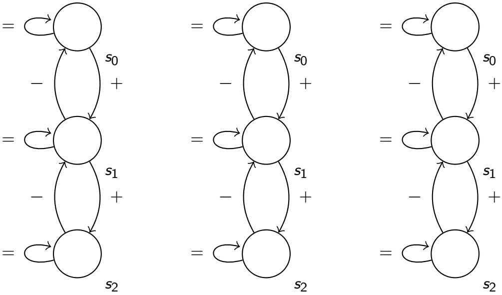

Example

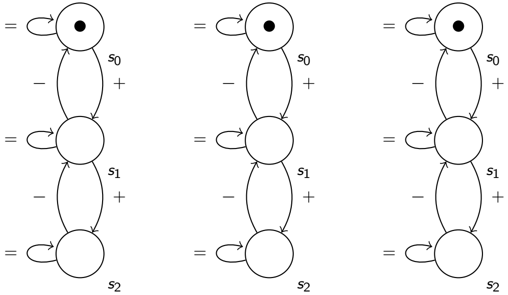

$$\rho = (s_0, s_0, s_0)$$

Example

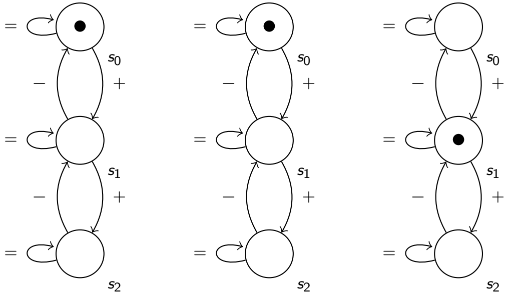

$$\rho = (s_0, s_0, s_0)(s_0, s_0, s_1)$$

Example

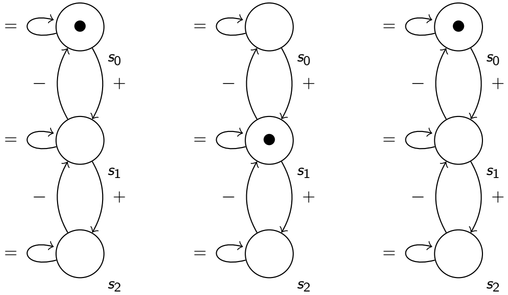

$$\rho = (s_0, s_0, s_0)(s_0, s_0, s_1)(s_0, s_1, s_0)$$

Example

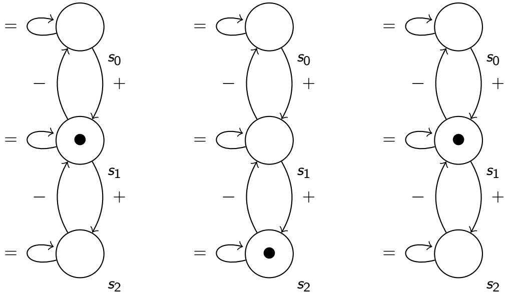

$$\rho = (s_0, s_0, s_0)(s_0, s_0, s_1)(s_0, s_1, s_0)(s_1, s_2, s_1)$$

Example

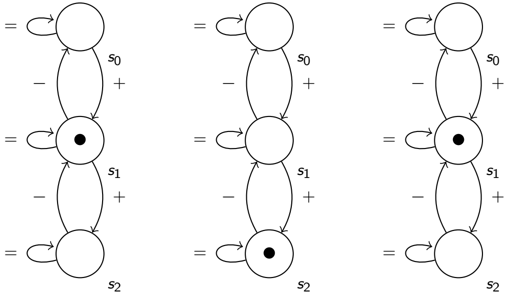

$$\rho = (s_0, s_0, s_0)(s_0, s_0, s_1)(s_0, s_1, s_0)(s_1, s_2, s_1)^\omega$$

Arena

An arena is a tuple $\langle \text{States}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab} \rangle$ where

- States is a finite set of states
- Agt is a finite set of agents
- Act is a finite set of actions
- $\text{Mov}\colon \text{States} \times \text{Agt} \to 2^{\text{Act}} \setminus \{\emptyset\}$ is the set of actions available to a given player in a given state
- $\text{Tab}\colon \text{States} \times \text{Act}^{\text{Agt}} \to \text{States}$ is a transition function that specifies the next state, given a state and an available action of each player

Histories, outcomes and strategies

Let $\mathcal{G} = \langle \text{States}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab} \rangle$ be an n-player arena with $\text{Agt} = [n] = \{0, ..., n-1\}$. Then

- A history is a finite sequence $s_0 s_1 ... s_\ell \in \text{States}^+$ so there exists $m_0, ... m_\ell \in \text{Act}^n$ where $\text{Tab}(s_i, m_i) = s_{i+1}$ for all $0 \leq i < \ell$
- Hist is the set of histories
- An outcome is an infinite sequence $s_0 s_1 ... \in \text{States}^\omega$ so there exists $m_0, m_1, ... \in \text{Act}^n$ where $\text{Tab}(s_i, m_i) = s_{i+1}$ for all $i \geq 0$
- Out is the set of outcomes
- A strategy for $j \in \text{Agt}$ is a mapping $\sigma : \text{Hist} \to \text{Act}$ such that $\sigma(h) \in \text{Mov}(\text{last}(h), j)$ for all $h \in \text{Hist}$
- A strategy profile is an n-tuple $(\sigma_0, ..., \sigma_{n-1})$ where σ_j is a strategy for player j for all $j \in [n]$

Winning conditions and Nash equilibria

- A winning condition is a subset $\Omega \subseteq \mathrm{Out}$ of the set of outcomes
- The unique outcome of a strategy profile σ from state s is denoted $\mathrm{Out}(s, \sigma)$
- Let Ω_j be a winning condition for each player j. Then a Nash equilibrium from state s is a strategy profile $\sigma = (\sigma_1, ..., \sigma_n)$ such that for every player j and every strategy σ'_j for player j we have

$$\mathrm{Out}(s, \sigma[\sigma_j \mapsto \sigma'_j]) \in \Omega_j \Rightarrow \mathrm{Out}(s, \sigma) \in \Omega_j$$

n-player game network

An n-player game network is a tuple $\mathcal{G} = \langle G, (\equiv_i)_{i \in [n]}, (\Omega_i)_{i \in [n]} \rangle$ where

- $G = \langle \text{States}, \{a\}, \text{Act}, \text{Mov}, \text{Tab} \rangle$ is a 1-player arena
- $\equiv_i$ is an equivalence relation on States^n for each $i \in [n]$ modelling partial information
- $\Omega_i \in (\text{States}^n)^\omega$ is a winning condition for each $i \in [n]$

Product game semantics

Semantics of $\mathcal{G}$ given by the product game
$\mathcal{G}' = \langle \text{States}', [n], \text{Act}, \text{Mov}', \text{Tab}', (\Omega_i)_{i \in [n]} \rangle$ where

- $\text{States}' = \text{States}^n$
- $\text{Mov}'((s_0, \ldots, s_{n-1}), i) = \text{Mov}(s_i, a)$
- $\text{Tab}'((s_0, \ldots, s_{n-1}), (m_i)_{i \in [n]}) = (\text{Tab}(s_0, m_0), \ldots, \text{Tab}(s_{n-1}, m_{n-1}))$

$\equiv_i$ induces equivalence classes called information sets
A strategy σ_i for player i is $\equiv_i$-realisable if $\sigma_i(h) = \sigma_i(h')$ for all $\rho, \rho' \in \text{Hist}$ with $\rho \equiv_i \rho'$.

Example

Card-game tournament

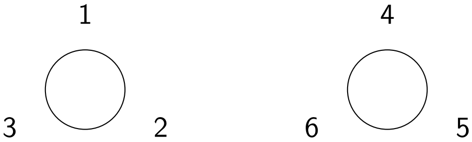

Define symmetry between player 1 and 5 by:
$\pi_{1,5}(1) = 5$, $\pi_{1,5}(2) = 6$, $\pi_{1,5}(3) = 4$
$\pi_{1,5}(4) = 1$, $\pi_{1,5}(5) = 2$, $\pi_{1,5}(6) = 3$

Intuitively, $\pi_{i,j}(k) = \ell$ means that k has the same role from i's point of view as ℓ does from j's point of view.

Symmetric game networks

A game network $\mathcal{G} = \langle G, (\equiv_i)_{i\in[n]}, (\Omega_i)_{i\in[n]} \rangle$ is symmetric if there are permutations $\pi_{i,j}$ of $[n]$ for all $i, j \in [n]$ such that for all $i, j, k \in [n]$

1. $\pi_{i,j}(i) = j$, $\pi_{i,i}$ is the identity and $\pi_{k,j} \circ \pi_{i,k} = \pi_{i,j}$
2. Observations are compatible with symmetry: $t \equiv_i t'$ iff $t(\pi_{i,j}^{-1}) \equiv_j t'(\pi_{i,j}^{-1})$
3. Winning conditions are compatible with symmetry: $\rho \in \Omega_i$ iff $\rho(\pi_{i,j}^{-1}) \in \Omega_j$

where $t(\pi) = (t_{\pi(0)}, \ldots, t_{\pi(n-1)})$ when $t = (t_0, \ldots, t_{n-1}) \in \text{States}^n$ and π is a permutation of $[n]$.

$(\pi_{i,j})_{i,j\in[n]}$ is called a symmetric representation of $\mathcal{G}$

Symmetric strategy profiles

Let $\mathcal{G}$ be a symmetric n-player game network with symmetric representation $\pi = (\pi_{i,j})_{i,j \in [n]}$.
A strategy profile $\sigma = (\sigma_i)_{i \in [n]}$ is symmetric wrt π if

- σ_i is $\equiv_i$-realisable for all $i \in [n]$
- $\sigma_i(\rho) = \sigma_j(\rho(\pi_{i,j}^{-1}))$ for all $i, j \in [n]$ and every history ρ

Lemma
If we fix a symmetric representation π and σ_0 is an $\equiv_0$-realisable strategy for player 0, then $\sigma = (\sigma_i)_{i \in [n]}$ defined by $\sigma_i(\rho) = \sigma_0(\rho(\pi_{i,0}^{-1}))$ for all ρ is symmetric wrt π.

Problems considered

The (symmetric) constrained existence problem

Input: Symmetric game network $\mathcal{G}$, a symmetric representation π, an initial configuration t, a set $L \subseteq [n]$ of losing players and a set $W \subseteq [n]$ of winning players

Output: YES, if there is a (symmetric) Nash equilibrium σ in $\mathcal{G}$ from t (wrt. π) such that players in L lose and players in W win. NO, otherwise.

If $L = W = \emptyset$ it is called the (symmetric) existence problem
If $W = [n]$ it is called the positive (symmetric) existence problems

Symmetric (contrained) existence is not easier

Proposition

(Constrained) existence can be reduced to symmetric (contrained) existence.

The reduction neither changes the number of players nor the types of objectives. It only increases the size of the arena polynomially.

Undecidability results

Result 1
The (constrained) existence of a (symmetric) Nash equilibrium for non-regular objectives in 2-player perfect information symmetric game networks is undecidable

Result 2
The (constrained) existence of a (symmetric) Nash equilibrium for LTL objectives in a 3-player symmetric game network is undecidable

Result 3
The (constrained) existence of a parameterized symmetric Nash equilibrium for LTL objectives in symmetric game networks is undecidable (even for memoryless strategies)

A decidability result

The (constrained) existence of a bounded-memory (symmetric) Nash equilibrium for LTL objectives in symmetric game networks is EXPSPACE-complete

Concluding remarks

We have

- Introduced a framework for modelling games with symmetry between players
- Shown that (constrained) existence of symmetric Nash equilibria is as hard as (constrained) existence of Nash equilibria
- Obtained preliminary results on (un)decidability and complexity of (constrained) existence of (symmetric) Nash equilibria in the model

However, there are many open questions to analyze.