Tutorial:
Complexity of Equilibria and Fixed Points: FIXP, FIXP_a, and linear-FIXP$(=$ PPAD$)$
(and some associated open problems)

Kousha Etessami

LFCS, School of Informatics
University of Edinburgh

A few provocative quotes

A few provocative quotes

> " *In general, a game may have several equilibria. Yet uniqueness is crucial . . .. Nash equilibrium makes sense only if each player knows which strategies the others are playing; if the equilibrium recommended by the theory is not unique, the players will not have this knowledge.* "
>
> – Robert J. Aumann (foreword to Harsanyi & Selten's book)

A few provocative quotes

> “ *In general, a game may have several equilibria. Yet uniqueness is crucial . . .. Nash equilibrium makes sense only if each player knows which strategies the others are playing; if the equilibrium recommended by the theory is not unique, the players will not have this knowledge.* ”
>
> – Robert J. Aumann (foreword to Harsanyi & Selten’s book)

> “ *In comparative statics . . . we study the response of our [market] equilibrium to designated changes in the parameters.* ”
>
> – Paul A. Samuelson (Foundations of Economic Analysis)

A few provocative quotes

“ *In general, a game may have several equilibria. Yet uniqueness is crucial . . .. Nash equilibrium makes sense only if each player knows which strategies the others are playing; if the equilibrium recommended by the theory is not unique, the players will not have this knowledge.* ”
– Robert J. Aumann (foreword to Harsanyi & Selten’s book)

“ *In comparative statics . . . we study the response of our [market] equilibrium to designated changes in the parameters.* ”
– Paul A. Samuelson (Foundations of Economic Analysis)

“ *Post the 2008-09 crisis, the world economy is pregnant with multiple equilibria.* ”
“ *. . . it may not take much . . . to move from the good to the bad equilibrium.* ”
– Olivier Blanchard, IMF Chief Economist (IMF Blog, 2011-13)

A few provocative quotes

> “A characteristic feature [of] economics is that for us the equations of equilibrium constitute the center of our discipline. By contrast, other sciences put more emphasis on the dynamic laws of change. The reason... is that economists are good at recognizing a state of equilibrium, but are poor at predicting precisely how an economy in disequilibrium will evolve...”
>
> – Mas-Colell, Whinston, & Green (Microeconomic Theory)

Ok, let's wish away the multiple equilibria, for now

What is the complexity of the following search problem?

Given a 2-player bimatrix game, Γ, with the promise that Γ has a unique Nash equilibrium (NE), compute that unique NE.

Answer:

Ok, let's wish away the multiple equilibria, for now

What is the complexity of the following search problem?

Given a 2-player bimatrix game, Γ, with the promise that Γ has a unique Nash equilibrium (NE), compute that unique NE.

Answer: We do not know. (It is in PPAD, but unlikely to be PPAD-hard.)

(N.B. Ruta Mehta has made nice progress recently toward this question. We will revisit it later, when we discuss open problems and conjectures.)

Ok, let's wish away the multiple equilibria, for now

What is the complexity of the following search problem?

Given a 2-player bimatrix game, Γ, with the promise that Γ has a unique Nash equilibrium (NE), compute that unique NE.

Answer: We do not know. (It is in PPAD, but unlikely to be PPAD-hard.)

(N.B. Ruta Mehta has made nice progress recently toward this question. We will revisit it later, when we discuss open problems and conjectures.)

What about for 3-player games with a unique NE?

Given a 3-player normal form game, Γ, with the promise that it has a unique NE, compute any vector with ℓ_∞-distance $\leq 1/2 - \epsilon$ from the unique NE.

Answer:

Ok, let's wish away the multiple equilibria, for now

What is the complexity of the following search problem?

Given a 2-player bimatrix game, Γ, with the promise that Γ has a unique Nash equilibrium (NE), compute that unique NE.

Answer: We do not know. (It is in PPAD, but unlikely to be PPAD-hard.)

(N.B. Ruta Mehta has made nice progress recently toward this question. We will revisit it later, when we discuss open problems and conjectures.)

What about for 3-player games with a unique NE?

Given a 3-player normal form game, Γ, with the promise that it has a unique NE, compute any vector with ℓ_∞-distance $\leq 1/2 - \epsilon$ from the unique NE.

Answer: This is "hard": even placing it in FNP would resolve long standing open problems in arithmetic-vs.-Turing complexity.(PosSLP-hard.)

What about for market equilibrium?

Given a Arrow-Debreu exchange economy, with *n* commodities, and with market excess demands given by nonlinear functions (satisfying Walras's law and homogeneity of degree 0), and with the promise that there is a unique (normalized) market price equilibrium, compute any vector with ℓ_∞-distance $\leq 1/2 - \epsilon$ from the unique market equilibrium.

Answer:

What about for market equilibrium?

Given a Arrow-Debreu exchange economy, with *n* commodities, and with market excess demands given by nonlinear functions (satisfying Walras's law and homogeneity of degree 0), and with the promise that there is a unique (normalized) market price equilibrium, compute any vector with ℓ_∞-distance $\leq 1/2 - \epsilon$ from the unique market equilibrium.

Answer: Again, this is "hard": PosSLP-hard.

What makes an equilibrium/fixed point problem "hard"??

Note: These problems are in general not NP-hard, because existence of a solution (equilibrium/fixed point), is guaranteed by a classic fixed point theorem (e.g., Brouwer's, Kakutani's, Banach's, Tarski's, . . .).

What makes an equilibrium/fixed point problem “hard”??

Note: These problems are in general not NP-hard, because existence of a solution (equilibrium/fixed point), is guaranteed by a classic fixed point theorem (e.g., Brouwer’s, Kakutani’s, Banach’s, Tarski’s, . . .).

- PPAD-hardness captures a combinatorial difficulty for computing/approximating an equilibrium or fixed point.

What makes an equilibrium/fixed point problem “hard”??

Note: These problems are in general not NP-hard, because existence of a solution (equilibrium/fixed point), is guaranteed by a classic fixed point theorem (e.g., Brouwer’s, Kakutani’s, Banach’s, Tarski’s, . . .).

- PPAD-hardness captures a combinatorial difficulty for computing/approximating an equilibrium or fixed point.
- But there can also be another, numerical, difficulty for approximating a (real-valued) equilibrium or fixed point, which is not captured by PPAD-hardness.

  It is captured by “PosSLP-hardness”.

These two kinds of difficulties are somewhat “orthogonal”.

$\mathrm{FIXP}_{(a)}$-complete problems have both of these difficulties.

Rich landscape within FIXP:

Numerical Difficulty

	No	P.G.–hard	PPAD–hard
PosSLP–hard *	approx–Recursive Markov chains exact–Branching processes exact–Branching–MDPs	exact–Branch–simple–stoc–game approx–Unique–nonlinear Brouwer fixed point ??approx–Unique–3–player–Nash?? exact–concurrent–stocastic–game exact–Shapley–stocastic–game	approx–3–player–Nash approx–nonlin.–Arrow–Debreu market equilibrium approx–nonlinear Brouwer fixed point FIXP_a–complete
* No	PIT / ACIT exact–linear–Arrow–Debreu market equilibrium approx–Branching–MDPs approx–Branching–process exact–MDPs	approx–Branch–simple–stoc–game approx–Shapley–stochastic–game exact–Unique–piecewise–linear Brouwer fixed point ??exact–Unique–2–player–Nash?? exact–Condon–simple–stoc–game exact–mean–payoff–game parity–game	exact–2–player–Nash exact–Arrow–Debreu market equilibrium with SPLC utilities exact–piecewise–linear Brouwer fixed point "Almost"–nonlinear–Brouwer fixed point "Almost"–(epsilon)–Nash for >= 3 players

Combinatorial Difficulty

Outline of tutorial

- Background: Games, Equilibria, Brouwer Fixed Points.
- "Almost" vs. "Near" approximation of Fixed Points.
- Scarf's classic algorithm for "Almost"-approximation of a fixed point.
- The complexity class PPAD, and "Almost" approximation.
- PPAD-completeness results for ϵ-almost-Nash, and 2-player-Nash.
- Hardness of "Near" approximation: arithmetic circuits & PosSLP.
- The complexity classes FIXP and FIXP$_a$.

  3-player Nash (approx-Nash) is FIXP$_{(a)}$-complete.
- linear-FIXP $=$ PPAD.
- Many other FIXP$_a$ approximation problems:
  Market price equilibria,
  (Branching) stochastic processes/games,
  Recursive Markov Chains, ...
- Open problems and future challenges.

Nash Equilibria

A mixed strategy profile x is called:

- a Nash Equilibrium (NE) if:
  $\forall$ players 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 (ϵ-almost-NE), for $\epsilon > 0$, if:
  $\forall$ players 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 (NE) if:
  $\forall$ players 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 (ϵ-almost-NE), for $\epsilon > 0$, if:
  $\forall$ players 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:

So, $F_\Gamma(x)$ is expressed by a formula using gates $\{+, -, \times, /, \max, \min\}$.

Question

What is the complexity of the following search problem:

> *("Near") 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 Nash Equilibrium x^* of Γ with:*

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

Note:

This is not the same thing as asking for an ϵ-almost-NE.

Almost vs. Near 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:

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

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

- (**Near**) 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 a ϵ-(Almost) fixed point 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 ϵ-almost fixed points.
5. They need not in general be anywhere near an actual fixed point.

Sperner’s Lemma

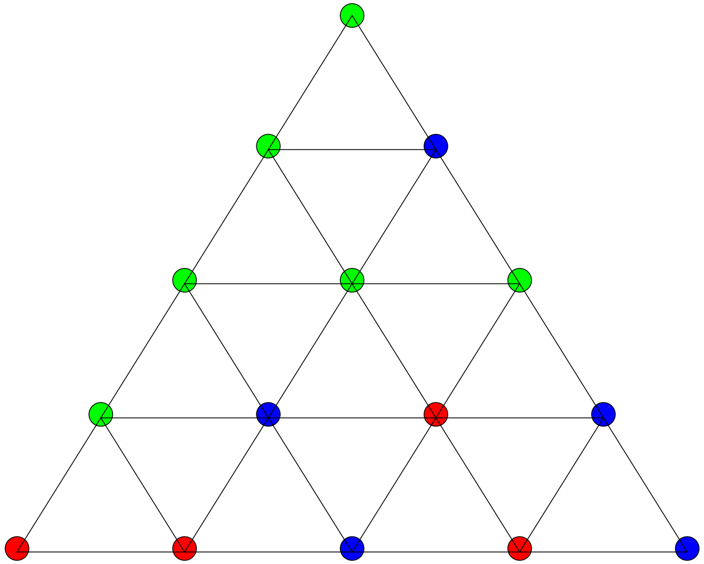

“Proof” of Sperner’s lemma

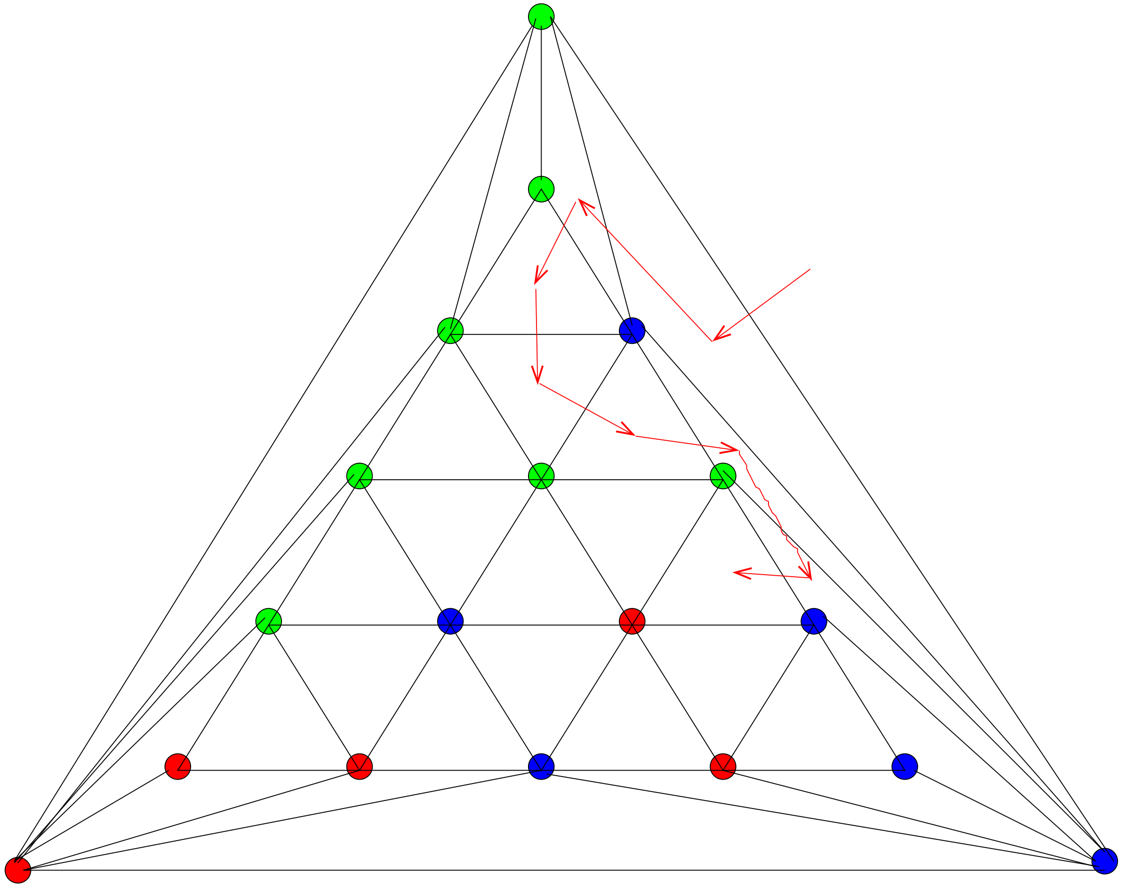

(Things are more involved in higher dimensions.)

The underlying "directed lines" parity argument in Scarf's algorithm

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

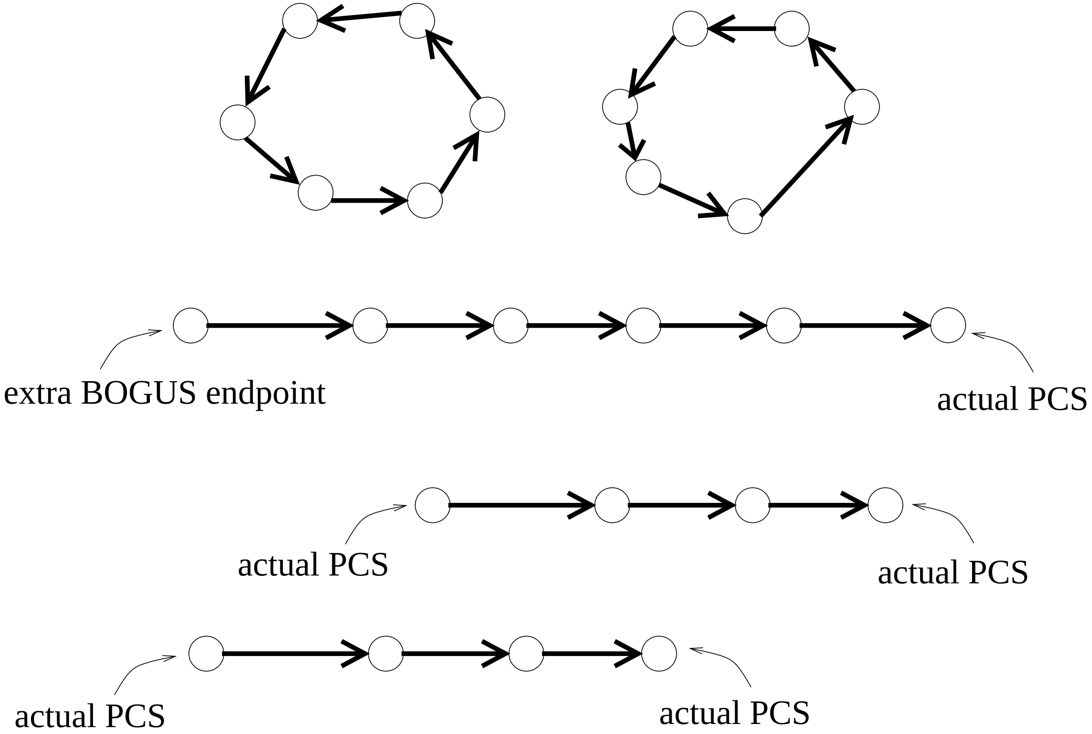

ϵ-almost-NEs are ϵ-almost-fixed points

Proposition

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

ϵ-almost-NEs are ϵ-almost-fixed points

Proposition

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

Thus, to compute an ϵ-almost-NE, simply apply Scarf's algorithm to F_Γ.

ϵ-almost-NEs are ϵ-almost-fixed points

Proposition

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

Thus, to compute an ϵ-almost-NE, simply apply Scarf's algorithm to F_Γ.

It also follows from this that computing a ϵ-almost-NE is in PPAD.

Standard definition of PPAD

Papadimitriou (1992) defined **PPAD**, based on the "directed line" parity argument, to capture (almost) 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.)

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

By Scarf's algorithm, computing a ϵ-almost-NE is in PPAD.

Theorem

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

But what if we want to do near approximation of a 3-player NE, or near approximation of a fixed point?

Scarf's algorithm does not in general yield something ϵ-near a fixed point.

Why care about near approximation of equilibria/fixed points?

For many problems, the goal is to approximate a specific quantity which happens to be given by the (unique) Brouwer fixed point of some function.

Examples:

–the value of Shapley's Stochastic Games (or Condon's Simple S.G.'s) ;
–the (optimal) extinction probability (or value) of a
 Branching (Markov Decision) Processes and Branching S. G.'s;
–the termination probability of a Recursive Markov chain;
–the unique market equilibrium in certain specific kinds of markets;
–the unique (refined) equilibrium of specific kinds of games;

In these contexts, an "almost" fixed point may tell us nothing about the unique fixed point that we are after.

A basic upper bound for Near ϵ-approximation of Nash

Proposition

Given game Γ and $\epsilon > 0$, we can ϵ-Near 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 Near 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 $\sum_{i=1}^{n} \sqrt{d_i} \leq k$.

Solvable in PSPACE.

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

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

PosSLP captures all of polynomial time in the unit-cost arithmetic RAM model of computation.

[Allender, Bürgisser, Kjeldgaard-Pedersen, Miltersen,2006] Gave a (Turing) reduction from Sqrt-Sum to PosSLP and showed both can be decided in the Counting Hierarchy: $P^{PP^{PP^{PP}}}$. Nothing better is known.

why isn't PosSLP easy??

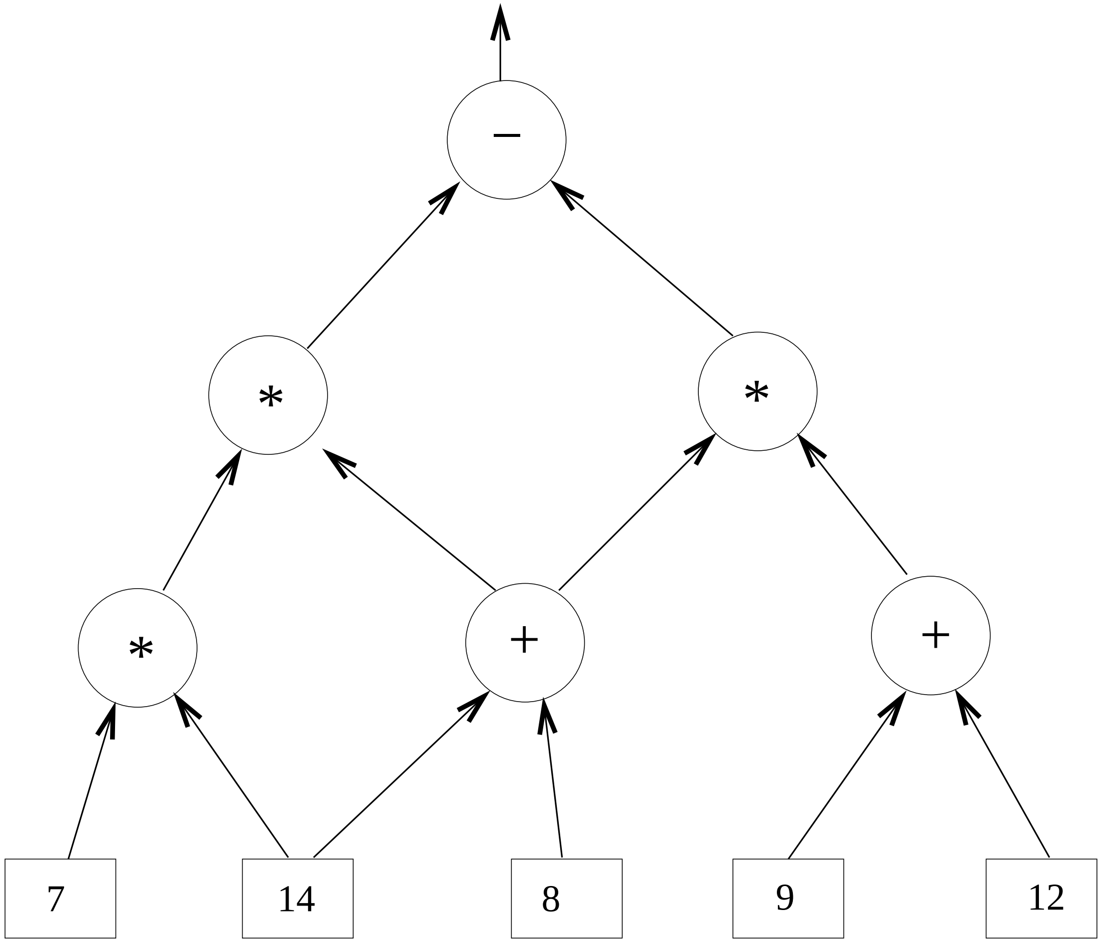

Note: even the much easier EquSLP ("equal to 0") is P-time equivalent to polynomial identity testing (PIT/ACIT), as shown by [ABKM'06].

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

Theorem ([E.-Yannakakis'07])

Any non-trivial Near approximation of an NE is **PosSLP***-hard.*

More precisely: for every fixed $\epsilon > 0$,
PosSLP *is P-time reducible to the following problem:*

Given a 3-player normal form game, Γ, *with the promise that:*

1. Γ *has a unique NE,* x^*, *which is fully mixed, and*
2. *In* x^*, *the probability that player 1 plays pure strategy* α *is either:*

$$\text{(a.)} \quad < \epsilon \quad , \qquad\qquad \text{or} \qquad\qquad \text{(b.)} \quad \geq (1 - \epsilon)$$

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

ϵ-almost-NEs can be very far from actual NEs

Proposition ([E.-Yannakakis'07])

There are constants $c, c' > 0$, such that for any $n \in \mathbb{N}_+$, there is a game, Γ_n, with encoding size $\Theta(n)$, which has a

$$\left(\frac{1}{2^{2^{c' n^c}}}\right)\text{-}almost\text{-}NE$$

which has ℓ_∞-distance 1 from the (unique) NE of Γ_n.

Note: This is the worst ℓ_∞-distance possible.

The complexity class FIXP (and FIXP_a)

FIXP (FIXP_a) is a class of real-valued (discrete) total search problems:

FIXP (FIXP_a)

- **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 can be much more general than $[0,1]^n$.)
- **Output:** Compute a (ϵ-near approximate) fixed point of F.

Close these problems under suitable (P-time) reductions.
The resulting class is called FIXP (FIXP_a).

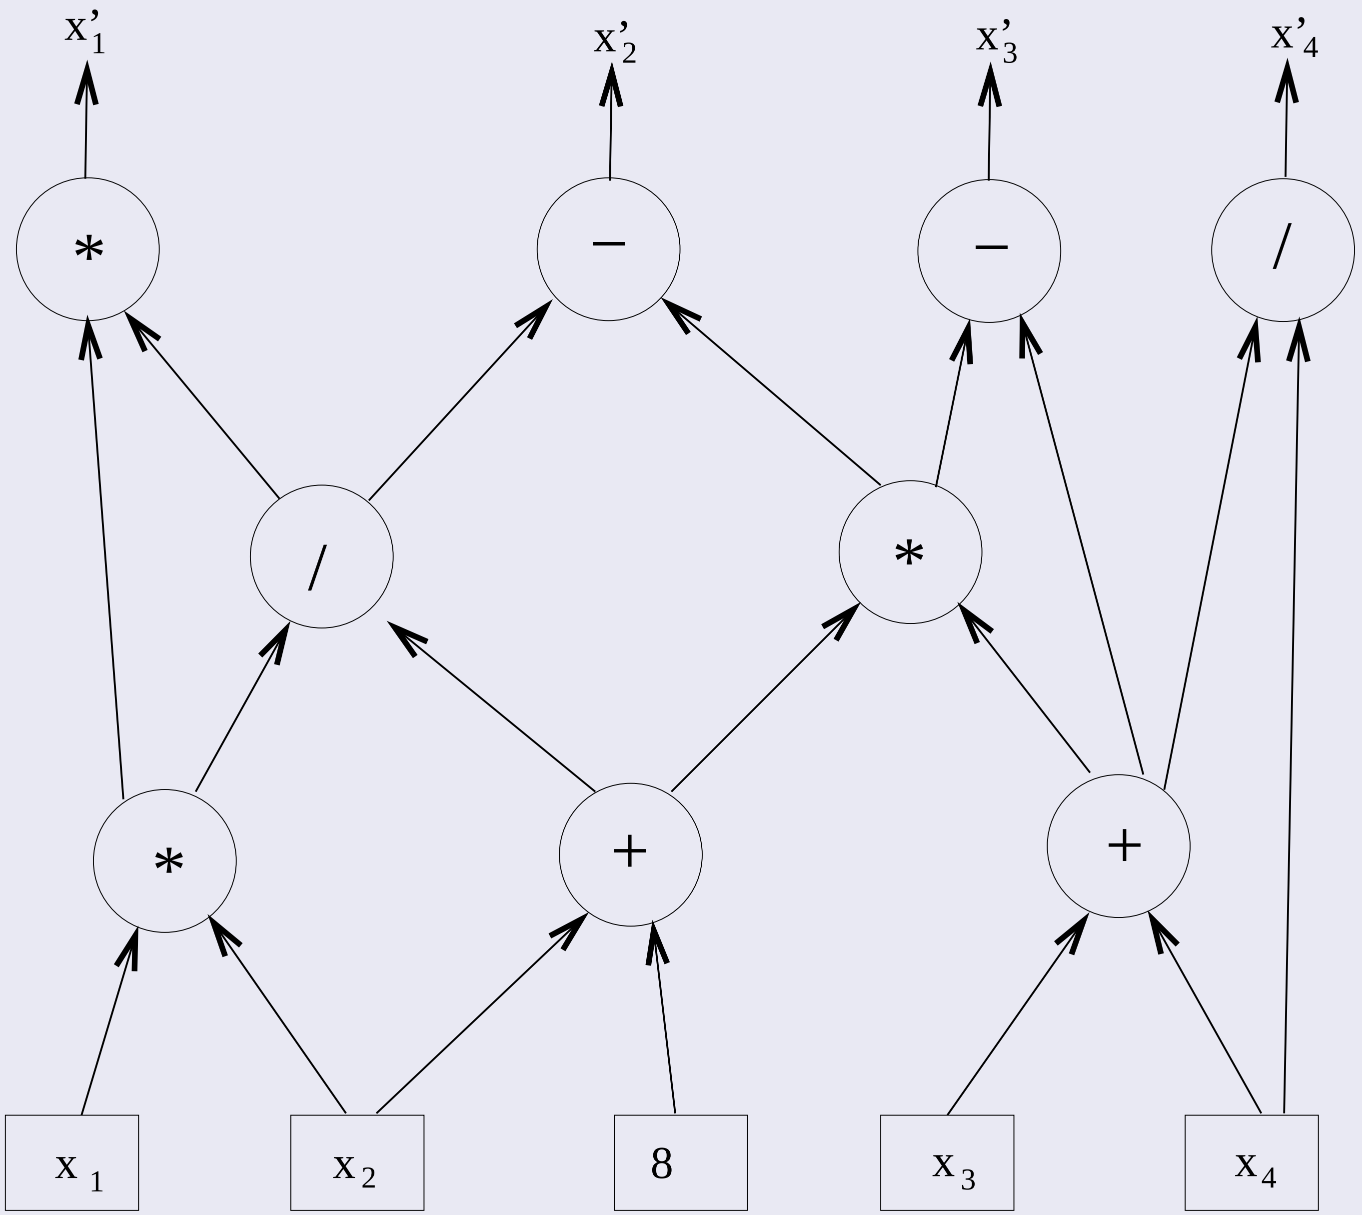
x'1
x'2
x'3
x'4
*
−
−
/
/
*
*
+
+
x 1
x 2
8
x 3
x 4

Nash is FIXP-complete

Theorem ([E.-Yannakakis'07])

Computing (ϵ-near approximating) a 3-player Nash Equilibrium given the game (and given $\epsilon > 0$) is:

FIXP-*complete (* FIXP_a*-complete, respectively).*

Theorem ([E.-Yannakakis'07])

The gates $\{+, *, \max\}$ *are sufficient to capture all of* $\text{FIXP}_{(a)}$.

Furthermore, allowing gates, $\sqrt[k]{\cdot}$*, for fixed* k*, does not add any power to* $\text{FIXP}_{(a)}$.

Very brief sketch of some proof ideas

- Suppose, given a FIXP circuit C, we can create a (3-player) game such that, in any NE, Player 1 plays strategy α with probability $> 1/2$ iff $C > 0$ and with probability $< 1/2$ iff $C < 0$.
  (Assume wlog that $C = 0$ can't happen.)
- Add an extra player with 2 pure strategies, who gets payoff 1 if it "guesses correctly" whether player 1 plays pure strategy α or not, and payoff 0 otherwise.
- In any NE, the new player will play one of its two pure strategies with probability 1.
  Deciding which of the two solves PosSLP.

A key ingredient in our proofs

Two beautiful results 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 separable-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".
- Transform circuit to a game with a large (but bounded) number of players, using suitable *gadgets*.
  All key gadgets can be derived from (Bubelis'79)'s constructions.

  (Alternatively, the gadgets of (Golberg-Papadimitriou'06), (Daskalakis-Golberg-Papadimitriou'06), (Chen-Deng'06) can also be used.)
- Reduce to 3-players: again use (Bubelis '79).

Alternative characterizations of PPAD

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

Theorem ([E.-Yannakakis'07])

The following are all P-time equivalent:

1. PPAD
2. linear-FIXP
3. *exact fixed point problem for "polynomial piecewise-linear functions".*
4. ϵ-*"Almost"-fixed point problem for "polynomially computable" and "polynomially continuous" functions,* $F_I(x)$, *given by input instance* I, *and given also* $\epsilon > 0$ *as input (in binary).*

Alternative characterizations of PPAD

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

Theorem ([E.-Yannakakis'07])

The following are all P-time equivalent:

1. PPAD
2. linear-FIXP
3. *exact fixed point problem for "polynomial piecewise-linear functions".*
4. *ϵ-"Almost"-fixed point problem for "polynomially computable" and "polynomially continuous" functions, $F_I(x)$, given by input instance I, and given also $\epsilon > 0$ as input (in binary).*
5. *[R. Mehta, 2014]:* *2-variable*-linear-FIXP *(!!)*

Recursive Markov Chains

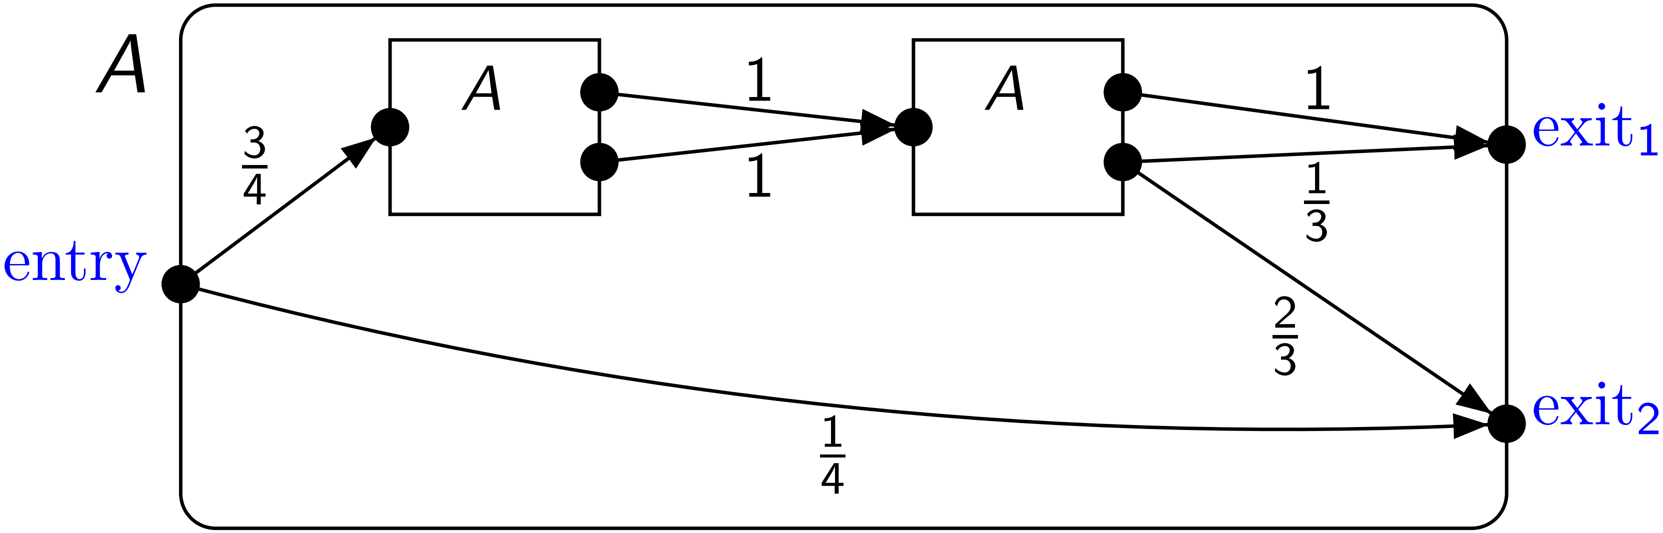

What is the probability of terminating at exit$_2$, starting at entry?

$$x_2 \quad =$$

Recursive Markov Chains

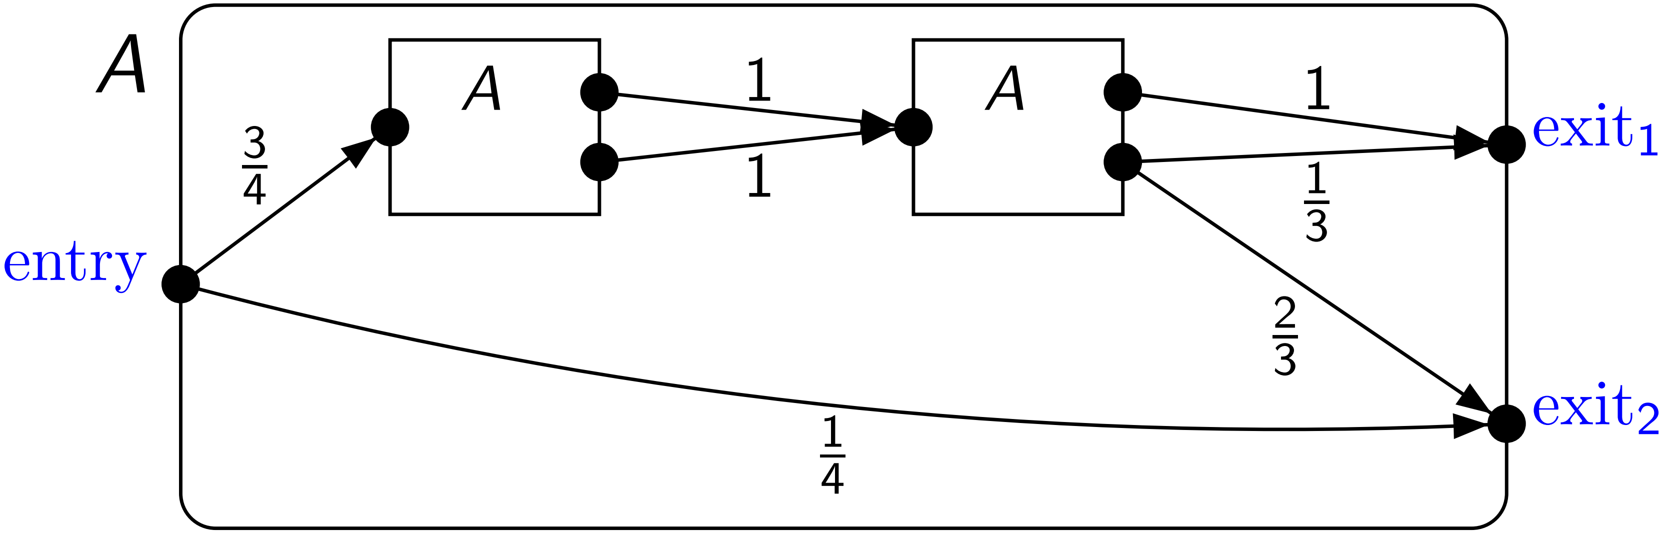

What is the probability of terminating at exit$_2$, starting at entry?

$$x_2 = \frac{1}{4} + \frac{1}{2}x_2^2 + \frac{1}{2}x_1x_2 \qquad \text{(Note: coefficients sum to } > 1\text{)}$$

$$x_1 = \frac{3}{4}x_1^2 + \frac{3}{4}x_2x_1 + \frac{1}{4}x_1x_2 + \frac{1}{4}x_2^2$$

Fact: The Least Fixed Point (LFP), $q^* \in [0,1]^n$, gives the termination probabilities.

approximation for Recursive Markov chains

Theorem

1. *[EY07]: Any non-trivial (near) approximation of the termination probabilities q^* of an RMC is PosSLP-hard.*

   In fact, deciding whether (a.) $q_1^* = 1$ or (b.) $q_1^* < \epsilon$, is PosSLP-hard.

2. *[ESY12]: ϵ-(near)-approximation of q^* is in FIXP$_a$.*

   (Can be reduced to a unique Brouwer fixed point problem.)

3. *[EY'05]: But there appears to be no combinatorial difficulty for approximating q^*: a decomposed Newton's method converges monotonically, starting from $\mathbf{0}$, to q^*.*

Branching Markov Decision Processes

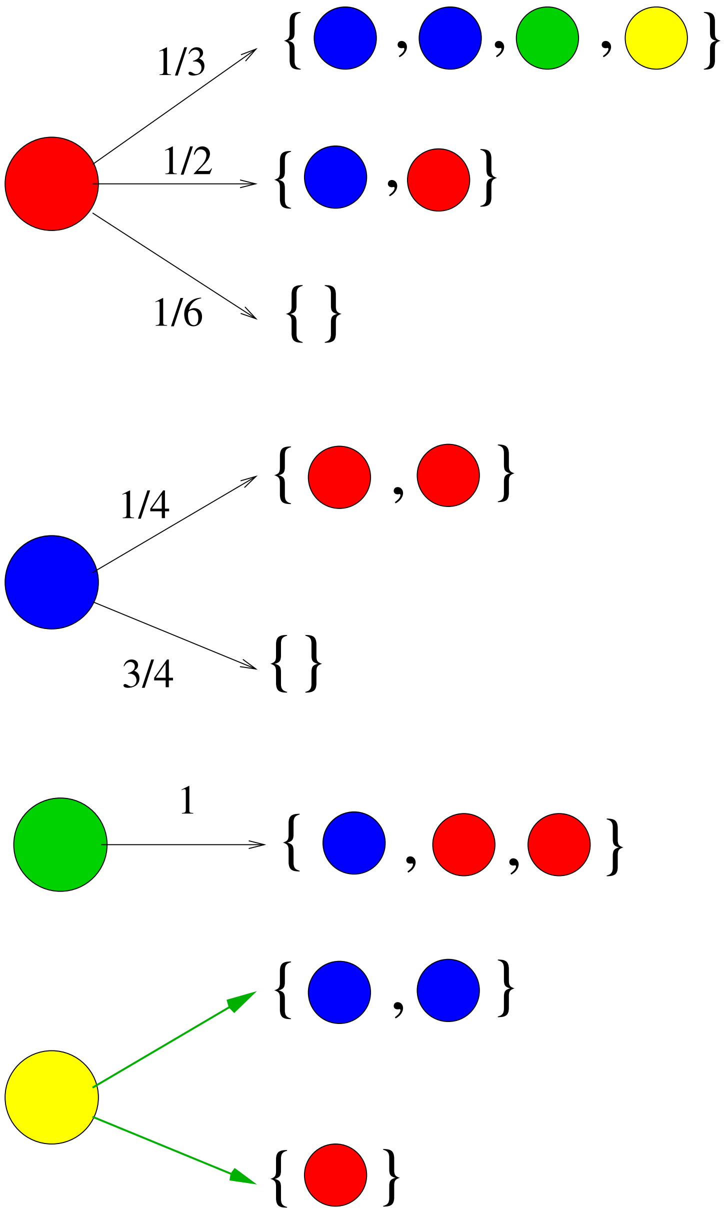

Branching Markov Decision Processes

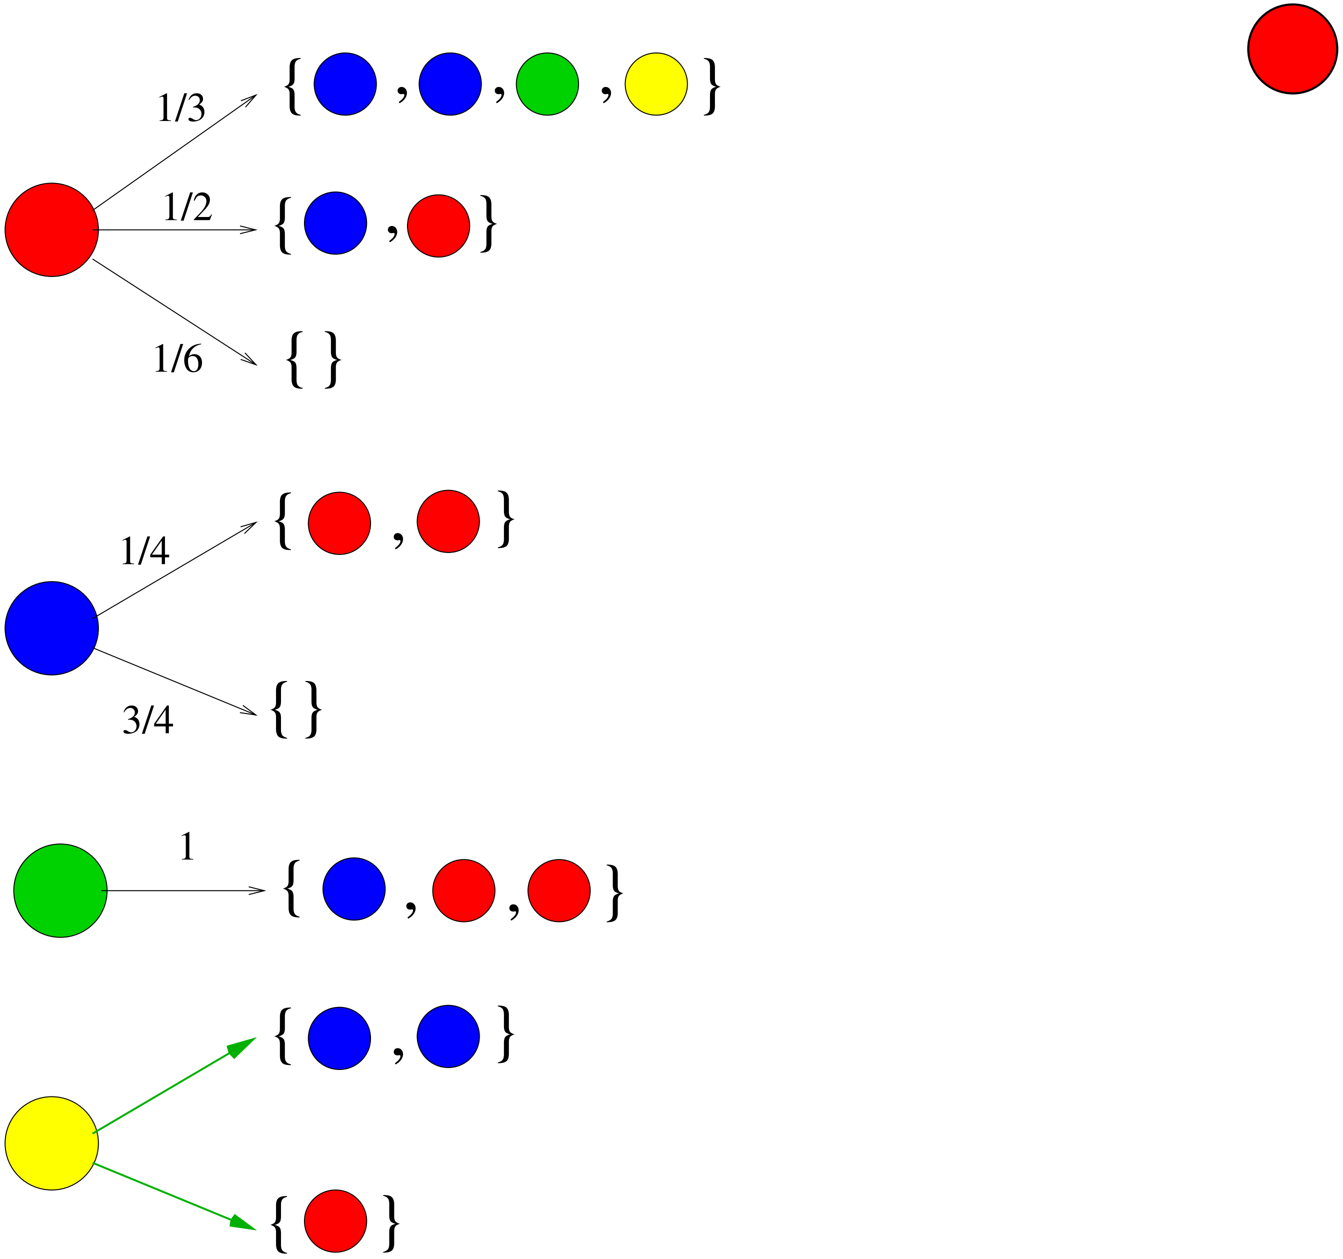

Branching Markov Decision Processes

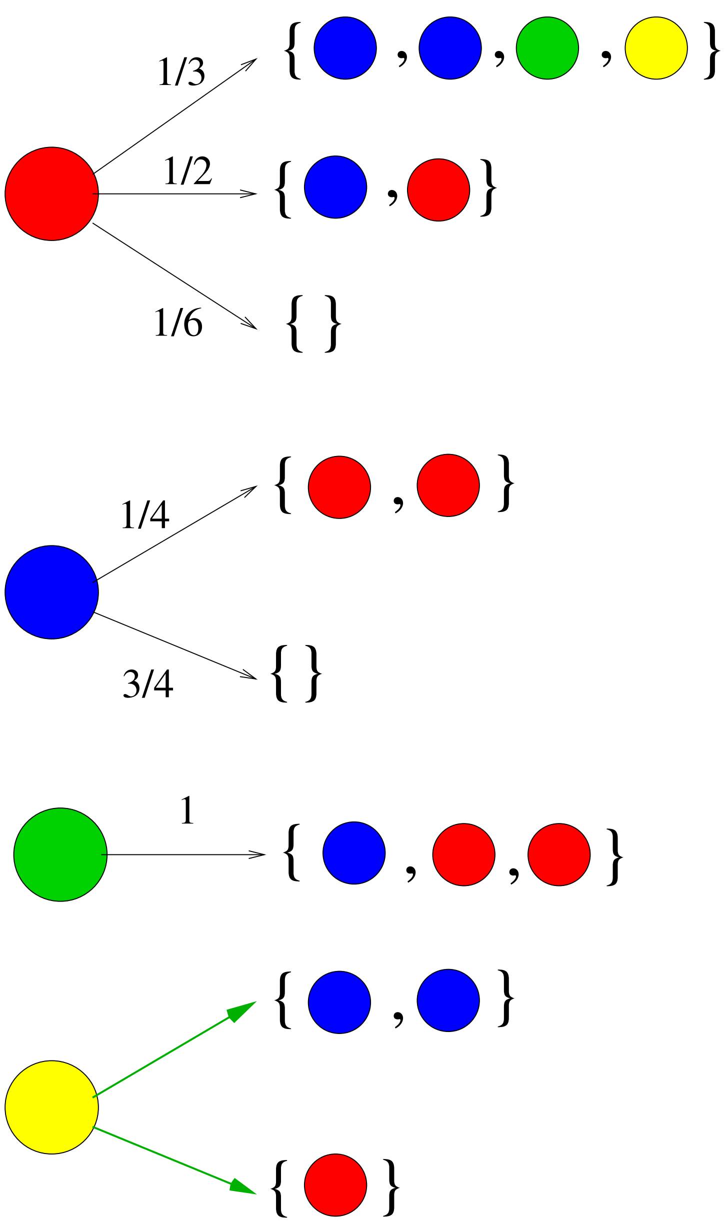

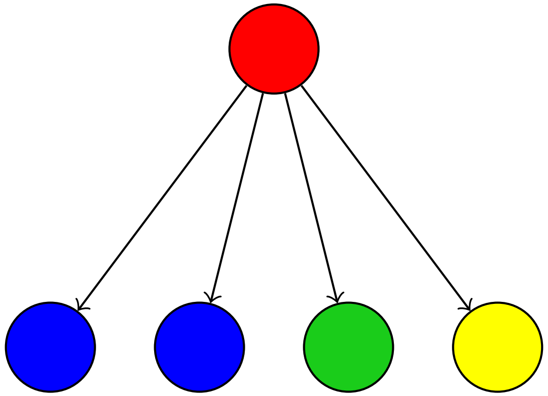

Branching Markov Decision Processes

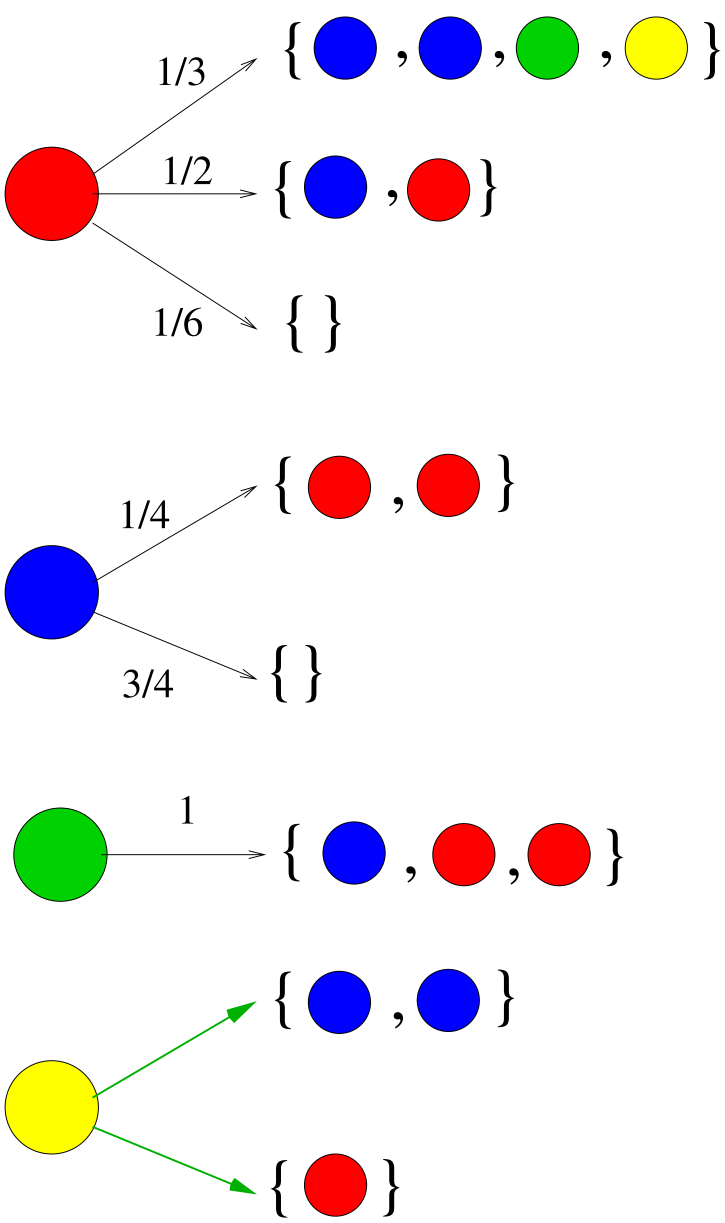

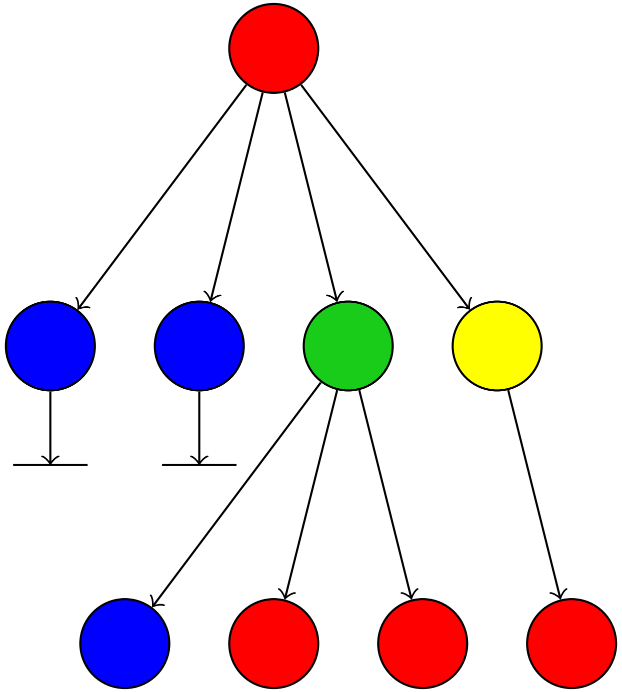

Branching Markov Decision Processes

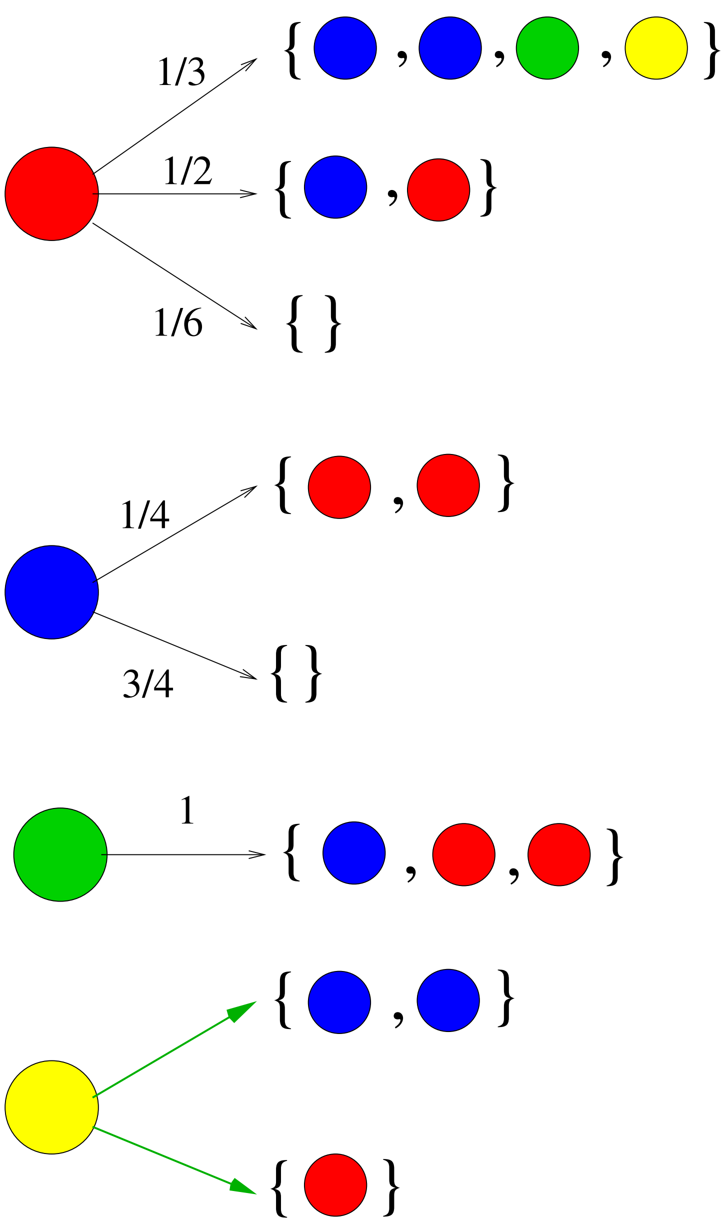

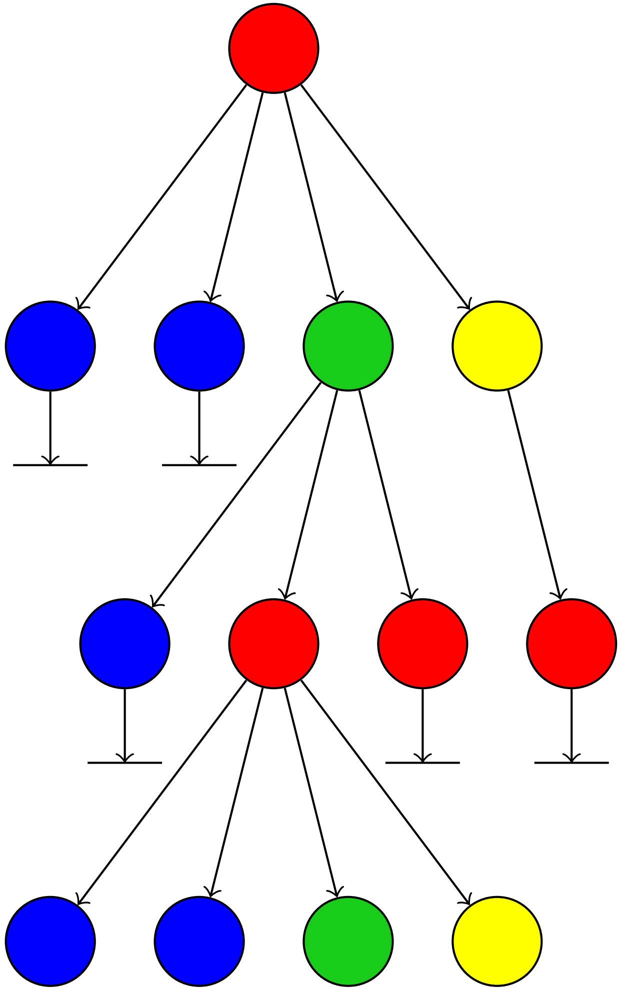

Branching Markov Decision Processes

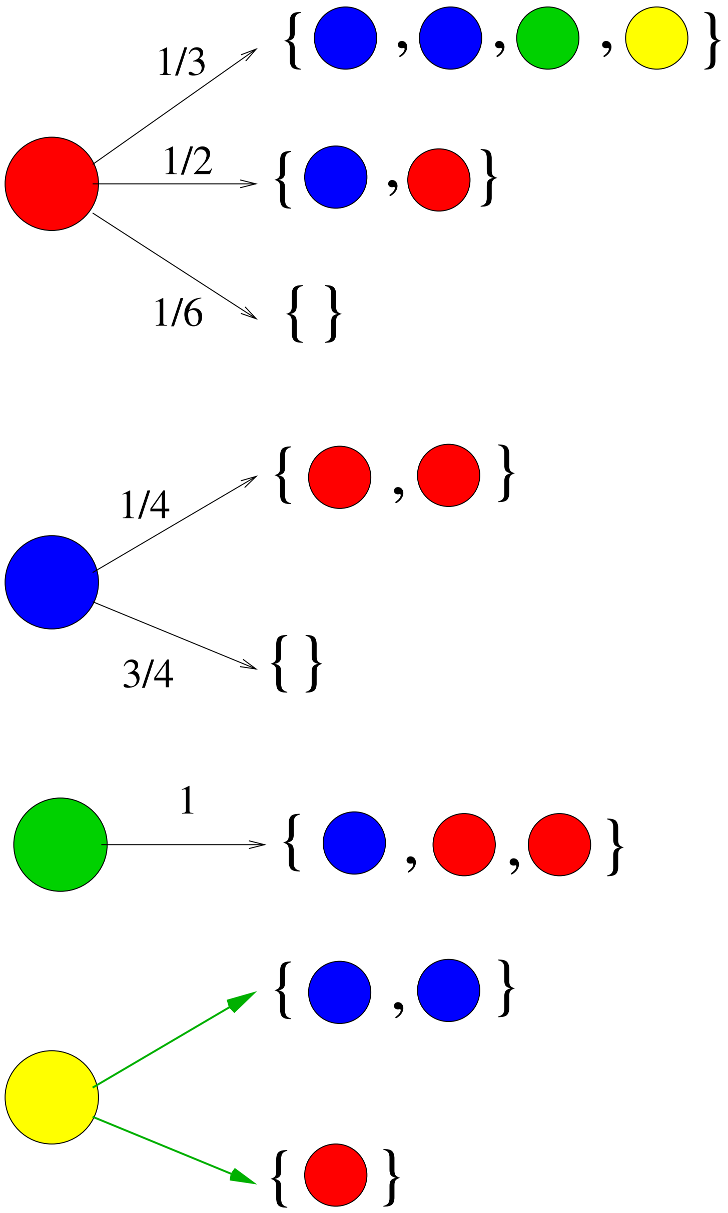

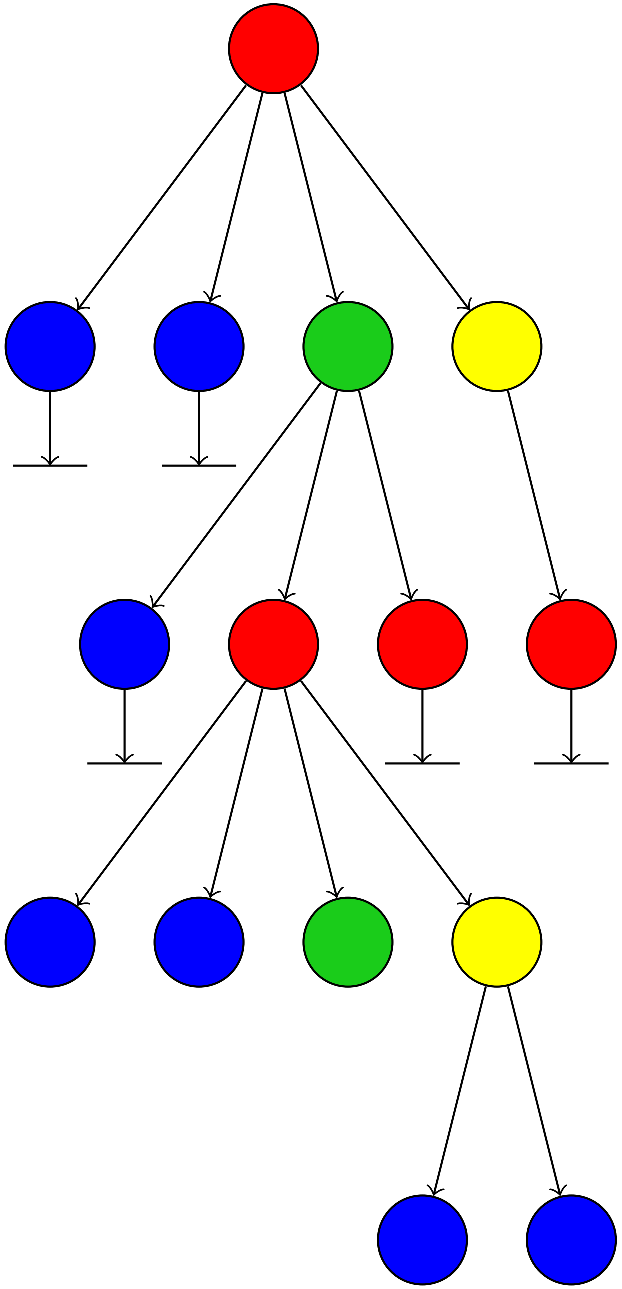

Branching Markov Decision Processes

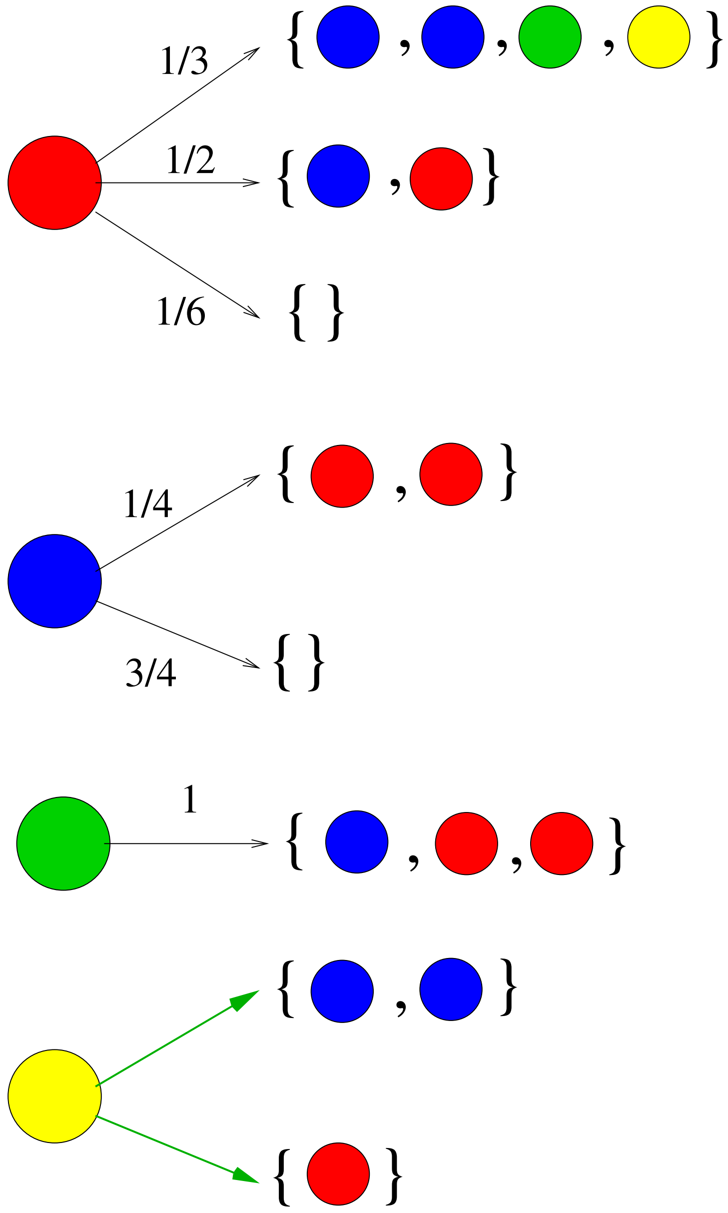

Question

What is the **maximum** probability of extinction, starting with one ● ?

Branching Markov Decision Processes

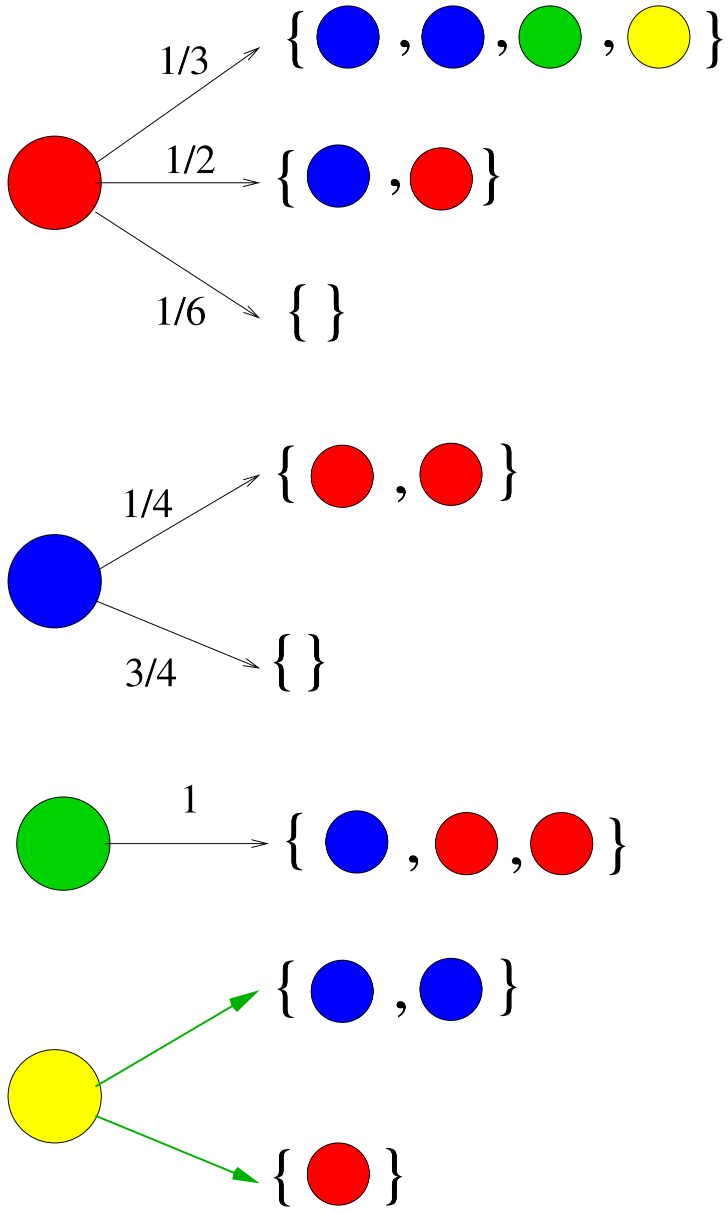

Question

What is the maximum probability of extinction, starting with one ● ?

$$
\begin{aligned}
x_R &= \frac{1}{3}x_B^2 x_G x_Y + \frac{1}{2}x_B x_R + \frac{1}{6}\\
x_B &= \frac{1}{4}x_R^2 + \frac{3}{4}\\
x_G &= x_B x_R^2\\
x_Y &=
\end{aligned}
$$

Branching Markov Decision Processes

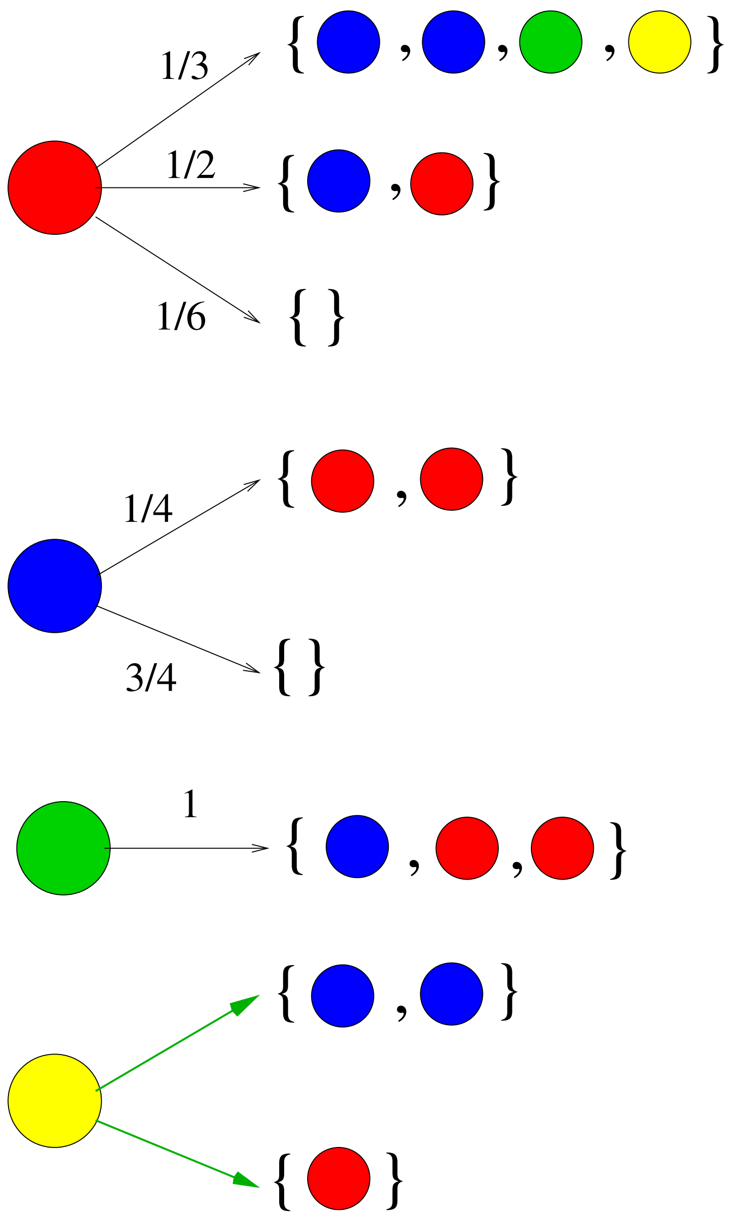

Question

What is the maximum probability of extinction, starting with one ●?

$$
\begin{aligned}
x_R &= \frac{1}{3}x_B^2 x_G x_Y + \frac{1}{2}x_B x_R + \frac{1}{6}\\
x_B &= \frac{1}{4}x_R^2 + \frac{3}{4}\\
x_G &= x_B x_R^2\\
x_Y &= \max\{x_B^2, x_R\}
\end{aligned}
$$

We get fixed point equations, $\bar{\mathbf{x}} = P(\bar{\mathbf{x}})$.

Fact [E.-Yannakakis'05]

The maximum extinction probabilities are the least fixed point, $\mathbf{q}^* \in [0,1]^3$, of $\bar{\mathbf{x}} = P(\bar{\mathbf{x}})$.

Branching Markov Decision Processes

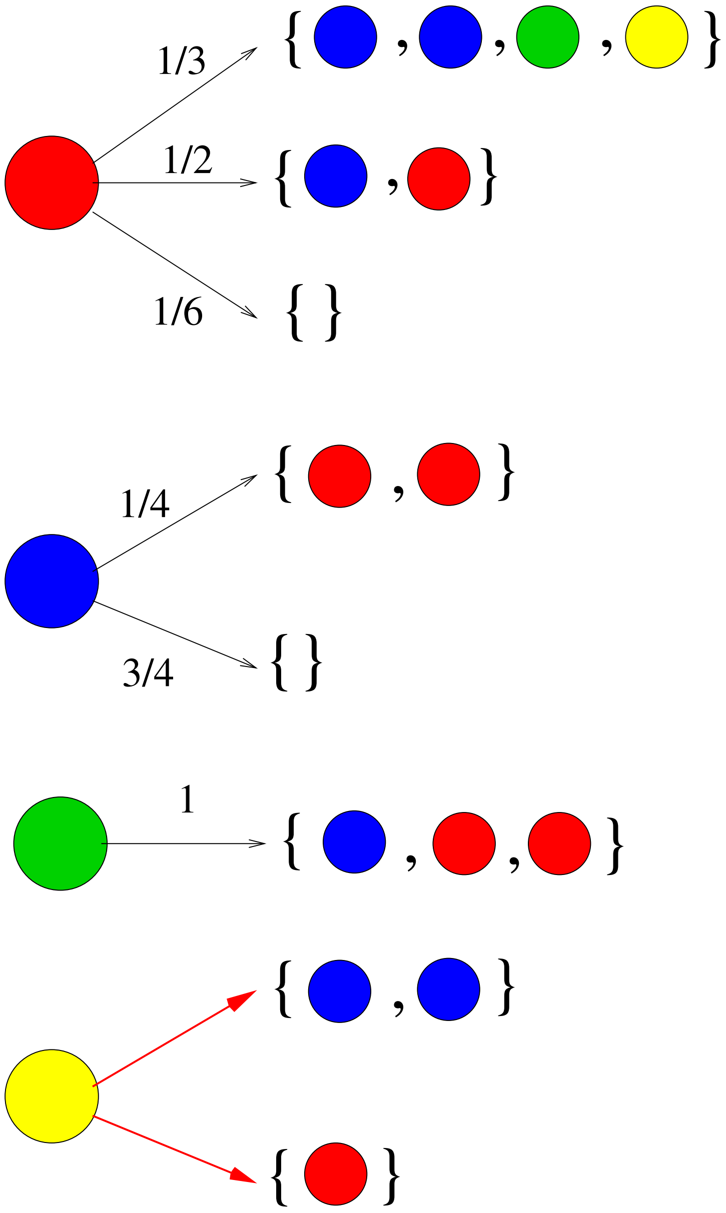

Question

What is the minimum probability of extinction, starting with one 🔴 ?

$$
\begin{aligned}
x_R &= \frac{1}{3}x_B^2 x_G x_Y + \frac{1}{2}x_B x_R + \frac{1}{6}\\
x_B &= \frac{1}{4}x_R^2 + \frac{3}{4}\\
x_G &= x_B x_R^2\\
x_Y &= \min\{x_B^2, x_R\}
\end{aligned}
$$

We get fixed point equations, $\bar{\mathbf{x}} = P(\bar{\mathbf{x}})$.

Fact [E.-Yannakakis'05]

The minimum extinction probabilities are the least fixed point, $\mathbf{q}^* \in [0,1]^3$, of $\bar{\mathbf{x}} = P(\bar{\mathbf{x}})$.

Max-Probabilistic Polynomial Systems of Equations

A Max-Probabilistic Polynomial System (maxPPS) is a system

$$\mathbf{x}_i = \max\{p_{i,j}(\mathbf{x}) : j = 1, \ldots, m_i\} \qquad i = 1, \ldots, n$$

of n equations in n variables, where each $p_{i,j}(x)$ is a probabilistic polynomial. We denote the entire system by:

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

Min-Probabilistic Polynomial Systems (minPPSs) defined similarly.

These are Bellman optimality equations for maximizing (minimizing) extinction probabilities in a BMDP.

We use max/minPPS to refer to either a maxPPS or an minPPS.

Basic properties of max/minPPSs, $\mathbf{x} = P(\mathbf{x})$

$P : [0,1]^n \to [0,1]^n$ defines a monotone map on $[0,1]^n$.

Proposition. [E.-Yannakakis'05]

- *Every max/minPPS, $x = P(x)$ has a least fixed point, $q^* \in [0,1]^n$.*
- $q^* = \lim_{k\to\infty} P^k(\mathbf{0})$.
- *q^* is the vector of optimal extinction probabilities for the BMDP.*
- *[EY'07] Deciding whether $q_1^* > 1/2$ is PosSLP-hard.*
- *[ESY'12] ϵ-Near approximation of q^* is in* FIXP_a.

P-time approximation for BMDPs and max/minPPSs

Theorem ([E.-Yannakakis'06])

Given a BMDP, deciding whether the optimal (max or min) extinction probability is $q_i^* = 1$ *is in P-time.*

Reduces to a spectral radius optimization problem for non-negative matrices (solvable using LP).

Theorem ([E.-Stewart-Yannakakis,2012])

Given a BMDP, or max/minPPS, $\mathbf{x} = P(\mathbf{x})$*, with LFP* $\mathbf{q}^* \in [0,1]^n$*, we can compute a rational vector* $\mathbf{v} \in [0,1]^n$ *such that*

$$\|\mathbf{v} - \mathbf{q}^*\|_\infty \leq 2^{-j}$$

in time polynomial in the encoding size $|P|$ *of the equations, and in* j*.*

We establish this via a Generalized Newton's Method that uses linear programming in each iteration.

Branching Simple Stochastic Games (BSSGs)

Both Max and Min types (two players): their goal is to maximize (minimize) extinction probability. We again get max-&-min-PPS equations whose LFP gives the game value.

Condon's finite-state Simple Stochastic Games (SSGs) are a special case.

Theorem

Given a BSSG,

1. *[EY'06]: deciding whether extinction value* $q_1^* = 1$ *is in* **NP** $\cap$ **coNP**. *And it is at least as hard as computing the exact value of Condon's finite-state SSG.*
2. *[ESY'12]: Given* $\epsilon > 0$, *computing a vector* $v \in [0,1]^n$, *such that* $\|v - q^*\|_\infty \leq \epsilon$, *is in* FIXP_a, *and in* **PLS**.

(But we still do not know whether it is in PPAD.*)*

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 ([E.-Yannakakis'07])

For Shapley's stochastic games:

1. *Computing the game value is in* FIXP.
2. *The (Near) 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, ϵ-"Almost" 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 computing the value of Condon's SSGs is in PPAD. But this can be shown more easily by observing unique linear-FIXP equations for the value of SSGs. (Cf. also, [Juba, MSc. thesis, 2005].)

Market (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 $\alpha > 0$, $p \geq 0$, $g_i^l(\alpha p) = g_i^l(p)$. (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.)

Theorem

Computing (approximating) a price equilibrium for an exchange economy with demands given by $\{+, -, *, /, \max\}$*-circuits is* $\mathrm{FIXP}_{(a)}$*-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 (Uzawa (1962)): given Brouwer function
$F : \Delta_n \mapsto \Delta_n$, define total excess demand function $g : \Delta_n \mapsto \mathbb{R}^n$ by

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

$g(p)$ satisfies excess demand axioms. The price equilibria of $g(p)$ are the fixed points of $F(p)$. □

Conclusions: Some open problems

Open Problem 1: unique fixed points and Nash equilibria

Question 1: What is the complexity of computing the NE of a 2-player normal form game with a unique NE?

Conclusions: Some open problems

Open Problem 1: unique fixed points and Nash equilibria

Question 1: What is the complexity of computing the NE of a 2-player normal form game with a unique NE?

Conjecture 1: *At least as hard as computing the value of Condon's SSGs, and, more generally, at least as hard as computing the fixed point of any* linear-FIXP *function with a unique fixed point.*

Remark: This does not follow from PPAD-completeness results. Both [Chen-Deng'06]'s and [Daskalakis-Goldberg-Papadimitriou'06]'s reductions go through ϵ-NEs, so uniqueness is lost.

Remark: [R. Mehta, 2014]: new PPAD-completeness proof, reduces unique-linear-FIXP to 2-player games with a convex set of NEs.

Conclusions: Some open problems

Open Problem 1: unique fixed points and Nash equilibria

Question 1: What is the complexity of computing the NE of a 2-player normal form game with a unique NE?

Conjecture 1: *At least as hard as computing the value of Condon's SSGs, and, more generally, at least as hard as computing the fixed point of any* linear-FIXP *function with a unique fixed point.*

Remark: This does not follow from PPAD-completeness results. Both [Chen-Deng'06]'s and [Daskalakis-Goldberg-Papadimitriou'06]'s reductions go through ϵ-NEs, so uniqueness is lost.

Remark: [R. Mehta, 2014]: new PPAD-completeness proof, reduces unique-linear-FIXP to 2-player games with a convex set of NEs.
Remark: For 3 or more players ([E.-Yannakakis'07]) our FIXP-completeness reductions similarly almost preserve uniqueness (but not quite!): 1-to-1 correspondence between fixed points & player 1's mixed strategies in NEs.

Open problems

In many settings, one can establish the existence of a unique market (price) equilibrium.

One classic setting is a Arrow-Debreu exchange economy satisfying weak gross substitutes (WGS).

[Arrow-Block-Hurwicz'1959] showed these have a unique price equilibrium.

[Codenotti et. al., 2005] showed that one can compute a ϵ-"Almost"-equilibrium for a WGS economy in P-time.

Open Problem 1b: unique price equilibrium for WGS economies

Question 1b: Can one ϵ-Near-approximate the (unique) price equilibrium for a WGS exchange economy in P-time?

Open problems

In many settings, one can establish the existence of a unique market (price) equilibrium.

One classic setting is a Arrow-Debreu exchange economy satisfying weak gross substitutes (WGS).

[Arrow-Block-Hurwicz'1959] showed these have a unique price equilibrium.

[Codenotti et. al., 2005] showed that one can compute a ϵ-"Almost"-equilibrium for a WGS economy in P-time.

Open Problem 1b: unique price equilibrium for WGS economies

Question 1b: Can one ϵ-Near-approximate the (unique) price equilibrium for a WGS exchange economy in P-time?

(Such a Near approximation might already be PosSLP-hard.
We don't know.)

Open problems

Open problem 2: complexity of PosSLP and unit-cost-P-time

Question 2: Can we obtain any better upper bounds for PosSLP??

Open problems

Open problem 2: complexity of PosSLP and unit-cost-P-time

Question 2: Can we obtain any better upper bounds for PosSLP??

Here is one basic (and probably bad) idea:
Given a $\{+,-,*\}$-circuit, C, guess a monotone $\{+,*\}$-circuit, C', as a "Witness of positivity", and verify that $C - C' = 0$ in **co-RP**.
(Checking equality to 0 is PIT-equivalent ([ABKM'06]).)

Conjecture 2: This does not work. In other words (surely!) $\exists$ a family of positive integers, $\langle A_n \rangle_{n \in \mathbb{N}}$, such that A_n has encoding size $O(n)$ as a $\{+,-,*\}$-circuit, but requires size $2^{\Omega(n)}$ monotone $\{+,*\}$-circuits.

Current state of knowledge is abismal. (We don't even know super-linear lower bounds.)
This is despite the fact that [Valiant'79] proved an exponential lower bound for monotone polynomials.
(This doesn't imply a lower bound in the integer setting.)

Open problems

Open problem 2: complexity of PosSLP and unit-cost-P-time

Question 2: Can we obtain any better upper bounds for PosSLP??

Open problems

Open problem 2: complexity of PosSLP and unit-cost-P-time

Question 2: Can we obtain any better upper bounds for PosSLP??

Definition: call a circuit, C', quasi-monotone if it consists of some squared $\{+,*,-\}$-subcircuits, which are inputs to a monotone $\{+,*\}$-circuit.

Same idea: Given a $\{+,-,*\}$-circuit, C, guess a pair of quasi-monotone circuits C', C'' as a "witness of positivity" for C, & verify the equality $((C'+1)*C - C'') = 0$ in **co-RP**. Checking equality is in fact PIT-equivalent ([Allender,et. al.'06]).

Conjecture 3: This works. There always exists poly-sized witness quasi-monotone circuits. More formally: *For every positive integer expressed by a $\{+,-,*\}$-circuit, C, there are quasi-montone circuits C' and C'' of size poly$(|C|)$, such that val$((C'+1)*C) =$ val(C'').*

Remark: This would imply that PosSLP $\in$ **MA**, and *if* we also knew that PIT $\in$ **P**, then it would further imply PosSLP $\in$ **NP**.

Open problems

Open problem 3: how many variables are needed for FIXP??

Recall: [Mehta'14]: linear-FIXP = 2-variable-linear-FIXP

Question 3: Is FIXP = 3-variable-FIXP ?? (or k-variable-FIXP, for any fixed k?)

Note: If boundedly many variables suffice, it requires k-variable circuits, not formulas: fixed points of k-variable formulas, for fixed k, can be approximated in P-time (using decision procedures for the existential theory of reals).

Open problems

Open problem 4: complexity of FIXP and $\mathbf{NP}_{\mathbb{R}}$

Question 4: Can we get any better upper bound than PSPACE for FIXP_a?

Open problems

Open problem 4: complexity of FIXP and $\mathbf{NP}_{\mathbb{R}}$

Question 4: Can we get any better upper bound than PSPACE for FIXP$_a$?

Conjecture 4 (wildly optimistic wishful thinking):

The existential theory of reals is decidable in $\mathbf{NP}^{PosSLP}$.

Remark: Would imply the (discrete) BSS class $\mathbf{NP}_{\mathbb{R}}$ is equal to $\mathbf{NP}^{PosSLP}$.

Would also imply that FIXP$_a \subseteq \mathbf{FNP}^{PosSLP}$.

Conjectures 3 & 4 together would imply (discrete) $\mathbf{NP}_{\mathbb{R}} \subseteq \mathbf{PH}$.