The Computational Complexity of Nash Equilibria and Fixed Points of Algebraic Functions

Kousha Etessami

LFCS, School of Informatics
University of Edinburgh

Open Lectures, Warsaw University
May 29-30, 2009

(based on joint work with Mihalis Yannakakis, Columbia U.)

some standard complexity classes

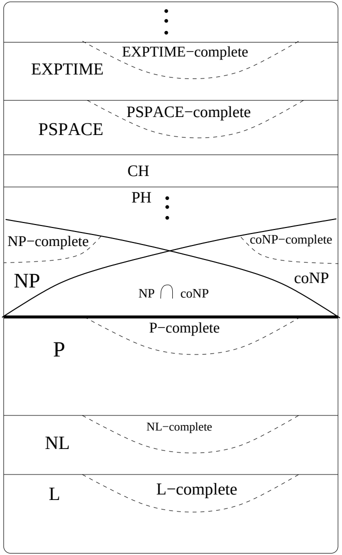

- The most important dichotomy in complexity theory is, of course:

  NP-hard ("intractable") vs. in **P** ("tractable")

- The most important dichotomy in complexity theory is, of course:

  NP-hard ("intractable") vs. in **P** ("tractable")

- But several fundamental problems in game theory, economics, and probability theory, and other fields, have resisted such a classification:

  they are neither known to be **NP**-hard, nor known to be in **P**.

- The most important dichotomy in complexity theory is, of course:

  NP-hard ("intractable") vs. in **P** ("tractable")

- But several fundamental problems in game theory, economics, and probability theory, and other fields, have resisted such a classification:

  they are neither known to be **NP**-hard, nor known to be in **P**.

- This hasn't stopped complexity theorists from trying to "classify" them.

Appetizer

What is the complexity of each of the following search problems?

Appetizer

What is the complexity of each of the following search problems?

a. (Nash, 1950) Given a finite game, and $\epsilon > 0$, compute a vector x' (a mixed strategy profile) that is within distance ϵ of some (exact) Nash Equilibrium.

Appetizer

What is the complexity of each of the following search problems?

a. (Nash, 1950) Given a finite game, and $\epsilon > 0$, compute a vector x' (a mixed strategy profile) that is within distance ϵ of some (exact) Nash Equilibrium.

b. (Shapley, 1953) Given an instance of Shapley's stochastic game, and $\epsilon > 0$, approximate the *value* of the game to within distance ϵ.

Appetizer

What is the complexity of each of the following search problems?

a. (Nash, 1950) Given a finite game, and $\epsilon > 0$, compute a vector x' (a mixed strategy profile) that is within distance ϵ of some (exact) Nash Equilibrium.

b. (Shapley, 1953) Given an instance of Shapley's stochastic game, and $\epsilon > 0$, approximate the *value* of the game to within distance ϵ.
Note:
`Parity-Games` $\leq_p$ `Mean-Payoff-Games` $\leq_p$
`Simple-Stochastic-Games` $\leq_p$
`Approximate-Shapley's-Stochastic-Games`

Appetizer

What is the complexity of each of the following search problems?

a. (Nash, 1950) Given a finite game, and $\epsilon > 0$, compute a vector x' (a mixed strategy profile) that is within distance ϵ of some (exact) Nash Equilibrium.

b. (Shapley, 1953) Given an instance of Shapley's stochastic game, and $\epsilon > 0$, approximate the *value* of the game to within distance ϵ.
Note:
`Parity-Games` $\leq_p$ `Mean-Payoff-Games` $\leq_p$
`Simple-Stochastic-Games` $\leq_p$
`Approximate-Shapley's-Stochastic-Games`

c. (Kolmogorov, 1947) Given a multi-type Branching Process, and $\epsilon > 0$, approximate its *extinction probability* within distance ϵ.

Question:

What do these three problems have to do with each other?

Question:

What do these three problems have to do with each other?

Hint

They are all fixed point problems for algebraically defined functions. Respectively:

a. Brouwer
b. Banach
c. Tarski

Question:

What do these three problems have to do with each other?

Hint

They are all fixed point problems for algebraically defined functions. Respectively:

a. Brouwer
b. Banach
c. Tarski

But are they related in terms of computational complexity?

Question:

What do these three problems have to do with each other?

Hint

They are all fixed point problems for algebraically defined functions. Respectively:

a. Brouwer
b. Banach
c. Tarski

But are they related in terms of computational complexity?

Yes!

Outline of lectures

- Background: Games, Nash Equilibria, Brouwer Fixed Points.
- Weak vs. Strong approximation of Fixed Points.
- Scarf's classic algorithm, and its complexity implications.
- The complexity class PPAD, and weak approximation.
- PPAD-completeness results for ϵ-Nash, and 2-player Nash.
- Hardness of strong approximation: square-root-sum & arithmetic circuits.
- A new complexity class: **FIXP**. Nash is FIXP-complete.
- linear-FIXP $=$ PPAD.
- Other FIXP problems:
  price equilibria, *stochastic games*, *branching processes*...
- Conclusions and future challenges.

Finite Games

A finite (normal form) game, Γ, consists of:

- A set $N = \{1, \ldots, n\}$ of players.
- Each player $i \in N$ has a finite set $S_i = \{1, \ldots, m_i\}$ of (pure) strategies. Let $S = S_1 \times S_2 \times \ldots \times S_n$.
- Each player $i \in N$, has a payoff (utility) function:

$$u_i : S \mapsto \mathbb{Q}$$

- A mixed strategy, $x_i = (x_{i,1}, \ldots, x_{i,m_i})$, for player i is a probability distribution over S_i.
  A profile of mixed strategies: $x = (x_1, \ldots, x_n)$
  Let X denote the set of all profiles.
- The expected payoff for player i:

$$U_i(x) = \sum_{s=(s_1,\ldots,s_n)\in S} \left(\prod_{k=1}^{n} x_{k,s_k}\right) u_i(s)$$

- Let x_{-i} denote everybody's strategy in x except player i's.
  Let $(x_{-i}; y_i)$ denote the new profile: $(x_1, \ldots, x_{i-1}, y_i, x_{i+1}, \ldots, x_n)$.

Nash Equilibria

A mixed strategy profile x is called:

- a Nash Equilibrium if:
  $\forall\ i$, and all mixed strategies y_i: $U_i(x) \geq U_i(x_{-i}; y_i)$

  In other words: *No player can increase its own payoff by unilaterally switching its strategy.*

- a ϵ-Nash Equilibrium, for $\epsilon > 0$, if:
  $\forall\ i$, and all mixed strategies y_i: $U_i(x) \geq U_i(x_{-i}; y_i) - \epsilon$

  In other words: *No player can increase its own payoff by more than* ϵ *by unilaterally switching its strategy.*

Nash Equilibria

A mixed strategy profile x is called:

- a Nash Equilibrium if:
  $\forall\ i$, and all mixed strategies y_i: $U_i(x) \geq U_i(x_{-i}; y_i)$

  In other words: *No player can increase its own payoff by unilaterally switching its strategy.*

- a ϵ-Nash Equilibrium, for $\epsilon > 0$, if:
  $\forall\ i$, and all mixed strategies y_i: $U_i(x) \geq U_i(x_{-i}; y_i) - \epsilon$

  In other words: *No player can increase its own payoff by more than ϵ by unilaterally switching its strategy.*

Theorem (Nash, 1950)

Every finite game has a Nash Equilibrium.

Nash's proof

Brouwer's fixed point theorem

Every continuous function $F : D \mapsto D$ from a compact convex set $D \subseteq \mathbb{R}^m$ to itself has a fixed point: $x^* \in D$, such that $F(x^*) = x^*$.

- The NEs of a finite game, Γ, are precisely the fixed points of the following Brouwer function $F_\Gamma : X \mapsto X$:

$$F_\Gamma(x)_{(i,j)} = \frac{x_{i,j} + \max\{0, g_{i,j}(x)\}}{1 + \sum_{k=1}^{m_i} \max\{0, g_{i,k}(x)\}}$$

where $g_{i,j}(x) \doteq U_i(x_{-i}; j) - U_i(x)$.

Note: $g_{i,j}(x)$ are polynomials in the variables in x, and they measure:

"how much better off would player i be if it switched to pure strategy j?"

A basic computational question

Question

What is the complexity of the following search problem:

> **("Strong") ϵ-approximation of a Nash Equilibrium:**
> *Given a finite (normal form) game, Γ, with 3 or more players, and given $\epsilon > 0$, compute a rational vector x' such that there is some (exact!) Nash Equilibrium x^* of Γ so that:*
>
> $$\|x^* - x'\|_\infty < \epsilon$$

Note:

This is NOT the same thing as asking for an ϵ-Nash Equilibrium.

Weak vs. Strong approximation of Fixed Points

- 2-player finite games always have rational NEs, and there are algorithms for computing an exact rational NE in a 2-player game (Lemke-Howson'64).
- For games with ≥ 3 players, all NEs can be irrational (Nash,1951). So we can't hope to compute one "*exactly*".

Two different notions of ϵ-approximation of fixed points:

- (**Weak**) Given $F : \Delta_n \mapsto \Delta_n$, compute x' such that:

$$\|F(x') - x'\| < \epsilon$$

- (**Strong**) Given $F : \Delta_n \mapsto \Delta_n$, compute x' s.t. there exists x^* where $F(x^*) = x^*$ and:

$$\|x^* - x'\| < \epsilon$$

Scarf's classic algorithm

Scarf (1967) gave a beautiful algorithm (refined by Kuhn and others) for computing (weak!) ϵ-fixed points of a given Brouwer function $F : \Delta_n \mapsto \Delta_n$:

1. Subdivide the simplex Δ_n into "small" subsimplices of diameter $\delta > 0$ (δ depending on ϵ and on the "modulus of continuity" of F).
2. Color every *vertex*, $\mathbf{z}$, of every subsimplex with a color i such that $\mathbf{z}_i > 0$ & $F(\mathbf{z})_i \leq \mathbf{z}_i$.
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 $\forall x, y \in \Delta_n$, $||x - y||_\infty < \delta \;\Rightarrow\; ||F(x) - F(y)||_\infty < \epsilon/2n$, then all points in a panchromatic subsimplex are weak ϵ-fixed point.
5. They need NOT in general be anywhere near an actual fixed point.

Sperner's Lemma

$\Delta_n = \{x \in [0,1]^n \mid \sum_{i=1}^n x_i = 1\}$.
If V is the set of vertices of a simplicial subdivision of Δ_n, we call a function $f : V \mapsto \{1, \ldots, n\}$ a legal coloring if $\forall x \in V$, $f(x) \in \{i \mid x_i > 0\}$.

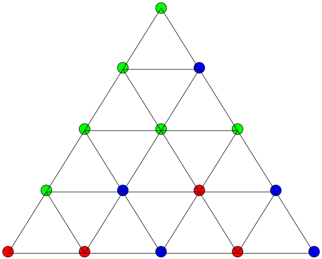

Sperner's Lemma (1928)

If vertices of a simplicial subdivision of Δ_n are legally colored, there must be at least one panchromatic subsimplex (in fact, an odd number).

"Proof" of Sperner's lemma

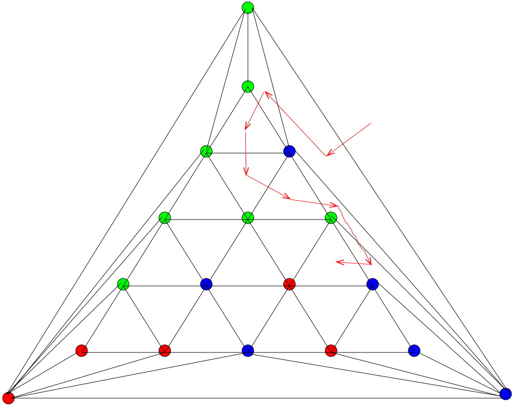

(The same combinatorial argument was also used by (Lemke-Howson'64) for an algorithm for computing a 2-player Nash Equilibrium.)

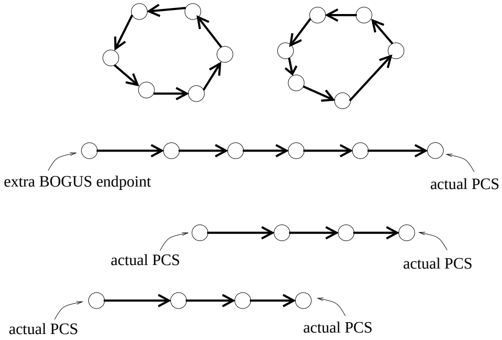

Implicit assumptions: when is Scarf's algorithm applicable?

To use Scarf's algorithm (in a reasonably efficient way) we are making several assumptions. Suppose $F : \Delta_n \mapsto \Delta_n$ is given in a (unspecified) form that requires m bits to describe.

1. $F(x)$ should be polynomial-time computable for given rational vector x. I.e., the time to compute $F(x)$ should be polynomial in both m and the encoding size of x.
2. We should have a "tractable" simplicial subdivision of Δ_n: the subsimplices and their vertices must have polynomial encoding size (in m and $size(\epsilon)$), and must yield a P-time algorithm (in m and $size(\epsilon)$) for starting at the extra bogus endpoint, and for traversing "on the fly" a single directed edge of the (implicit) line graph whose nodes are subsimplices.
3. Finally, $F(x)$ should be polynomially continuous, meaning there is a polynomial $q(r)$ such that for $\epsilon > 0$, if $\delta = 1/2^{q(m+size(\epsilon))}$, then $\forall x, y \in \Delta_n$, $\|x - y\| < \delta \Rightarrow \|F(x) - F(y)\| < \epsilon$.

Assumptions (1. – 3.) do not guarantee that Scarf's algorithm will run in P-time.
They just guarantee that each step (each edge traversal) is P-time, and we will eventually halt at a panchromatic subsimplex such that every point inside that subsimplex is a weak ϵ-fixed point of F.
(But it can potentially take exponentially many traversal steps in the encoding size m and in $size(\epsilon)$, because there can be exponentially many subsimplices. Indeed, such worst-case exponential examples exist.)

ϵ-NEs are weak ϵ-fixed points

Proposition

For finite games, Γ, computing an ϵ-NE is P-time equivalent to computing a weak ϵ-fixed point of Nash's function F_Γ.

Thus, to compute an ϵ-NE, we can simply apply Scarf's algorithm to F_Γ. The functions F_Γ satisfy all the implicit assumptions (1.–3.) for applicability of Scarf's algorithm. (The compact convex domain X has tractable simplicial subdivisions too.)

Question

What does all this tell us about the complexity of computing an ϵ-NE?

The complexity class PPAD

Papadimitriou (1992) defined **PPAD**, based on the "directed line" parity argument, to capture (approximate) Nash and Brouwer, etc...

Definition

PPAD is the class of search problems polynomial-time reducible to:
Directed line endpoint problem: Given two boolean circuits, S ("Successor") and P ("Predecessor"), each with n input bits and n output bits, such that $P(0^n) = 0^n$, and $S(0^n) \neq 0^n$, find a n-bit vector, $\mathbf{z}$, such that either: $P(S(\mathbf{z})) \neq \mathbf{z}$ or $S(P(\mathbf{z})) \neq \mathbf{z} \neq 0^n$.
(By the directed line parity argument such a $\mathbf{z}$ exists (for inconsistent P and S it exists trivially).)

PPAD lies somewhere between (the search problem versions of) P and NP.

By Scarf's algorithm, computing a ϵ-NE is in PPAD.
Can we do better??

No. Computing (ϵ-)NEs is hard for PPAD:

Theorem

1. *[Daskalakis-Goldberg-Papadimitriou'06][Chen-Deng'06]:*
   Computing a ϵ-NE for a 3 player game is PPAD-complete.
2. *[Chen-Deng'06]:*
   Computing an exact (rational) NE for a 2 player game is PPAD-complete.

But what if we want to approximate exact NEs for games with ≥ 3 players and to approximate exact fixed points?
I.e., what if we want to do strong approximation of fixed points?
(**Warning:** Scarf's algorithm does not in general yield strong ϵ-fixed points.

Why care about strong approximation of fixed points?

- It can be argued (Scarf (1973) implicitly did) that for many applications in economics weak ϵ-fixed points of Brouwer functions are sufficient.
- However, many important problems boil down to a fixed point computation for which weak ϵ-FPs are useless, unless they also happen to be strong ϵ-FPs.

  Examples:

  –Shapley's Stochastic Games;

  –Condon's (1992) Simple Stochastic Games;

  –Kolmogorov's multi-type Branching Processes;

  (and Recursive Markov Chains, and Recursive Stochastic Games, ...)

Proposition

Given game Γ and $\epsilon > 0$, we can Strong ϵ-approximate a NE in **PSPACE**.

Proof.

For Nash's functions, F_Γ, the expression

$$\exists \mathbf{x}(\mathbf{x} = F_\Gamma(\mathbf{x}) \wedge \mathbf{a} \leq \mathbf{x} \leq \mathbf{b})$$

can be expressed as a formula in the Existential Theory of Reals (ETR).
So we can Strong ϵ-approximate an NE, $x^* \in \Delta_n$, in **PSPACE**, using $\log(1/\epsilon)n$ queries to a PSPACE decision procedure for ETR ([Canny'89],[Renegar'92]).
(These are deep, but thusfar impractical algorithms.) □

Can we do better than **PSPACE**?

two "hard" problems

Sqrt-Sum

The square-root sum problem is the following decision problem:

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 solvable in PSPACE.

Open problem ([GareyGrahamJohnson'76]) whether it is solvable even in NP (or even the polynomial time hierarchy).

PosSLP

Given an arithmetic circuit (Straight Line Program) over basis $\{+, *, -\}$ with integer inputs, decide whether the output is > 0.

[Allender et. al.'06] Gave a (Turing) reduction from Sqrt-Sum to PosSLP and showed both can be decided in the Counting Hierarchy:

$$P^{PP^{PP^{PP}}}$$

why isn't PosSLP easy??

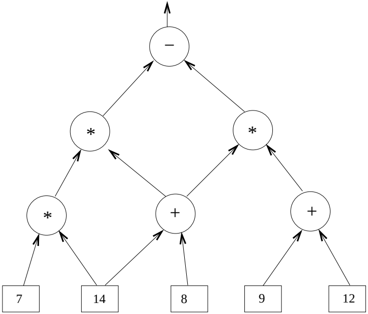

Remarks

- Deciding **Sqrt-Sum** is critical to exact geometric computations in Euclidean space.
  In particular, whether exact Euclidean-TSP is in NP hinges on Sqrt-Sum.
- Every discrete decision problem in P-time in the unit-cost arithmetic RAM model, i.e., the (discrete, rational) Blum-Shub-Smale class $\mathbf{P}_{\mathbb{R}}$, is P-time (Turing) reducible to **PosSLP**.
  So, **PosSLP** captures discrete problems in $\mathbf{P}_{\mathbb{R}}$.
- Testing whether a sum of square roots is equal to k is decidable in P-time (e.g., (Borodin-Fagin-Hopcroft-Tompa, 1985).
- Testing $= 0$ for $\{+, *, -\}$ arithmetic circuits (much easier than PosSLP) is known to be in **coRP** (Schönhage,1979).
  Whether it is in P-time is already a well-known open problem: it is equivalent to polynomial identity testing.

Sqrt-Sum & PosSLP $\leq_p$ approximation of actual NE

Theorem

Any non-trivial approximation of an actual NE is both **Sqrt-Sum***-hard and* **PosSLP***-hard.*

More precisely: for every $\epsilon > 0$*, both* **Sqrt-Sum** *and* **PosSLP** *are P-time reducible to the following problem:*

Given a 3-player (normal form) game, Γ*, with the property that:*

1. *in every NE, player 1 plays exactly the same mixed strategy,* x_1^**, and*
2. *the probability,* $x_{1,1}^*$*, with which player 1 plays it first pure strategy is either:*

$$(a.) \quad = 0 \quad , \qquad\qquad or \qquad\qquad (b.) \quad \geq (1 - \epsilon)$$

Decide which of (a.) or (b.) is the case.

Question

How far can an ϵ-NE be from an actual NE?

Answer: Very far!

First, a seemingly contrary fact:

Fact

For every continuous function $F : \Delta \mapsto \Delta$, and every $\epsilon > 0$, there exists a $\delta > 0$, such that a weak δ-fixed point of F is a strong ϵ-fixed point of F.

But this is non-constructive! It uses a compactness argument. (Bolzano-Weierstrass.)
(Indeed, compactness and the Sperner Lemma argument together easily yield a proof of Brouwer's fixed point theorem.)

From a constructive, computational perspective, this is certainly NOT the full story.

ϵ-NEs can be very far from actual NEs

Theorem

- *For every n, there exists a 4-player game Γ_n of size $O(n)$ with an ϵ-NE, x', where $\epsilon = \frac{1}{2^{2^{\Omega(n)}}}$, and yet x' has distance 1 in l_∞ to any actual NE. (Thus worst possible distance in l_∞.)*
- *The same holds for 3 players, but with distance 1 replaced by distance $(1 - \delta)$, for any fixed constant $\delta > 0$ (and even for $\delta = 2^{-poly(n)}$).*

Question

Is that the smallest ϵ (in terms of the game size n) for which an ϵ-NE has "large" distance to any actual NE?

Conjecture

Essentially yes. Meaning for large enough n, you can't have, say, a ϵ-NE where $\epsilon = \frac{1}{2^{2^{n^{\omega(1)}}}}$, without it being close (say within distance $1/poly(n)$), of an actual NE.

A new complexity class: FIXP

Consider the following class of fixed point problems:

FIXP

- **Input:** algebraic circuit (straight-line program) over basis $\{+, *, -, /, \max, \min\}$ with rational constants, having n input variables and n outputs, such that the circuit represents a continuous function $F : [0,1]^n \mapsto [0,1]^n$.
  (The domain $[0,1]^n$ can be allowed to be much more general.
  It can any convex polytope defined by given linear inequalities, or it can even be an ellipsoid domain. See our paper.)
- **Output:** Compute (or strong ϵ-approximate) a fixed point of F.

We close these problems under suitable P-time reductions.
Call the resulting class **FIXP**.

We shall see that many interesting problems besides Nash are in **FIXP**.

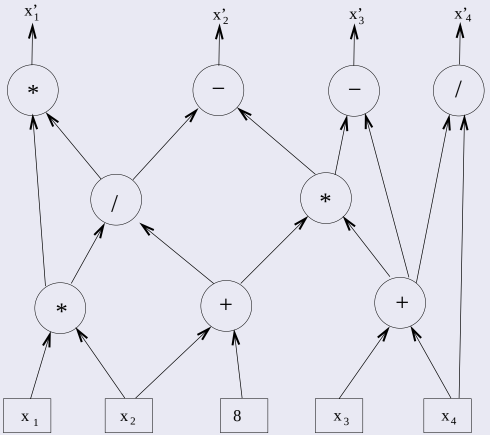
x'1
x'2
x'3
x'4
*
−
−
/
/
*
*
+
+
x 1
x2
8
x3
x4

Nash is FIXP-complete

Theorem

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

It is complete in several senses:

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

Completeness holds under very restrictive P-time (real valued) search problem reductions where the "solution recovery" function g is linear.

Note that containment in FIXP follows from Nash's functions F_Γ.

Very brief sketch of some proof ideas

- 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 payoff 1 if it "guesses correctly" whether player 1 plays pure strategy A or not, and payoff 0 otherwise.
  In any NE, the new player will play one of its two strategies with probability 1.
  Deciding which of the two solves Sqrt-Sum.
- What about equality? We don't have to worry about it because $\sum_i \sqrt{d_i} = k$ is P-time decidable ([BFHT'85]).

A key ingredient in our proofs

Two beautiful gems by Bubelis:

Theorem (Bubelis, 1979)

1. *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.*
2. *There is a general polynomial-time reduction from n-player games to 3-player games.*
   Such that you can easily recover a (real valued) NE of the n-player game as a linear function of a given NE in the resulting 3-player game.

Many details in the proof of FIXP-completeness:

- A series of transformations to get circuits into a "normal form" with additional "conditional assignment gates".
- Transform circuit to a game with a large (but bounded) number of players, using suitable *gadgets*.
  Key gadgets can be derived from (Bubelis'79)'s constructions.
  (Alternatively, the gadgets of (Golberg-Papadimitriou'06), (Daskalakis-Golberg-Papadimitriou'06) can also be used.)
- Reduce to 3-players: again uses (Bubelis '79).

Another FIXP-complete problem: Price Equilibria

Price Equilibria in Exchange Economies

- An idealized exchange economy with n agents and m commodities.
- Each agent j starts off with an initial endowment of commodities $w_j = (w_{j,1}, \ldots, w_{j,m})$.
- For a given price vector, $p \geq 0$, each agent j has an demand function $d_i^j(p)$ for commodity i.
  It will choose its demands to maximize its utility using the budget obtained by selling all its endowment w_j at the price vector p.
  Under certain conditions (e.g., continuity and strict quasi-concavity of utilitity functions) demands are uniquely determined continuous functions of the utilities of the agents.

Excess demands, Walras's law, etc.

- From the demand functions we directly get excess demand functions: $g_i^j(p) = d_i^j(p) - w_{j,i}$, for agent j and commodity i.
- The total excess demand for commodity i is $g_i(p) = \sum_j g_i^j(p)$.
- Excess demands are continuous and satisfy economically justified axioms:
  - (Homogeneous of degree 0): For all $\alpha > 0$, $p \geq 0$, $g_i^j(\alpha p) = g_i^j(p)$. (So, we can w.l.o.g. consider only "normalized" price vectors in Δ_m.)
  - (Walras's law): $\sum_i p_i g_i(p) = 0$.

Excess demand functions can be quite arbitrary continuous functions (Sonnenschein-Mantel-Debreu,1973-74).

Price Equilibrium

A vector of prices $p^* \geq 0$ such that $g_i(p^*) \leq 0$ for all i ($= 0$ if $p_i^* > 0$).

Theorem ((Arrow-Debreu'54) proved a much more general fact)

Every exchange economy has a price equilibrium.

The proof is via Brouwer's fixed point theorem. (And for more general market equilibrium results (including with production, etc.), it is via the closely related Kakutani fixed point theorem.)

Proposition

Computing price equilibria in exchange economies where excess demands are given by algebraic circuits over $\{+, *, -, /, \max, \min\}$ is FIXP-complete.

Proof.

One direction of proof is via the following variant of Nash's function:

$$H(p)_i = \frac{p_i + \max\{0, g_i(p)\}}{1 + \sum_{j=1}^{m} \max\{0, g_j(p)\}}$$

where $g_i(x)$ is the total excess demand for commodity i.
The (Brouwer) fixed points of $H(p)$ are the price equilibria of the economy.
The other direction follows from Uzawa (1962):
For a Brouwer function $F : \Delta_n \mapsto \Delta_n$, define total excess demand function $g : \Delta_n \mapsto \mathbb{R}^n$ by

$$g(p) = F(p) - (\frac{\langle p, F(p) \rangle}{\langle p, p \rangle})p$$

A new characterization of PPAD

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

Theorem

The following are all equivalent:

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

Corollary

`Simple-Stochastic-Games` *(and* `Parity Games`*, etc.) are in PPAD.*

sketch proof that PPAD $\leq$ linear-FIXP

Computing a 2-player NE (exactly) is PPAD-complete (Chen-Deng'06).
So we only need to give a reduction from two player NE to linear-FIXP.
Nash's functions F_Γ are already non-linear even for 2 players.
Is there a different, $\{+, \max\}$ function for 2-player NEs??
Yes!

(Gul-Pearce-Staccetti, 1993) describe a fixed point approach for NEs.
By examining carefully what they do, one can derive the follow function:

1. let $x'_{i,j} := x_{i,j} + U_i(x_{-i};j)$.

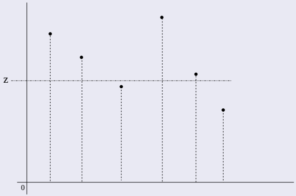

2. "*project*" the vector x'_i onto the simplex Δ_{m_i}, for every player i.

Fact

The fixed points of this function are the NEs.

Can "projection" be computed with a linear-FIXP function?
Yes, ... with the help of *sorting networks*.

From this revised function for n-player NEs we also obtain:

Theorem

The basis $\{+, *, \max\}$ *is sufficient to capture all of* **FIXP**.

Simple Stochastic Games

Simple Stochastic Games (SSGs) (Condon,1992) are 2-player games on directed graphs:

- some nodes are *random* (V_{rand}), some belong to Player 1 (V_1), some to Player 2 (V_2). There is a designated *goal* node, t.
- Starting at a vertex, players choose edges out of nodes belonging to them. Edges out of random nodes are chosen randomly according to a probability distribution.
- Player 1 wants to maximize the probability of reaching t. Player 2 wants to minimize it.

Deciding whether the *value* of these (zero-sum) games is $\geq 1/2$ is in **NP**$\cap$**coNP** (Condon'92).

Note: At least as hard as Parity Games, and Mean Payoff Games.

SSGs are in PPAD

Fixed point equations for x_u, the *value* of these games starting at vertex u:
$x_t = 1$
$x_u = \sum_v p_{u,v} x_v$, for $u \in V_{rand}$
$x_u = \max\{x_v \mid (u,v) \in E\}$, for $u \in V_1$
$x_v = \min\{x_v \mid (u,v) \in E\}$, for $u \in V_2$
These are piecewise-linear, but can have multiple fixed points. But it is possible to "preprocess" them so that they have a unique fixed point which gives the value of the game starting at each vertex.

Theorem

Simple stochastic game are in linear-FIXP, and thus in PPAD.

Note: Weak ϵ-fixed points are useless here and easy to compute (exercise).

(Juba, MSc thesis, 2005) observed SSGs $\in$ PPAD, but his proof had a gap related to weak vs. strong approx. and misinterpretating (Papadimitriou, 1992).

Shapley's Stochastic Games (Shapley, 1953)

2-player, zero-sum, imperfect information, discounted stochastic games.

1. finite state space, finite move alphabet.
2. Starting in a given state, at each round both players (independently), choose a move, or a probability distribution on moves. Their joint move determines a probability distribution on the next state, and a reward to player 1.
3. The rewards after each round are discounted by given factor $0 < \beta < 1$, and the total discounted reward to player 1 is sum $\sum_i \beta^i r_i$.

The *value* of Shapley's games (which can be irrational) can be characterized by fixed point equations, $\mathbf{x} = P(\mathbf{x})$, where $P(\mathbf{x})$ is a *contraction map*.

There is a unique *Banach fixed point* (which can be irrational), which yields the game value starting at each state.

Theorem

For Shapley's stochastic games:

1. *Computing the game value is in FIXP.*
2. *The (strong) approximation problem for the game value is in PPAD.*
3. *The decision problem (is the game value $\geq r$?) is SqrtSum-hard.*

Proof.

Sketch Proof of part (2.): $P(\mathbf{x})$ is a "fast enough" contraction mapping. For such mappings, Weak ϵ-fixed points are "close enough" to the actual Banach fixed point. $P(\mathbf{x})$ is a Brouwer function on a "not too big" domain. Thus: apply Scarf's algorithm to $P(\mathbf{x})$. □

Note: this also implies Condon's Simple Stochasic Games are in PPAD.

multi-type Branching Processes

Branching processes, studied in the 19th century by Galton and Watson.
multi-type Branching Processes (mt-BPs) defined by Kolmogorov and studied by him and Sevastyanov ('47-'51) and others.
mt-PBs have a huge literature in probability, population genetics,...

1. A population of *individuals*. Each individual has one of a fix set of *types*.
2. In each generation, each individual of each type "gives birth" to a number of individuals (a multi-set) of different types, according to a probability distribution on multi-sets, determined by its type.

Question

Starting from one entity of a given type, will the population eventually go extinct with probability $\geq 1/2$?

(Whether it will *almost surely* go extinct is decidable in P-time (E.-Y.'05).)

The extinction problem for mt-BPs is in FIXP

- The extinction probabilities are the *Least Fixed Point* (LFP) solution of a *monotone* system of nonlinear polynomial equations, $\mathbf{x} = P(\mathbf{x})$. (The LFP exists, by Tarski's (Tarski-Knaster) fixed point theorem.)
- The LFP can be irrational, and the associated decision problems are SqrtSum-hard and PosSLP-hard ([EY05,EY07]).

Theorem

The mt-BP extinction problem is in FIXP.

Proof.

The LFP can be "*isolated*" as the unique fixed point of FIXP function. □

Note: mt-BP extinction ≡ 1-exit *Recursive Markov Chain* termination ≡ Stochastic-Context-Free Grammar termination.

Conclusions, and many, many open questions

- Can strong approximation of NEs be done in better than **PSPACE**? Is it hard for a standard complexity class like **NP**? (NP-hardness would imply the "*rational*" BSS class $\mathbf{NP}_{\mathbb{R}}$ contains both NP and coNP. That's an open problem.)
- Can we obtain any better upper bound for the Sqrt-Sum and PosSLP problems than the Counting Hierarchy?
- Basic practical question: Is there an algorithm that given a game & $\epsilon > 0$:
  1. is guarranteed to output a point x within distance ϵ of some actual NE, and
  2. performs "reasonably well" in practice?

K. Etessami and M. Yannakakis, "On the complexity of Nash Equilibria and Other Fixed Points", FOCS'07.

(See full version of paper at: `http://homepages.inf.ed.ac.uk/kousha`)