Algorithmic Game Theory - Complexity Classes and Nash

Branislav Bošanský

Czech Technical University in Prague

branislav.bosansky@agents.fel.cvut.cz

March 4, 2019

Complexity Classes

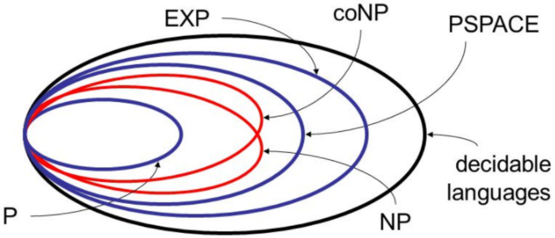

all containments believed to be proper

Main Complexity Classes in Algorithmic Game Theory

$\mathcal{P}$

- polynomial problems (linear programming of polynomial size, etc.)
- computing a Nash Equilibrium in zero-sum games (normal-form, extensive-form)
- computing a Correlated Equilibrium in general-sum games (normal-form[1])
- computing a Stackelberg Equilibrium in general-sum games (normal-form, simple security games)
- computing equilibria in many of the instances from succinctly represented games (we shall see)

[1] open for extensive-form games

Main Complexity Classes in Algorithmic Game Theory

$\mathcal{NP}$

- NP-hard problems, (mixed-integer linear programs, etc.)
- computing some specific Nash Equilibrium in general-sum games (normal-form, extensive-form)
- computing some specific Correlated-based Equilibrium in general-sum games (extensive-form games)
- computing a Stackelberg Equilibrium in general-sum games (extensive-form games)

Main Complexity Classes in Algorithmic Game Theory

the search problem problem that asks for *any Nash Equilibrium* is a different, potentially an 'easier' problem

there is a finer description of the complexity classes

Main Complexity Classes in Algorithmic Game Theory

1. equilibria are guaranteed to exist (i.e., total problems TFNP $\subseteq$ FNP ("a function extension of a decision problem in NP")
2. we can search for them
   - pure strategy profiles
   - support enumeration

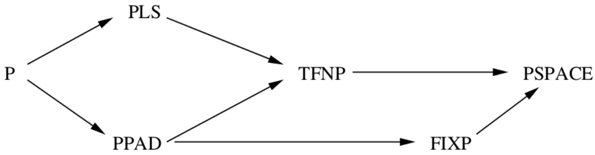

Fig. 2 – Relations between the complexity classes. [2]

[2] Figure from M. Yannakakis "Equilibria, fixed points, and complexity classes" Computer Science Review (3) 71–85, 2009

Polynomial Local Search ($\mathcal{PLS}$)

- consider an instance I of an optimization problem, $S(I)$ is a set of candidate solutions, $p_I(s)$ is a cost (or utility) associated with candidate $s \in S(I)$ that has to be minimized (or maximized, respectively)
- each candidate $s \in S(I)$ has a neighborhood $N_I(s) \subseteq S(I)$
- a candidate s is locally optimal (cost-wise) if

$$p_I(s) \leq p_I(s') \; \forall s' \in N_I(s)$$

- Sol(I) is a set of locally optimal solutions
- every step of the algorithm (generating starting solution, computing the cost, getting a better neighbor) is polynomial, but there can be an exponential number of steps

Polynomial Local Search ($\mathcal{PLS}$) (2)

- several well-known problems of this kind
- finding a local optimum in Traveling Salesman Problem, Max Cut, Max Sat, ...
  - we define a neighborhood function (e.g., *2-Opt*) and perform a greedy search
- finding a stable configuration of a neural network
- finding a pure equilibrium when it is guaranteed to exist
  - computing an optimal strategy in simple stochastic games, where pure stationary strategy is known to be optimal (the problem is in PLS, but it is open whether it is in P, or not)
  - some other variants of stochastic games (mean payoff/parity games with no chance)

From PLS to PPAD

Searching for pure Nash equilibria is not sufficient.

Pure equilibria do not have to exist.

What can we search for in mixed strategies?

Is this problem in PLS? Can we redefine the problem of finding a mixed NE as a PLS?

Nash and Fixed Points

Theorem (Brouwer's Fixed Point Theorem)

Let X be a convex and compact set in a n-dimensional Euclidean space, and let $f : X \to X$ be a continuous function. Then there exists a point $x \in X$ such that $f(x) = x$. Such a point is called a fixed point of f.

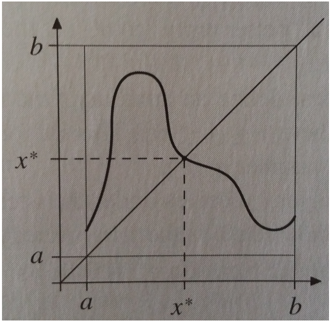

(Brief) Proof of Existence of Nash in Finite Games using Brouwer's Fixed Point Theorem

Let Σ be a mixed strategy profile $\Sigma = \Sigma_1 \times \Sigma_2 \times \ldots \Sigma_n$ (i.e., convex and compact subset of Euclidean space). We will define a (continuous) function $f : \Sigma \rightarrow \Sigma$ and show that every fixed point of f is an equilibrium of the game.

We can use *regret* function that specifies how much player i can gain by switching to pure strategy j:

$$g_i^j(\sigma) := \max\left\{0, u_i(s_i^j, \sigma_{-i}) - u_i(\sigma_i, \sigma_{-i})\right\}$$

We want to get $g_i^j(\sigma) = 0 \; \forall j \in \mathcal{S}_i$ (no regret).

(Brief) Proof of Existence of Nash in Finite Games using Brouwer's Fixed Point Theorem (2)

We can now define function f, such that if σ is an equilibrium profile, then $f(\sigma) = \sigma$;

$$f_i^j(\sigma) := \frac{\sigma_i(s_i^j) + g_i^j(\sigma)}{1 + \sum_{k=1}^{m_i} g_i^k(\sigma)}$$

We need to show the other implication. For fixed point σ it holds

$$g_i^j(\sigma) = \sigma_i(s_i^j) \sum_{k=1}^{m_i} g_i^k(\sigma)$$

(Brief) Proof of Existence of Nash in Finite Games using Brouwer's Fixed Point Theorem (3)

Suppose that fixed point σ is not an equilibrium. There must exist a pure strategy $l \in \{1, \ldots, m_i\}$ such that $g_i^l(\sigma) > 0$ and consequently from the previous slide we know that $\sigma_i(s_i^l) > 0$. Now

$$u_i(\sigma) = \sum_{j=1}^{m_i} \sigma_i(s_i^j) u_i(s_i^j, \sigma_{-i}) \tag{1}$$

$$0 = \sum_{j=1}^{m_i} \sigma_i(s_i^j) \Big(u_i(s_i^j, \sigma_{-i}) - u_i(\sigma) \Big) \tag{2}$$

$$0 = \sum_{j:\sigma_i(s_i^j)>0} \sigma_i(s_i^j) g_i^j(\sigma), \tag{3}$$

where the last summand is strictly positive due to pure strategy l, which is a contradiction. $\square$

Generalization of Nash's Theorem

The set of strategies of game G does not have to be a probability distribution, but generally a convex set (polytope; recall *convex games*).

We can create an auxiliary game G', where pure strategies will be vertexes of the polytope of the convex game and use the original Nash's Theorem.

Finally, we translate the equilibrium strategies from G' to G and show that they must form an equilibrium in G.

Discretized variant – Sperner's Lemma and Scarf's algorithm

Sperner's Lemma (2D):
Given a triangle ABC, and a triangulation T of the triangle, the set S of vertices of T is colored with three colors in such a way that:

1. A, B, and C are colored 1, 2, and 3 respectively
2. Each vertex on an edge of ABC is to be colored only with one of the two colors of the ends of its edge. For example, each vertex on AC must have a color either 1 or 3.

Then there exists a triangle from T, whose vertices are colored with the three different colors.

More precisely, there must be an odd number of such triangles.

Discretized variant – Sperner's Lemma and Scarf's algorithm

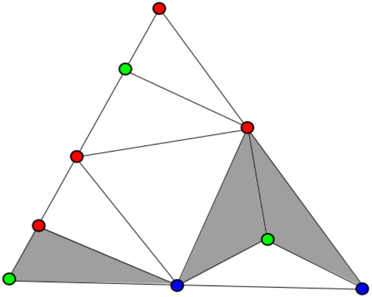
[3]

[3] Figure from Wikipedia.

Discretized variant – Sperner's Lemma and Scarf's algorithm

Scarf's Algorithm for approximating fixed points of Brouwer function F: $\Delta_n \to \Delta_n$:

1. Subdivide the simplex Δ_n into "small" subsimplices of diameter $\delta > 0$ (depending on the "modulus of continuity" of F, and on $\epsilon > 0$).
2. Color every vertex, z, of every subsimplex with a color $i = \min\{i \mid z_i > 0 \,\&\, F(z)_i \leq z_i\}$.
3. By Sperner's Lemma there must exist a panchromatic subsimplex. (And the proof provides a way to "navigate" toward such a simplex.)
4. Fact: If $\delta > 0$ is chosen such that $\delta \leq \epsilon/2n$ and $\forall x, y \in \Delta_n, ||x - y||_\infty < \delta \to ||F(x) - F(y)||_\infty < \epsilon/2n$, then all the points in a panchromatic subsimplex are weak ϵ-fixed points.

[4]

[4] From slides by K. Etessami: Tutorial on GAMES' 08

Discretized variant – Sperner's Lemma and Scarf's algorithm

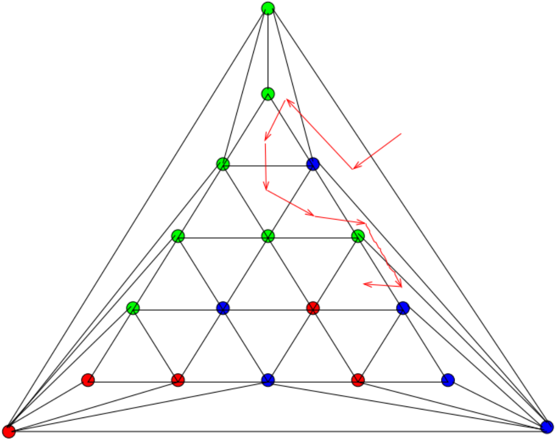

5

[5]Figure from slides by K. Etessami: Tutorial on GAMES' 08.

Polynomial Parity Arguments on Directed graphs (PPAD)

there is a set of candidate solutions $S(I)$ for an instance I and polynomial-time algorithms:

- compute an initial candidate solution $s_0 \in S(I)$
- given I and s test whether $s \in S(I)$ and if so compute a successor $\mathsf{succ}_I(s) \in S(I)$ and a predecessor $\mathsf{pred}_I(s) \in S(I)$
  - $\mathsf{pred}_I(s_0) = s_0$
  - $\mathsf{succ}_I(s_0) \neq s_0$
  - $\mathsf{pred}_I(\mathsf{succ}_I(s_0)) = s_0$

$\mathsf{Sol}(I)$ is a set of nodes where indegree $+$ outdegree $= 1$.

PPAD – End of the Line (2)

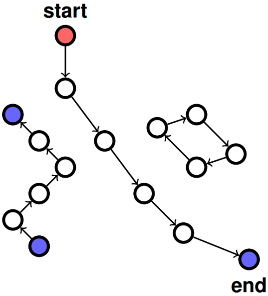

[6]

Given a graph G of indegree/outdegree at most 1, and a strat node of indegree 0 and outdegree 1, find another node degree of 1.

Blue nodes are $\mathsf{Sol}(I)$.

[6]Figure from R. Savani "Polymatrix Games" Tutorial at WINE 2015

PPAD and equilibria

Theorem ([1],[5])

It is PPAD-complete to compute an exact Nash equilibrium of a bimatrix game.

An alternative proof of the existence of a Nash equilibrium is based on Lemke-Howson algorithm.

PPAD and equilibria (2)

Many follow-up results of the completeness theorem:

- computing an exact Nash equilibrium for a two-player extensive-form game is PPAD-complete [4]
- computing an exact Nash equilibrium for a two-player normal-form game is PPAD-complete even if all the payoffs are 0 and 1 (so called *win-lose games*) [3]
- computing ϵ-Nash equilibrium for an n-player game is PPAD-complete [2][7]

[7] Approximation in a weak sense.

FIXP

What are the computational challenges moving to n-player games when computing Nash equilibria?

- knowing the support does not help in computing one
- Nash equilibria use in general irrational numbers

1	Left	Right
Top	3, 0, 2	0, 2, 0
Bottom	0, 1, 0	1, 0, 0

2	Left	Right
Top	1, 0, 0	0, 1, 0
Bottom	0, 3, 0	2, 0, 3

$p(L) = (-13 + \sqrt{601})/24$ ...

Theorem (Bubelis 1979)

Every real algebraic number can be "encoded" in a precise sense as the payoff to player 1 in a unique NE of a 3-player game.

FIXP

K. Etessami and M. Yannakakis [6] defined a new class FIXP:
Input: an *algebraic circuit* over basis $\{+, *, -, /, \max, \min\}$ with rational constants, having n input variables and n outputs, such that the circuit represents a continuous function $F : [0,1]^n \to [0,1]^n$.
Output: Compute (or strong ϵ-approximate) a fixed point of F.

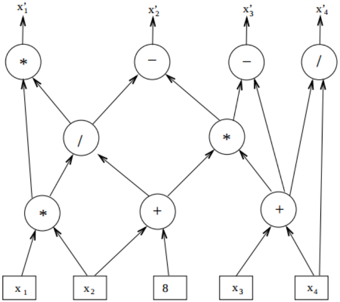

FIXP

The most famous problem in this class is the **square-root sum problem**:

Sqrt-Sum

Given $(d_1, \ldots, d_n) \in \mathbb{N}^n$ and $k \in \mathbb{N}$, decide whether $\sum_{i=1}^{n} \sqrt{d_i} \leq k$.

It is known to be solvable in PSPACE.

Theorem (Etessami and Yannakakis, 2007)

Any non-trivial approximation of an actual NE solves Sqrt-Sum

FIXP

Theorem (Etessami and Yannakakis, 2007)

Computing a 3-player Nash Equilibrium is **FIXP***-complete.*

The completeness holds in several senses:

- exact (real-valued) computation;
- strong ϵ-approximation,
- decision version of the problem – given a game $\mathcal{G}$, rational value $q \in \mathbb{Q}$, and coordinate i: if for all NEs x^*, $x_i^* \geq q$, then answer "Yes"; if for all NEs x^*, $x_i^* < q$, then the answer is "No". Otherwise, any answer is fine.

FIXP-completeness

Proof Sketch:

- Suppose we could create a (3-player) game such that, in any NE, Player 1 plays strategy A with probability $> 1/2$ iff $\sum_i \sqrt{d_i} > k$ and with probability $< 1/2$ iff $\sum_i \sqrt{d_i} < k$. (Suppose equality can't happen.)
- Add an extra player with 2 strategies, who gets high payoff if it "guesses correctly" whether player 1 plays pure strategy A, and low payoff otherwise.
- In any NE, the new player will play one of its two strategies with probability 1.

Deciding which solves SqrtSum

PPAD as an algebraic circuit

Theorem (Etessami and Yannakakis, 2007)

Let linear-FIXP denote the subclass of FIXP where the algebraic circuits are restricted to basis $\{+, \max\}$ *and multiplication by rational constants only. Then, the following are all equivalent:*

1. *PPAD*
2. *linear-FIXP*
3. *exact fixed point problems for "polynomial piecewise-linear functions"*

References I

(besides the books)

X. Chen and X. Deng.
Settling the complexity of two-player nash equilibrium.
In *IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 261–272, 2006.

X. Chen, X. Deng, and S.-H. Teng.
Computing Nash equilibria: Approximation and smoothed complexity.
In *Proc. 47th IEEE FOCS*, 2006.

X. Chen, S.-H. Teng, and P. Valiant.
The approximation complexity of winlose games.
In *Proc. 18th ACM SODA*, 2007.

C. Daskalakis, A. Fabrikant, and C. H. Papadimitriou.
The Game World Is Flat: The Complexity of Nash Equilibria in Succinct Games.
In *ICALP*, pages 513–524, 2006.

References II

C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou.
The Complexity of Computing a Nash Equilibrium.
In *Proceedings of the 38th annual ACM symposium on Theory of computing*, 2006.

K. Etessami and M. Yannakakis.
On the complexity of nash equilibria and other fixed points.
In *FOCS*, 2007.

Z. H. Gumus and C. A. Floudas.
Global optimization of mixed-integer bilevel programming problems.
Computational Management Science, 2:181–212, 2005.