
**On the complexity of approximating
exact Fixed Points:
Nash Equilibria, Stochastic Games,
and Recursive Markov Chains**

Kousha Etessami
U. of Edinburgh

Mihalis Yannakakis
Columbia U.

Algorithmic Game Theory Workshop, Warwick
March 26, 2007

Appetizer

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.

Weak vs. Strong approximation of Fixed Points

The NEs of a finite game, Γ , are the Brouwer fixed points of $F_\Gamma : \Delta_n \mapsto \Delta_n$ (Nash, 1951).

(Recall: $F_\Gamma(\mathbf{x})_{(i,j)} \doteq \frac{x_{i,j} + \max\{0, g_{i,j}(\mathbf{x})\}}{1 + \sum_{k=1}^{m_i} \max\{0, g_{i,k}(\mathbf{x})\}}$, where $g_{i,j}(\mathbf{x})$ are polynomials in \mathbf{x} .)

For ≥ 3 players, all NEs can be irrational. So we can't 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 $\|x^* - x'\| < \epsilon$.
-

some facts about the Weak vs. Strong distinction

Fact: For a large class of Brouwer functions¹
Weak ϵ -approximation is P-time reducible to Strong ϵ -approximation

Fact: For finite games, Γ , 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, apply Scarf's algorithm (SPERNER) to F_Γ .
This yields a Weak ϵ -FP of F_Γ .

So, computing ϵ -NEs is in PPAD, and of course PPAD-complete ([DasGolPap'06]),
and even computing exact NEs for 2 players is PPAD-complete ([CheDen'06]).

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

¹namely, all "polynomially continuous" functions. These include Nash's functions, and the other explicit classes of functions we will discuss.

Scarf and Nash

Scarf's algorithm treats $F(\mathbf{x})$ as a black-box, only evaluating it at various points. For such "oracle" algorithms, it is known that no number of "adaptive" queries suffice to Strong ϵ -approximate some FP.

Of course, Nash's functions are not black-box oracles.

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 (1967) implicitly did) that for many applications in economics and elsewhere 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* and *Stochastic Games*,....
so I will make a (brief) detour...
-

And now for something completely different:
What is a Recursive Graph?




Question: Is it possible to reach **b** from **a**?

Such information can easily be computed in P-time.

Recursive Graphs are abstract models of procedural programs with recursion. They are expressively equivalent to Pushdown Systems, and there has been very extensive work on their algorithmic analysis in verification research.


What is a Recursive Markov Chain?



Question: What is the probability of eventually reaching **b** from **a**?
Is there an efficient algorithm for computing such probabilities?

- The special case of 1-exit RMCs (1-RMCs) already captures some classic probabilistic models: Multi-Type Branching Processes and Stochastic Context-Free Grammars.
- A restricted subclass of 1-RMCs captures Random Walks with Back-Buttons, a model of “web surfing” studied by ([Fagin,Karlin,Kleinberg,et. al.'01]).

Let's calculate this termination probability



Let x be the (unknown) probability that starting at **a** (in the empty calling context) we will eventually reach **b** (in the empty calling context) and terminate.

An equation for x :
$$x = (1/2)x^2 + 1/4$$

Note: this is a nonlinear equation with two solutions: $x = 1 \pm \frac{1}{\sqrt{2}}$.


The *least* solution, let's call it the *Least Fixed Point* (LFP), is: $x^* = (1 - \frac{1}{\sqrt{2}})$.

Fact: This is the probability we are after.

(In particular, termination probabilities can be irrational.)

The non-linear system associated with an RMC

Let $x_{(f,u,z)}$ denote the (unknown) probability that, in “component” f , starting at u (with empty call stack) we eventually reach exit z (with empty call stack).



What is $x_{(f,z,z)}$?

$$x_{(f,z,z)} = 1$$

What is $x_{(f,a,z)}$?

$$x_{(f,a,z)} = \frac{1}{3}x_{(f,h,z)} + \frac{2}{3}x_{(f,(b1,c),z)}$$

What is $x_{(f,(b1,c),z)}$?

$$x_{(f,(b1,c),z)} = x_{(g,c,d)}x_{(f,(b1,d),z)} + x_{(g,c,e)}x_{(f,(b1,e),z)}$$

These “patterns” cover all cases, yielding a system of polynomial equations:

$$\mathbf{x} = P(\mathbf{x})$$

Basic facts about the system $\mathbf{x} = P(\mathbf{x})$

- The coefficients in $P(\mathbf{x})$ are non-negative, and $P : \mathbb{R}^n \mapsto \mathbb{R}^n$ defines a monotone operator mapping $D \subseteq [0, 1]^n$ to itself.

By a Tarski-Knaster argument, $P()$ has a Least Fixed Point, \mathbf{x}^* in $[0, 1]^n$,

- **Theorem:** The **LFP**, $\mathbf{x}^* = \lim_{k \rightarrow \infty} P^k(\mathbf{0})$, is the vector of termination probabilities.
 - Can we compute \mathbf{x}^* efficiently? Again, we can express the formula: $\exists \mathbf{x}(\mathbf{x} = P(\mathbf{x}) \wedge \mathbf{a} \leq \mathbf{x} \leq \mathbf{b})$ in ETR. Thus, deciding exact queries about \mathbf{x}^* , and Strong ϵ -approximation of it, are in **PSPACE**.
 - (We know a lot more about numerical computation of \mathbf{x}^* another talk!)
 - **Note:** Weak ϵ -FPs of $P(\mathbf{x})$ are useless.
-

RMCs and the Square-Root Sum problem

The square-root sum problem (**Sqrt-Sum**) is the following decision problem:

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

It is known to be solvable in PSPACE 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.)


Theorem: **Sqrt-Sum** is P-time reducible to the following problems:

1. Given a 1-exit RMC, and a rational p , decide whether $x_{(1,en,ex)}^* \geq p$.
2. Given a 2-exit RMC, decide whether $x_{(1,en,ex_1)}^* = 1$.
3. **NEW:** ([EY'07,unpublished]) Given a 2-exit RMC, Strong ϵ -approximate \mathbf{x}^* . (!!)

Let's extend RMCs to RMDPs and RSSGs

Recursive Markov Decision Processes (RMDPs): some nodes are *controlled*.


Recursive Simple Stochastic Games (RSSGs): some nodes belong to Player 1 (controller), others to Player 2 (adversary).



RSSGs strictly generalize Condon's finite-state Simple Stochastic Games.

Termination questions for general RMDPs and RSSGs are undecidable, but decidable for the special case of 1-exit RMDPs and RSSGs.....

1-exit RSSGs and nonlinear min/max equations



What is $x_{(f,z,z)}$?

$$x_{(f,z,z)} = 1$$

What is $x_{(f,a,z)}$?

$$x_{(f,a,z)} = \frac{1}{3}x_{(f,h,z)} + \frac{2}{3}x_{(f,(b1,c),z)}$$

What is $x_{(f,(b1,c),z)}$?

$$x_{(f,(b1,c),z)} = x_{(g,c,d)}x_{(f,(b1,d),z)}$$

What is $x_{(f,h,z)}$?

$$x_{(f,h,z)} = \max_{\{\text{neighbors } v \text{ of } h\}} x_{(g,v,z)}$$

What is $x_{(f,(b1,d),z)}$?

$$x_{(f,(b1,d),z)} = \min_{\{\text{neighbors } v \text{ of } (b1,d)\}} x_{(g,v,z)}$$

We get a new system: $\mathbf{x} = P(\mathbf{x})$.

Facts about $P(\mathbf{x})$ (déjà vu)

- $P : \mathbb{R}^n \mapsto \mathbb{R}^n$ is monotone on $[0, 1]^n$, and has a Least Fixed Point.
- **Theorem:** The LFP, $\mathbf{x}^* = \lim_{k \rightarrow \infty} P^k(\mathbf{0})$, is the vector of game values, for the termination game starting at each vertex of the 1-RSSG.

Again, the formula $\exists \mathbf{x}(\mathbf{x} = P(\mathbf{x}) \wedge \mathbf{a} \leq \mathbf{x} \leq \mathbf{b})$ is expressible in ETR, so we can Strong ϵ -approximate \mathbf{x}^* , in **PSPACE**.

- **Theorem:**
 1. The qualitative termination problem for 1-RSSGs, i.e., whether $\mathbf{x}_1^* = 1$, can be decided in **NP** \cap **coNP**.
 2. It is at least as hard as Condon's problem for finite SSGs: "Is the game value $\geq 1/2$?". (And we do not know a reduction in the other direction.)
-

Shapley's Stochastic Games

- As is well known, the quantitative termination problem for Condon's finite SSGs is reducible to that of "discounted" SSGs.
 - In turn, the discounted SSG problem reduces to the problem of ϵ -approximating the value of Shapley's games.
 - Shapley (1953) introduced discounted (finite-state) Stochastic Games, where both players choose moves independently and "concurrently" at each state.
 - The value of Shapley's games (which can be irrational, unlike Condon's games) is the unique (Banach) fixed point of a map $P : [0, 1]^n \mapsto [0, 1]^n$, which can again be described in ETR, and thus ϵ -approximated in **PSPACE**.
-

Shapley reduces to Nash, and to PPAD

Theorem: (new) ([EY'07,unpublished]) *ϵ -approximation of the value of Shapley's games reduces to computing a Weak ϵ -fixed point of $P(\mathbf{x})$, and is thereby in PPAD. (And thus reducible to 2-Nash!)*

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 $P(\mathbf{x})$. ■

Note: [Juba, Blum, and Williams, 2006, unpublished MSc. thesis] have already observed that the quantitative decision problem for Condon's SSGs is in PPAD.

(Their proof has a small gap, essentially because they didn't note the distinction between Weak vs. Strong approximate fixed points, but the gap can be patched.)

Finally, back to our original question:
Strong ϵ -approximation of Nash Equilibria

Theorem: (NEW)[EY'07, unpublished]

1. For every $\epsilon > 0$, **Sqrt-Sum** is P-time reducible to the following decision problem. Given a 3-player (normal form) game, Γ , with the property that:
 - (a) in every NE, player 1 plays exactly the same strategy, and
 - (b) in every NE, player 1 plays its first 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 $(1 - \epsilon)$?).
 2. Thus, if NEs can be Strongly approximated, even in NP, then **Sqrt-Sum** is in NP, and exact Euclidean-TSP is in NP, etc., etc.,
-

theorem continued.....

3. [Allender et. al., '06] Showed that **Sqrt-Sum** reduces to the following more general problem, which they showed lies in the 4th level of the *Counting Hierarchy* ($\mathbf{P}^{\mathbf{P}^{\mathbf{P}^{\mathbf{P}^{\mathbf{P}}}}}$):

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

4. **Theorem: (NEW)**

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.)

theorem continued.....

5. Theorem: (NEW)

Computing a (Strong ϵ -approximate) Nash Equilibrium for a 3-player game is complete for the following class of fixed point problems:

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, compute a (Strong ϵ -approximate) fixed point of F .

Let us call this class of fixed point search problems **FIXP**.

6. These functions include all the fixed point problems we have encountered: n -NASH, SSGs, RMCs, Shapley & “Concurrent” Stochastic Games, 1-RSSGs, 1-RCSGs, and obviously much more.....

However, to get containment in **FIXP** requires, for each such problem, to isolate (in P-time) the relevant fixed point (e.g., the LFP, etc.).....

concluding remarks

Our results open up a “Pandora’s Box” of 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. 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 appears known to be NP-hard.

- The following is **#P**-hard: given a 3-player normal form game, with the property that player 1 plays the same strategy in all NEs, and given $i \in \mathbb{N}$ (in binary), compute the i 'th bit of the probability with which player 1 plays strategy 1 in an NE.
-