On the complexity of Nash Equilibria and other Fixed Points

Kousha Etessami
U. of Edinburgh

Mihalis Yannakakis
Columbia U.

(based on a paper at FOCS'07)

Dagtuhl Seminar on Equilibrium Computation
November 19-23, 2007

Question: *What is the complexity of the following search problem?*

Given a finite game, and $\epsilon > 0$, compute a vector x' that has distance less than ϵ to some (exact!) Nash Equilibrium.

Let's restate this search problem more precisely:

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

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

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

Finite Games

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

1. A set $N = \{1, \ldots, n\}$ of players.

2. Each player $i \in N$ has a finite set $S_i = \{1, \ldots, m_i\}$ of (pure) *strategies*.
   Let $S = \Pi_{i=1}^{n} S_i$.

3. Each player $i \in N$, has a *payoff (utility) function* $u_i : S \mapsto \mathbb{Q}$.

notation: mixed strategies, expected payoffs, etc.

- A *mixed* (i.e., *randomized*) *strategy*, x_i, for player i is a probability distribution over its pure strategies S_i, i.e., a vector $x_i = (x_{i,1}, \ldots, x_{i,m_i})$, such that $x_{i,j} \geq 0$, and $\sum_{j=1}^{m_i} x_{i,j} = 1$.

  Let X_i denote the set of mixed strategies for player i.
  Let $X = \Pi_{i=1}^{n} X_i$ denotes the set of *profiles* of mixed strategies.

- The *expected payoff* for player i under profile $x \in X$, is:
  $U_i(x) = \sum_{s=(s_1,\ldots,s_n)\in S} (\Pi_{k=1}^{n} x_{k,s_k}) u_i(s)$.

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

  $z_i \in X_i$ is a *best response* for player i to x_{-i} if for all $y_i \in X_i$,
  $$U_i(x_{-i}; z_i) \geq U_i(x_{-i}; y_i).$$

Nash Equilibria

A profile of mixed strategies, $x \in X$, is a *Nash Equilibrium* if for every player i its mixed strategy, x_i, is a best response to x_{-i}. In other words, *no player can increase its own payoff by switching its strategy unilaterally*.
(x is an *ϵ-Nash Equilibrium*, for $\epsilon > 0$, if 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 proved this using *Brouwer's fixed point theorem:* Every continuous function $F : D \mapsto D$ from a compact convex set D to itself has a fixed point.
He showed that the NEs of a finite game, Γ, are the fixed points of the function $F_\Gamma : X \mapsto X$:

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

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

Weak vs. Strong approximation of Fixed Points

Games with 2 players always have rational NEs, and there are specialized 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' such that there exists x^* where $F(x^*) = x^*$ and $\|x^* - x'\| < \epsilon$.

Scarf's classic algorithm

Scarf (1967) gave a beautiful algorithm 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 $\delta > 0$ (depending on the "modulus of continuity" of F, and on $\epsilon > 0$).

2. Color every *vertex*, $\mathbf{z}$, of every subsimplex with a color i such that $F(\mathbf{z})_i \leq \mathbf{z}_i$.

3. By **Sperner's Lemma** there must exist a panchromatic subsimplex. (And the proof of Sperner's lemma provides a way to "navigate" toward such a simplex.)

4. *Fact:* If $\delta > 0$ is chosen such that
$\forall x, y \in \Delta_n,\ ||x - y||_\infty < \delta \ \Rightarrow\ ||F(x) - F(y)||_\infty < \epsilon/2n$, and $\delta \leq \epsilon/2n$,
then the points in a panchromatic subsimplex are weak ϵ-fixed points.

some facts about the Weak vs. Strong distinction

Fact: *For a large class of fixed point search problems*[1]
Weak ϵ-approximation is P-time reducible to Strong ϵ-approximation

Fact: *For finite games, computing an ϵ-Nash Equilibrium 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 apply Scarf's algorithm to F_Γ.
Papadimitriou (1992) defined a complexity class, PPAD, to capture Sperner, Scarf, and computation of fixed points and NEs.
(PPAD lies between (the search problem versions of) P and NP.)

So, computing ϵ-NEs is in PPAD. In fact, it is PPAD-complete ([DasGolPap'06]), and even computing a exact NE for 2 players is PPAD-complete ([CheDen'06]).

Warning: Scarf's algorithm does not in general yield Strong ϵ-fixed points.

[1] Namely, those with *polynomially continuous* Brouwer functions. These include Nash's functions and much more.

A basic upper bound for Strong ϵ-aprroximation of Nash

Fact: *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]). ∎

Can we do better than **PSPACE**?

Why care about strong approximation of fixed points?

- It can be argued (as Scarf (1973) implicitly did) that for many applications in economics weak ϵ-fixed points of Brouwer functions are sufficient.

- However, there are many important computational problems that boil down to a fixed point computation, and for which Weak ϵ-FPs are useless, unless they also happen to be Strong ϵ-FPs.

- Our understanding of these issues is informed by our work on *Recursive Markov Chains*, *Branching Processes*, and *Stochastic Games*,....
I will come back to these later.....

The Square-Root Sum problem

The square-root sum problem (**Sqrt-Sum**) 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 known to be solvable in PSPACE.
(Recently, the upper bound was improved by Allender et. al. [ABKM'06] to the 4th level of the *Counting Hierarchy*: $P^{PP^{PP^{PP}}}$.)

But it has been a major open problem ([GareyGrahamJohnson'76]) whether it is solvable even in NP.
(In particular, whether exact Euclidean-TSP is in NP hinges on this.)

Sqrt-Sum and approximation of Nash Equilibria

Theorem: For every $\epsilon > 0$, **Sqrt-Sum** is 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, and
2. in every NE, player 1 plays its first pure strategy either with probability 0 or with probability $\geq (1 - \epsilon)$,

decide which of the two is the case (i.e., 0 or at least $(1 - \epsilon)$?).

Thus, if we can do any non-trivial approximation of an actual NE, even in NP, then **Sqrt-Sum** is in NP, and exact Euclidean-TSP is in NP, etc., etc., ...

Brief ideas of proof

- Suppose we could create a (3-player) game such that, Player 1 plays strategy 1 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 right" whether player 1 played strategy 1 or not, and low payoff otherwise.

  In any NE, the new player will play one of its two strategies with probability 1. Deciding which will solve 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]).

proof ingredients, continued...

- **Theorem** (Bubelis, 1979) *Every real algebraic number can be "encoded" as the payoff to player 1 in a unique NE of a 3-player game.*

  More precisely, given any polynomial $f(z)$ with rational coefficients, and given rationals $a < b$ such that $f(a) < 0 < f(b)$, we can efficiently construct a 3-player game of size polynomial in the size of f, a, and b, such that Player 1 gets payoff α in some NE iff $f(\alpha) = 0$ and $a < \alpha < b$. Moreover, if α is the unique root between a and b, then there is a unique (fully mixed) NE.

- Several issues to resolve:
  (a) we need to transfer this to the probability of strategies, not payoffs.
  (b) $\sum_{i=1}^{n} \sqrt{d_i}$ has high algebraic degree ($\sim 2^n$). We can instead express each $\sqrt{d_i}$ as a "subgame", and use *matching pennies* and a direct *summing* construction.
  (c) Somehow we still have to make it back down to 3 players at the end.....

A harder arithmetic circuit decision problem

Allender et. al. [ABKM'06] Showed that **Sqrt-Sum** reduces to the following more general problem (which they showed lies in the *Counting Hierarchy*):

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

In fact, every *discrete* decision problem solvable in the Blum-Shub-Smale class $\mathbf{P}_{\mathbb{R}}$ is P-time (Turing) reducible to **PosSLP**. So, **PosSLP** captures discrete decision problems in $\mathbf{P}_{\mathbb{R}}$.

Theorem:
PosSLP is P-time reducible to Strong ϵ-approximation of 3-player NEs.

(More precisely, it reduces to the same 0 vs. $(1-\epsilon)$ choice problem as before.)

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

Answer: Very far!

Seemingly contrary to this suggestion, is the following basic 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 a non-constructive fact. From a quantitative, computational perspective, that is certainly NOT the full story:

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_∞) from any actual NE. (Same holds for 3 players, but with distance 1 replaced by distance $(1 - 2^{-poly})$.)

Question: Is that the smallest ϵ (in terms of the game size n) for which an ϵ-NE has distance (close to) 1 to actual NEs?

Conjecture: Essentially yes.

A new complexity class: FIXP

Consider the following class of fixed point problems:

We are given a continuous function $F : [0,1]^n \mapsto [0,1]^n$, presented as an algebraic circuit over the basis $\{+, *, -, /, \max, \min\}$, with rational constants, and we wish to compute (or Strong ϵ-approximate) a fixed point of F.

Let us close these problems under P-time reductions, and call the resulting class of fixed point search problems **FIXP**.

As we shall see, many interesting problems besides Nash fall into the class FIXP.

Nash is FIXP-complete

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

It is complete in several senses: "exact" (real valued) computation, strong ϵ-approximation, and an appropriate "decision" version of the problem.

Very brief outline of proof:

- A series of transformations to get circuits into a "normal form" with additional "conditional assignment gates".
- Transform circuit to a game with a bounded number of players using suitable *gadgets*. Some gadgets are from [GolPap06],[DasGolPap06] and some are new.
- Reduce to 3-players using/adapting another beautiful construction by Bubelis (1979): a P-time reduction from arbitrary games to 3-player games.

Another FIXP-complete problem: Price equilibria

- An idealized *exchange economy* with n agents and m commodities.
- For a given price vector, p, each agent l has an excess demand function $g_i^l(p)$ for commodity i. Excess demands satisfy certain axioms (e.g.. Walras's law).
- The total excess demand for commodity i is $g_i(p) = \sum_l g_i^l(p)$.
- *Price Equilibrium:* prices, p^* such that $g_i(p^*) \leq 0$ for all i (and $= 0$ if $p_i^* > 0$).
- **Fact** Every exchange economy has a price equilibrium. (Proof via Brouwer.)
- **Proposition** Computing Price Equilibria in exchange economies where excess demands are given by algebraic circuits over $\{+, *, -, /, \max, \min\}$ is FIXP-complete. (Follows from Uzawa (1962).)

So, what is 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"

(These always have rational fixed points of polynomial bit complexity.)

In fact, from the proofs it also follows that the smaller basis $\{+, *, \max\}$ and rational constants, suffices to capture **FIXP**.

proof that PPAD $\leq$ linear-FIXP

Computing a 2-player NE (exactly) is PPAD-complete, so we only need to give a reduction from two player NE to linear-FIXP.

Nash's functions F_Γ are non-linear even for 2 players.

There is a different fixed point function for NEs ([GPS'93]):

First, let $x'_{i,j} := x_{i,j} + U_i(x_{-i}; j)$.
Second, *"project"* the vector x'_i onto the simplex Δ_{m_i}, for every player i.

The fixed points of this function are the NEs.

Can we compute the "projection" with a linear-FIXP function?

Yes, ... and here *sorting networks* come into the picture.

Simple Stochastic Games

Simple Stochastic Games (SSGs) [Condon'92] 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**.

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, and so that the fixed point is the value of the game. Thus:
Theorem: Computing SSGs game values is in linear-FIXP and thus in PPAD.

[Juba, Blum, Williams, 2005, CMU MSc. thesis] already observed that the SSGs problem is in PPAD. (But their proof has an gap, related to not noting the distinction between weak and strong approximation.)

Shapley reduces to Nash

Shapley (1953) originally defined a richer class of stochastic games. SSGs are P-time reducible to Shapley's games.

Shapley's games have non-linear fixed point equations $\mathbf{x} = P(\mathbf{x})$ with a unique (in fact a *Banach*) fixed point (which can be irrational). They are easily in FIXP.

Theorem *Deciding whether the value of Shapley's stochastic games is* $\geq 1/2$ *is Sqrt-Sum-hard.* On the other hand....

Theorem ϵ*-approximation of the value of Shapley's games is in PPAD.*

Proof: $P(\mathbf{x})$ 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 $P(\mathbf{x})$. ∎

Another problem in FIXP: Branching processes

Branching processes, were originally studied in the 19th century by Galton and Watson.

Kolmogorov (1947) defined and studied *Multi-Type Branching Processes*(MT-BPs) with Sevastyanov and others. They have a huge literature in probability theory, population genetics, and many other areas.

1. A population of *individuals*. Each individual has one of a fix set of *types*.

2. In each generation, every individual of a given type "gives birth" to a number of individuals of different types, according to a probability distribution based on its type.

Question: Will the population go extinct with probability $\geq 1/2$?

This is a non-linear fixed point problem. It is SqrtSum-hard. In general, there are multiple fixed points, but the *least fixed point* (LFP) gives the extinction probabilities we are interested in (they can be irrational).

With some hard work, we can "isolate" the LFP as the unique fixed point of a Brouwer function. Thus:

Theorem: *The Multi-Type Branching Process extinction problem is in FIXP.*

The MT-BP extinction problem is equivalent to the 1-exit *Recursive Markov Chain* (RMC) termination problem.

Theorem *Any non-trivial approximation of the general (multi-exit) RMC termination problem is both SqrtSum-hard and PosSLP-hard.*

Concluding remarks

Our results raise many new questions:

- Can Strong approximation of NEs be done in anything better than **PSPACE**?

- Is Strong approximation of NEs hard for a standard complexity class like **NP**?

  There is some reason to suspect this will not be easy to show. When fixed points/equilibria are unique, these problems can be placed in the "*rational fragment of*" the Blum-Shub-Smale class $\mathbf{NP}_{\mathbb{R}} \cap \mathbf{coNP}_{\mathbb{R}}$, and nothing in that class is known to be NP-hard.

- <u>A basic practical question:</u> Is there any algorithm that, given a game and ϵ:
  1. is guarranteed to output a point x within distance ϵ of some actual NE, and
  2. performs "reasonably well" in practice?