Algorithms for some Infinite-State MDPs and Stochastic Games (invited tutorial)

Kousha Etessami

University of Edinburgh

LICS 2017
Reykjavik, Iceland, June 2017

Overview

- Over last $\sim$ 15 years, there's been a substantial body of research in verification & TCS on algorithms & complexity of analyzing & model checking infinite-state (but finitely-presented) Markov chains, Markov decision processes (MDPs), and stochastic games.

Overview

- Over last $\sim$ 15 years, there's been a substantial body of research in verification & TCS on algorithms & complexity of analyzing & model checking infinite-state (but finitely-presented) Markov chains, Markov decision processes (MDPs), and stochastic games.
- Many of these models add probabilistic/control/game behavior to some classic automata-theoretic or process-algebraic model (e.g., context-free grammars, pushdown automata, one-counter automata, BPPs, BPAs).

Overview

- Over last ~ 15 years, there's been a substantial body of research in verification & TCS on algorithms & complexity of analyzing & model checking infinite-state (but finitely-presented) Markov chains, Markov decision processes (MDPs), and stochastic games.
- Many of these models add probabilistic/control/game behavior to some classic automata-theoretic or process-algebraic model (e.g., context-free grammars, pushdown automata, one-counter automata, BPPs, BPAs).
- These models are also intimately related to some classic stochastic processes.

Overview

- Over last ~ 15 years, there's been a substantial body of research in verification & TCS on algorithms & complexity of analyzing & model checking infinite-state (but finitely-presented) Markov chains, Markov decision processes (MDPs), and stochastic games.
- Many of these models add probabilistic/control/game behavior to some classic automata-theoretic or process-algebraic model (e.g., context-free grammars, pushdown automata, one-counter automata, BPPs, BPAs).
- These models are also intimately related to some classic stochastic processes.
- In this tutorial I hope to give you a flavor of this research.

  (I can't be comprehensive: it is by now a very rich body of work.)

- I will focus mainly on a series of results I have been involved with, on algorithms & complexity of analyzing the following models:
  - **Multi-type Branching Processes** (≈ PCFGs ≈ pBPPs/pBPA), and their generalization to:
    Branching MDPs and **Branching Stochastic Games**.
  - **One-counter Markov Chains**(≈Quasi-Birth Death processes(QBDs)), and **one-counter MDPs/stochastic games**.
  - **Recursive Markov Chains** (≈ prob. Pushdown Systems (pPDSs)), and **Recursive MDPs/stochastic games**.
- A key aspect of some of these results: new algorithms & complexity bounds for computing the **least fixed point** solution for **monotone/probabilistic systems of (min/max)-polynomial equations**.

  Such equations arise for various stochastic models, MDPs (as their Bellman optimality equations), and stochastic games.

Multi-type Branching Processes (BPs) (Kolmogorov,1940s)

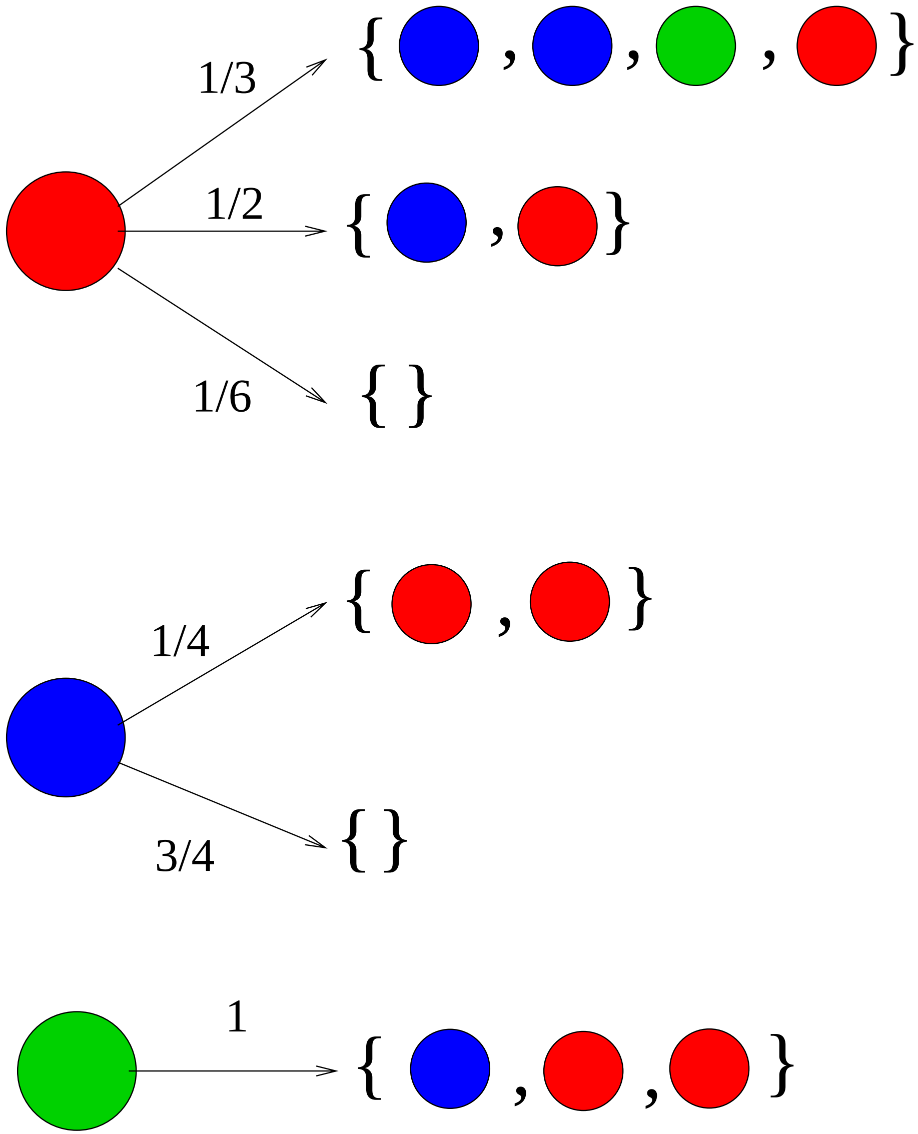

Multi-type Branching Processes (BPs) (Kolmogorov,1940s)

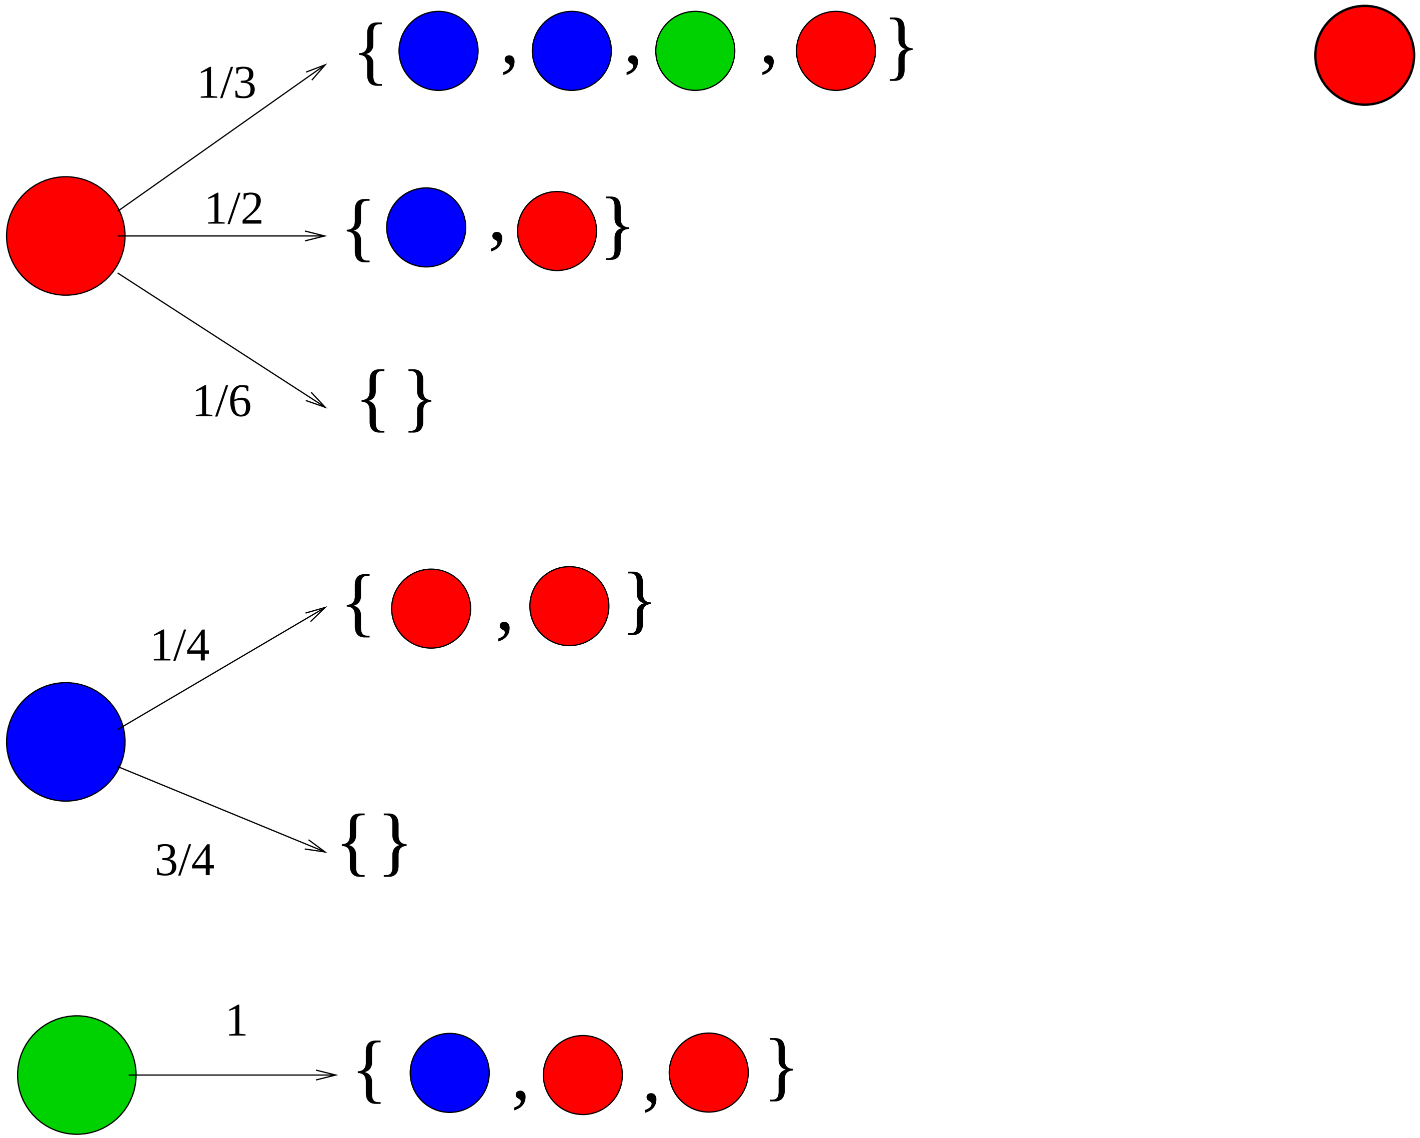

Multi-type Branching Processes (BPs) (Kolmogorov,1940s)

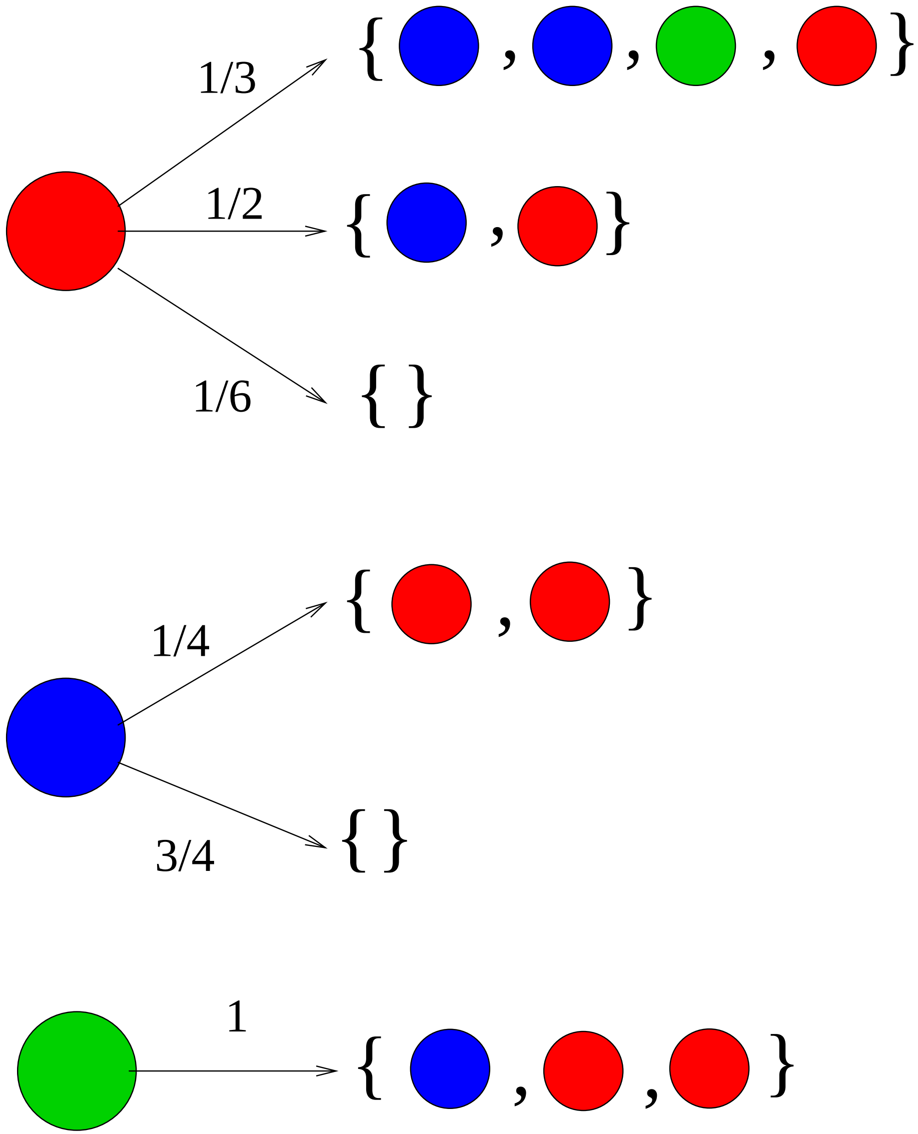

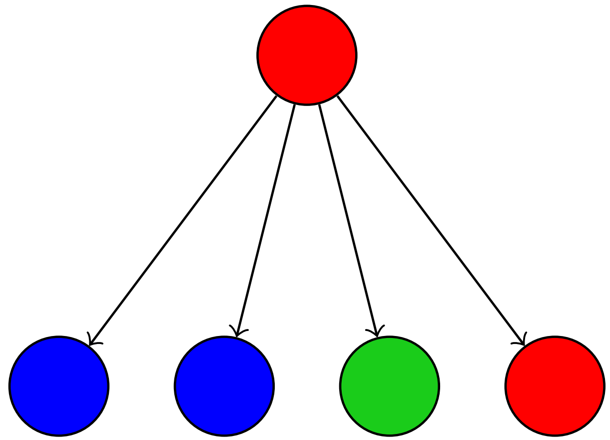

Multi-type Branching Processes (BPs) (Kolmogorov,1940s)

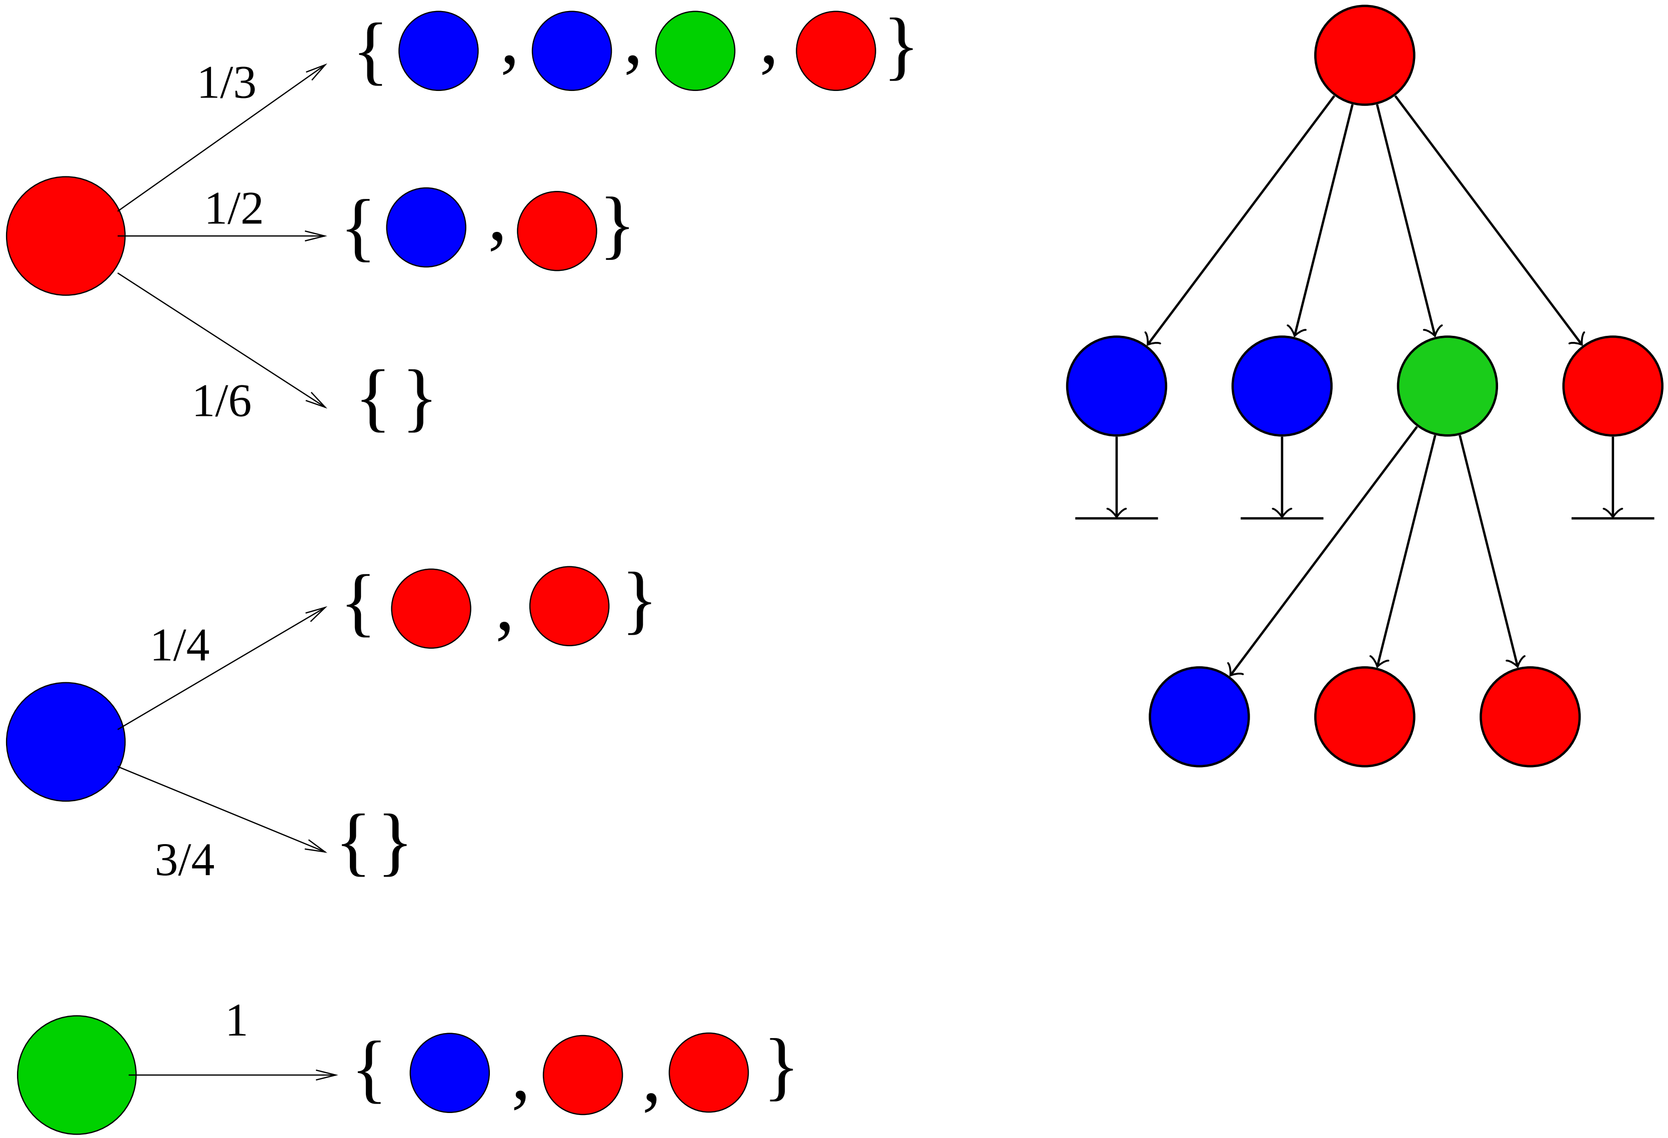

Multi-type Branching Processes (BPs) (Kolmogorov,1940s)

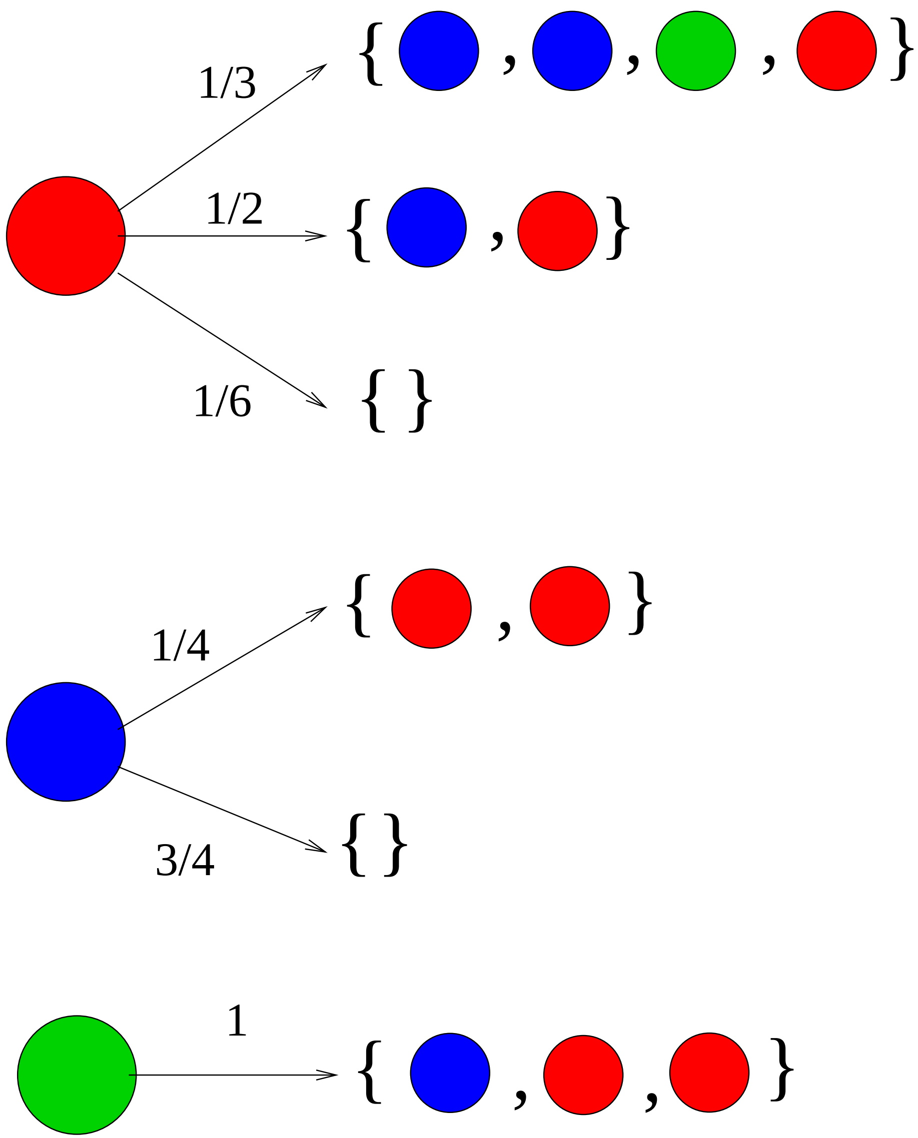

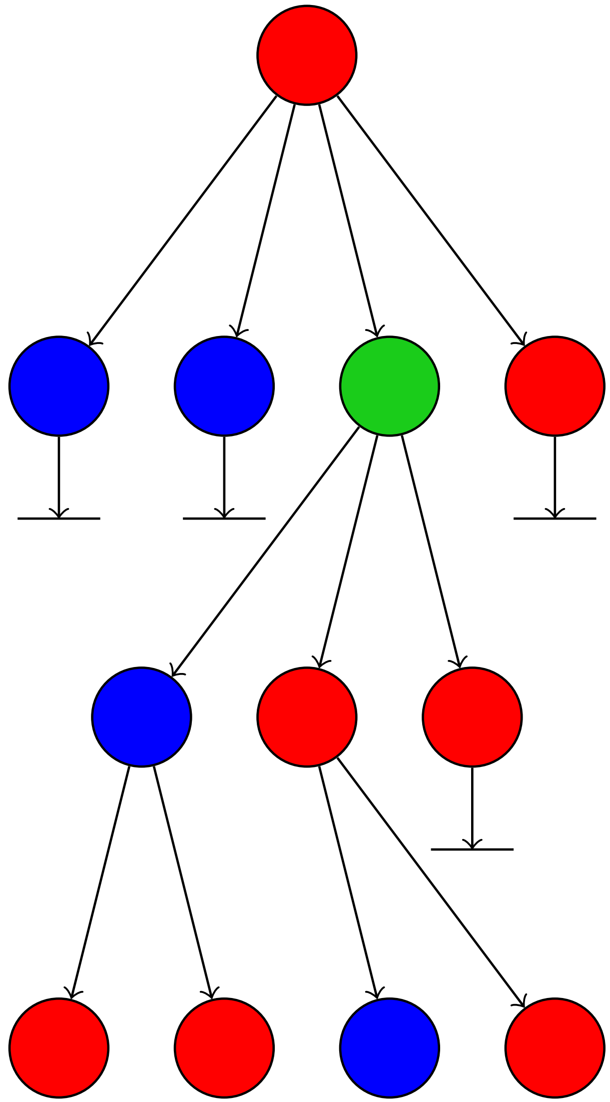

Multi-type Branching Processes (BPs) (Kolmogorov,1940s)

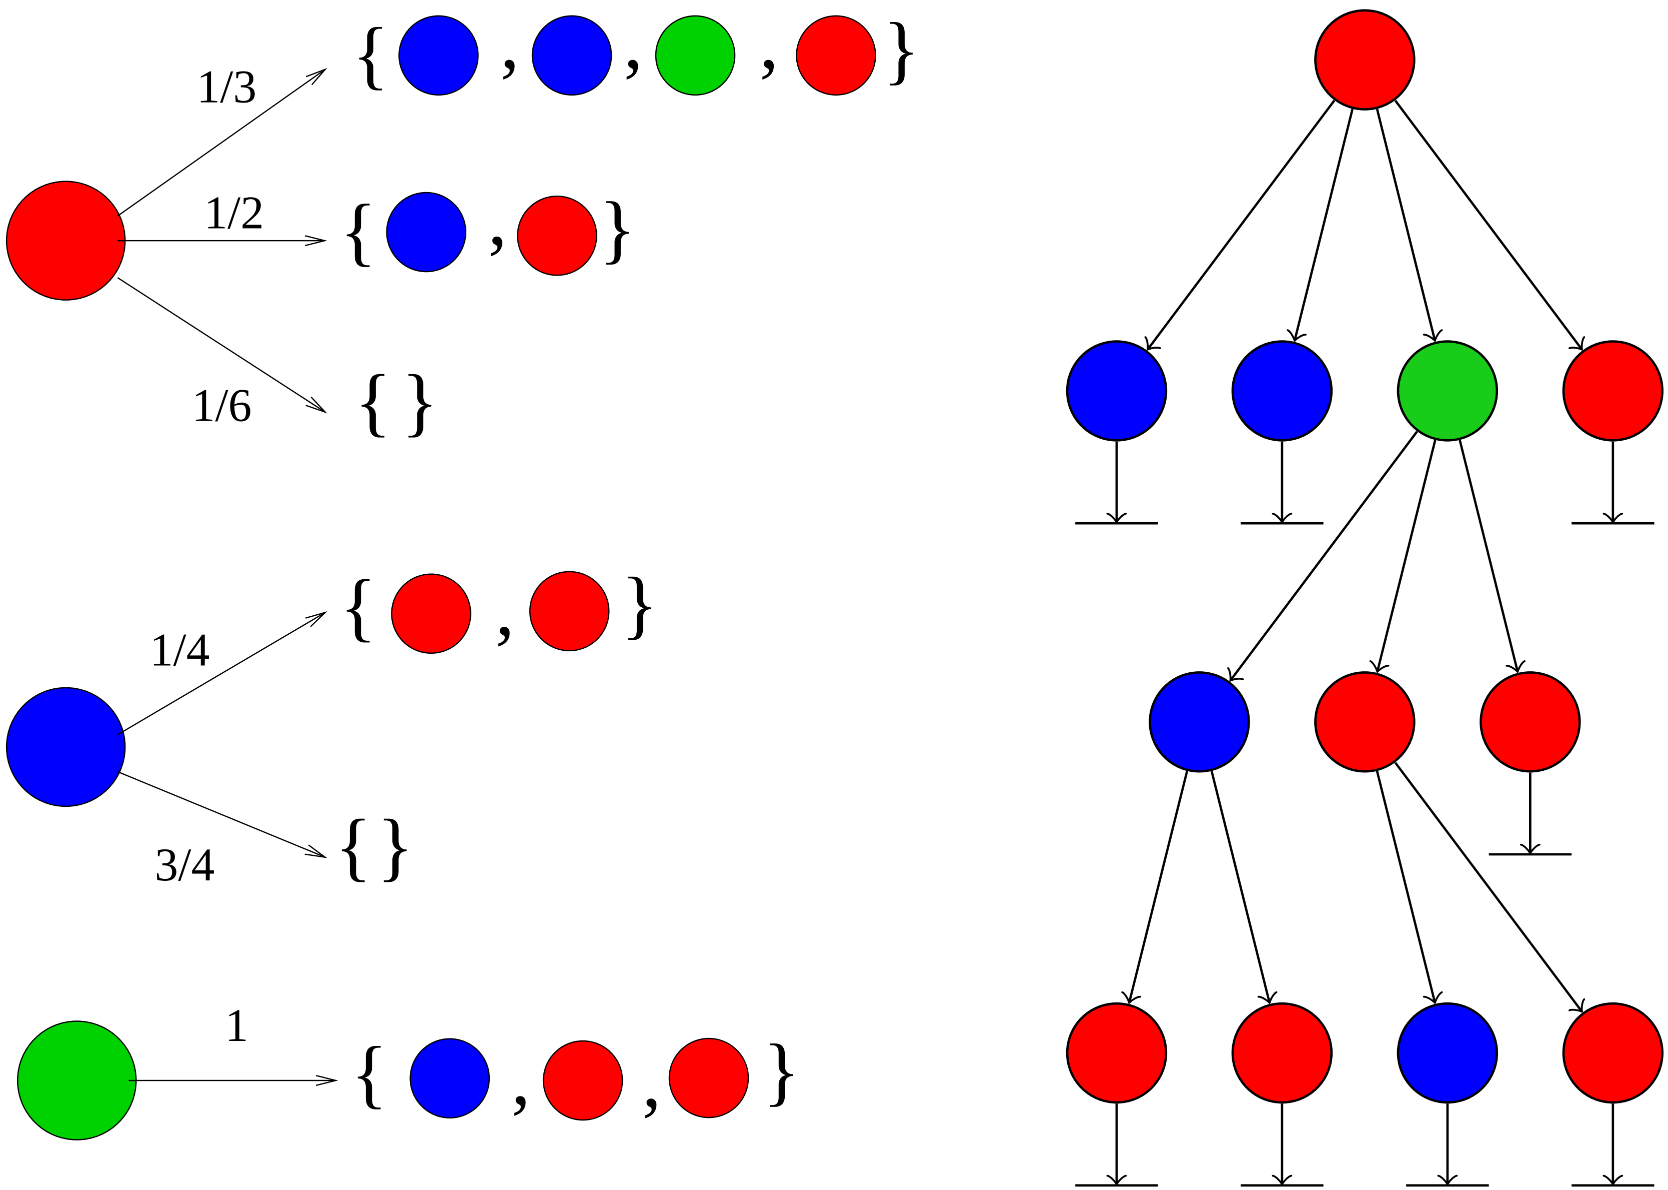

BPs are classic stochastic processes, studied for decades in probability theory, with many applications, eg.:
population biology, nuclear chain reactions, cancer tumor models, random graph theory, . . .

BPs are also "intimately related" to:

- probabilistic (stochastic) context-free grammars (PCFGs)
- probabilistic BPPs, and probablistic BPAs
- 1-exit recursive Markov chains
- 1-state probabilistic pushdown systems.

Nevertheless, some basic algorithmic questions about BPs remained open until recent years.

Probabilistic Context-Free Grammars (PCFGs)

$$R \xrightarrow{1/3} aBBcGaabR$$

$$R \xrightarrow{1/2} bcBbR$$

$$R \xrightarrow{1/6} \epsilon$$

$$B \xrightarrow{1/4} bbRRc$$

$$B \xrightarrow{3/4} a$$

$$G \xrightarrow{1} aBcRRb$$

Probabilistic Context-Free Grammars (PCFGs)

$$R \xrightarrow{1/3} aBBcGaabR$$

$$R \xrightarrow{1/2} bcBbR$$

$$R \xrightarrow{1/6} \epsilon$$

$$B \xrightarrow{1/4} bbRRc$$

$$B \xrightarrow{3/4} a$$

$$G \xrightarrow{1} aBcRRb$$

Question

What is the probability of termination, i.e., eventually generating a finite string, starting with non-terminal, R ?

(These probabilities are also known as the partition function of the PCFG.)

Multi-type Branching Processes (Kolmogorov,1940s)

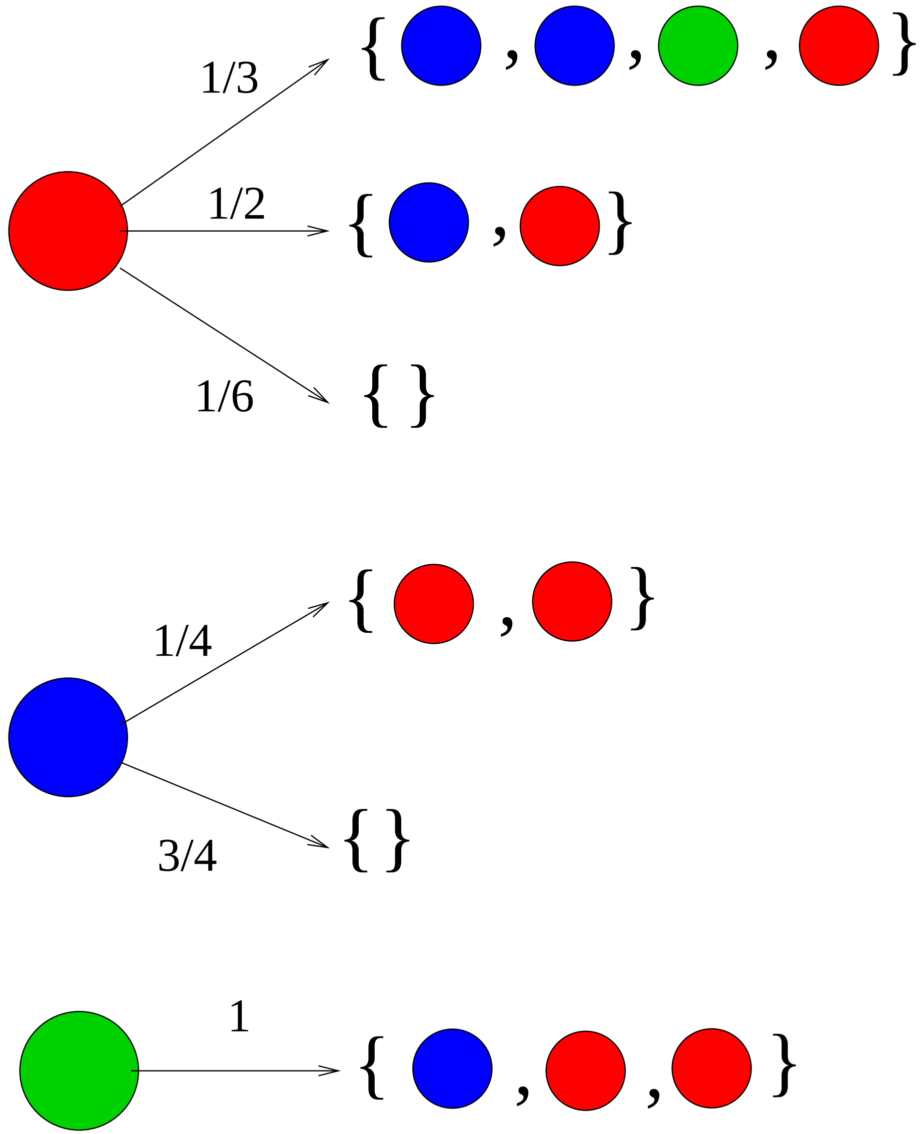

Question: What is the probability of eventual extinction, starting with one 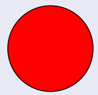 ?

Multi-type Branching Processes (Kolmogorov,1940s)

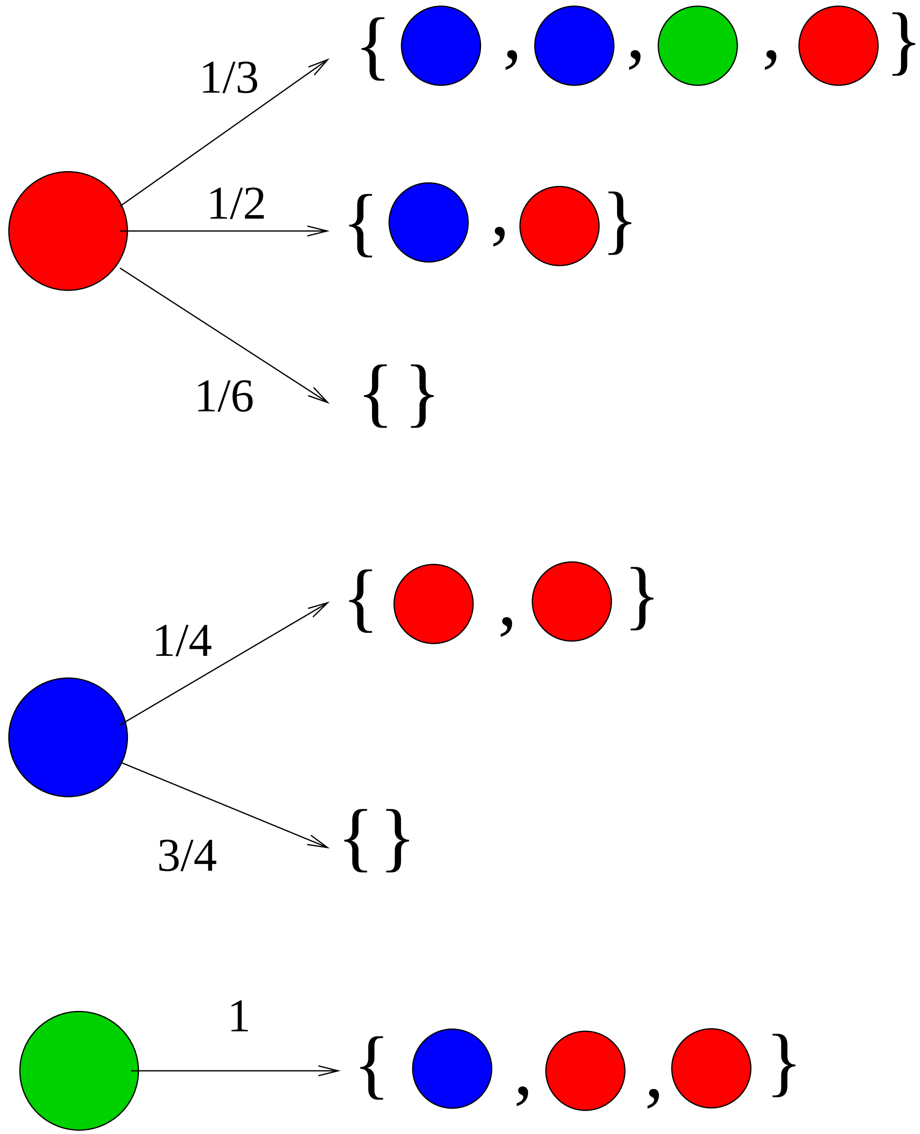

Question: What is the probability of eventual extinction, starting with one 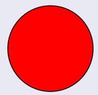 ?

$$x_R \quad =$$

Multi-type Branching Processes (Kolmogorov,1940s)

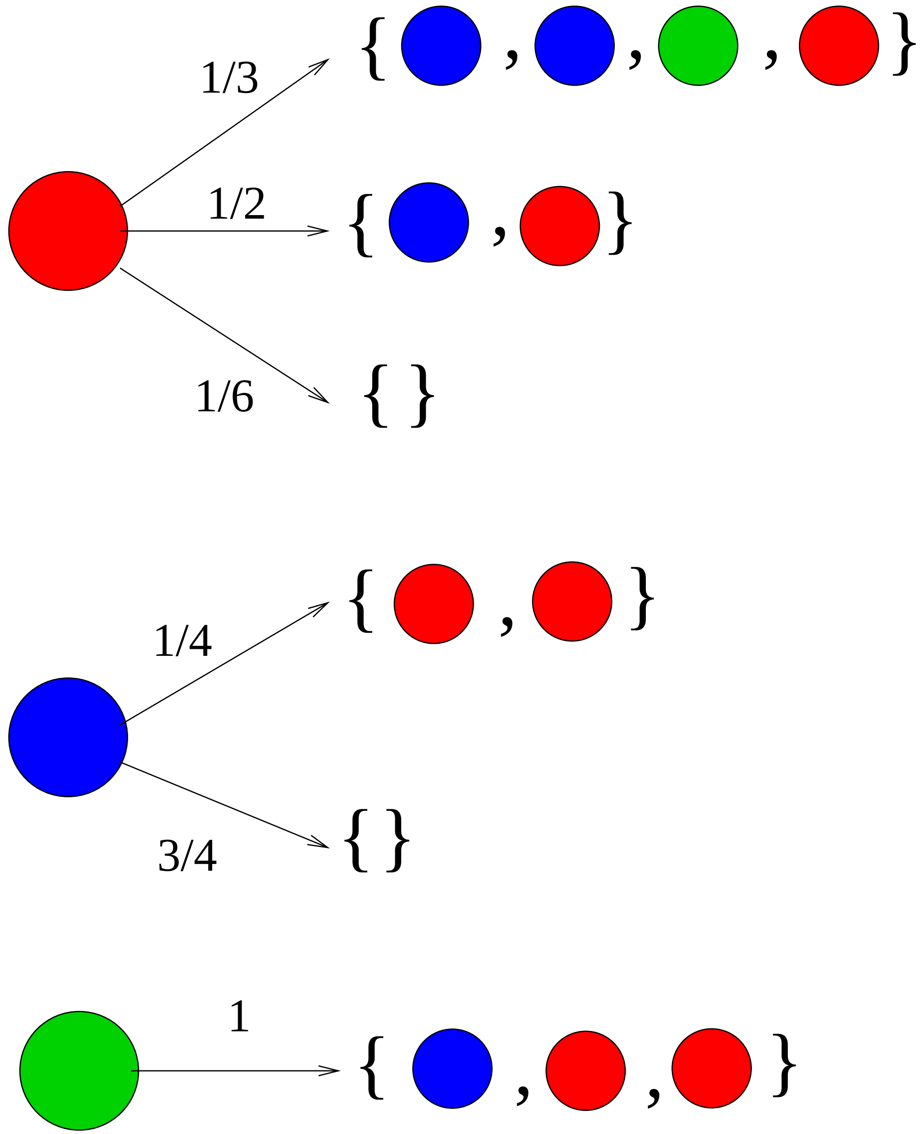

Question: What is the probability of eventual extinction, starting with one 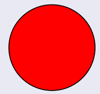 ?

$$x_R = \frac{1}{3} x_B^2 x_G x_R + \frac{1}{2} x_B x_R + \frac{1}{6}$$

Multi-type Branching Processes (Kolmogorov,1940s)

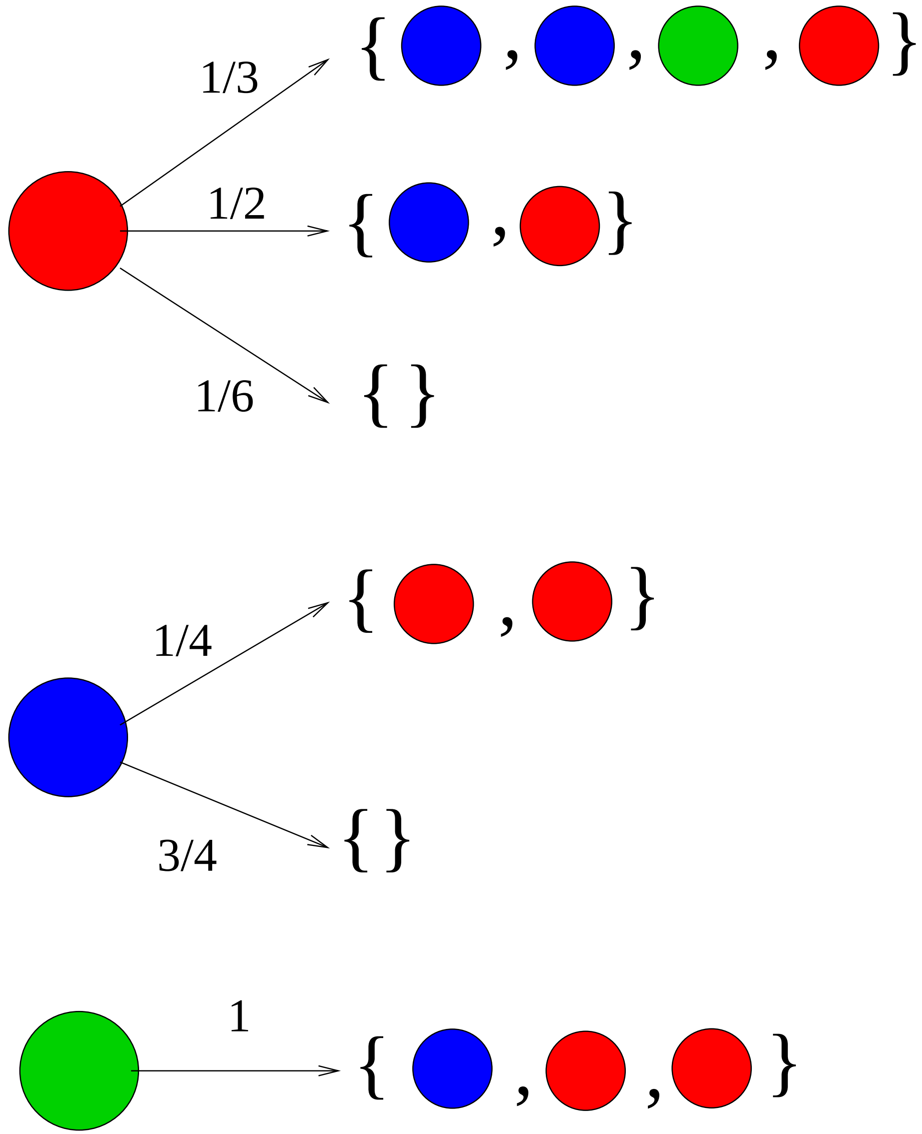

Question: What is the probability of eventual extinction, starting with one 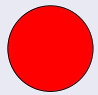 ?

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

Multi-type Branching Processes (Kolmogorov,1940s)

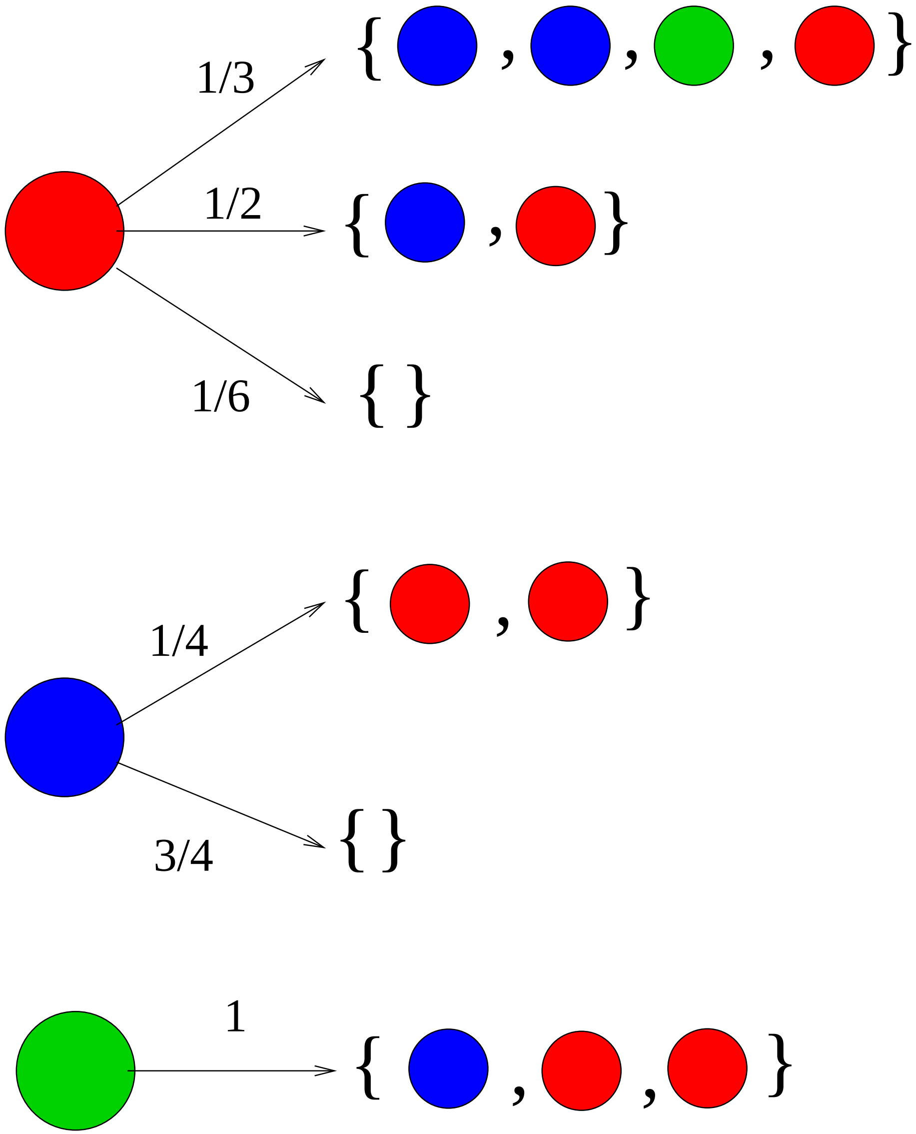

Question: What is the probability of eventual extinction, starting with one ?

$$x_R = \frac{1}{3}x_B^2 x_G x_R + \frac{1}{2}x_B x_R + \frac{1}{6}$$

$$x_B = \frac{1}{4}x_R^2 + \frac{3}{4}$$

$$x_G = x_B x_R^2$$

We get nonlinear fixed point equations:

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

Multi-type Branching Processes (Kolmogorov,1940s)

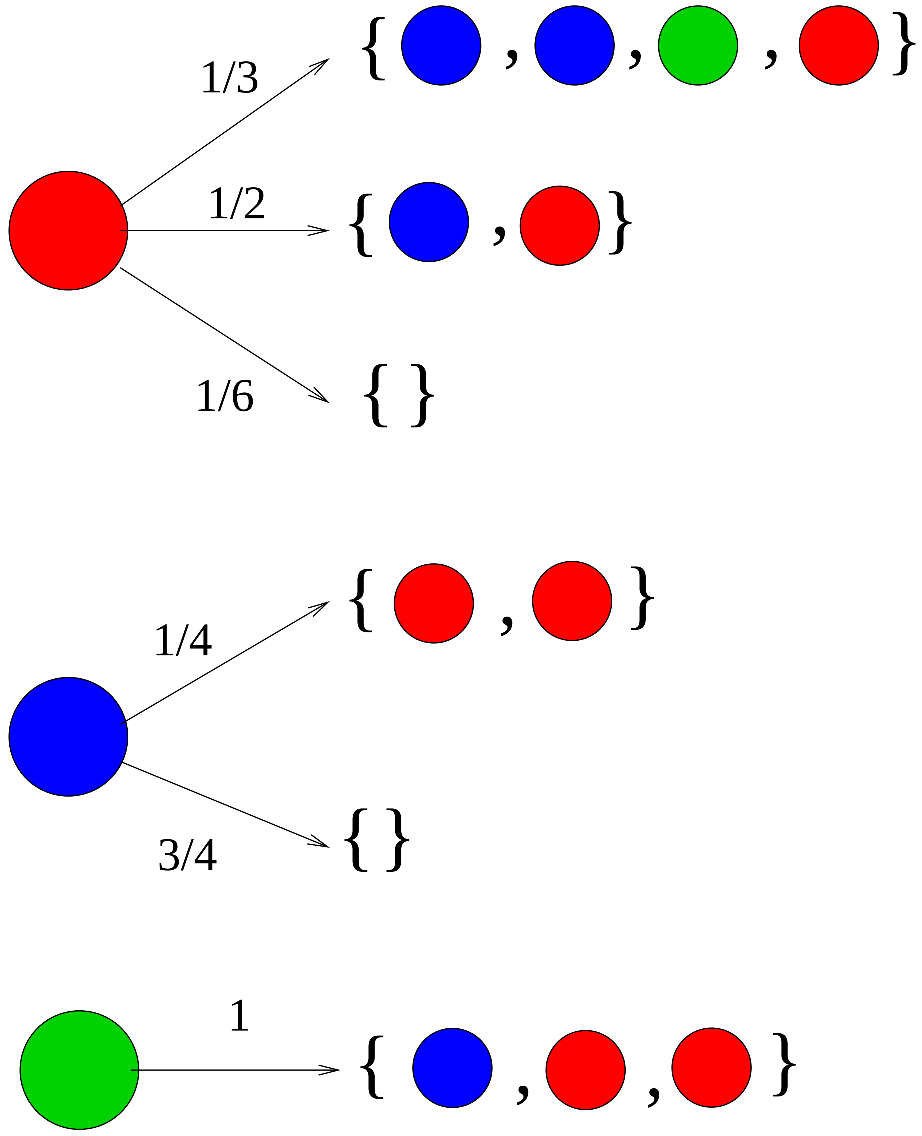

Question: What is the probability of eventual extinction, starting with one 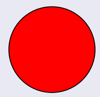 ?

$$x_R = \frac{1}{3}x_B^2 x_G x_R + \frac{1}{2}x_B x_R + \frac{1}{6}$$

$$x_B = \frac{1}{4}x_R^2 + \frac{3}{4}$$

$$x_G = x_B x_R^2$$

We get nonlinear fixed point equations:

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

Fact

The 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

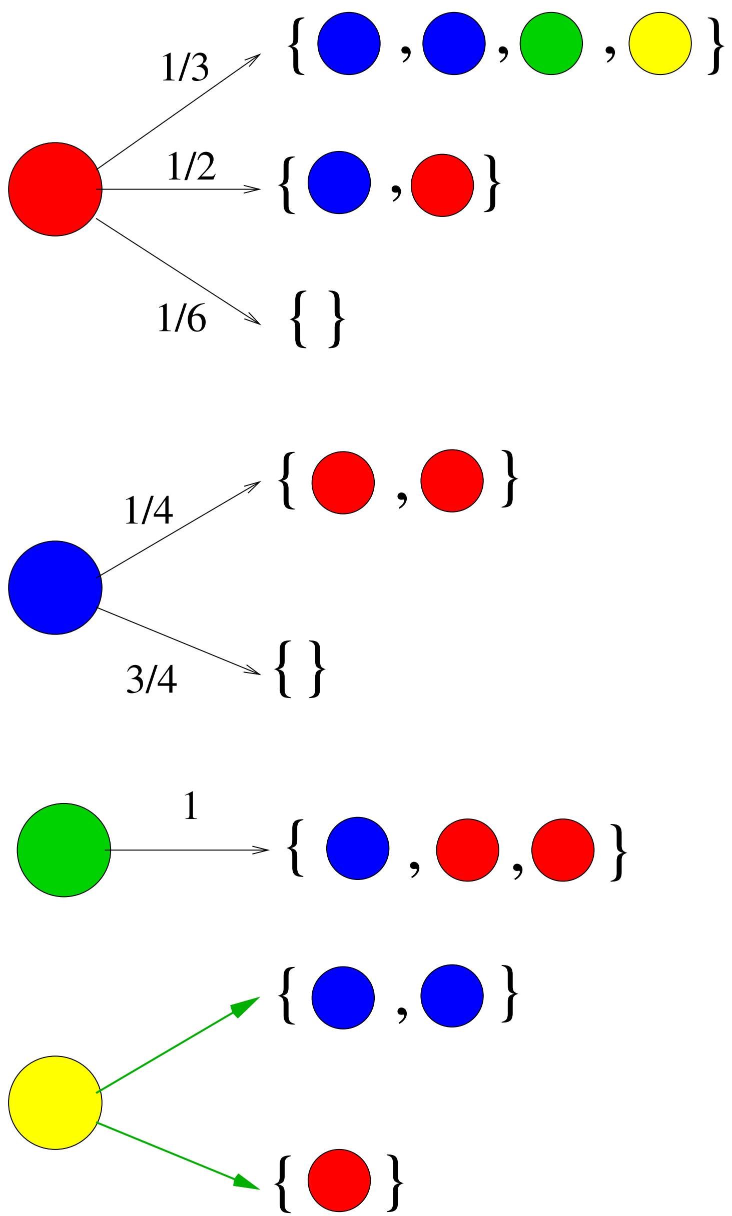

Branching Markov Decision Processes

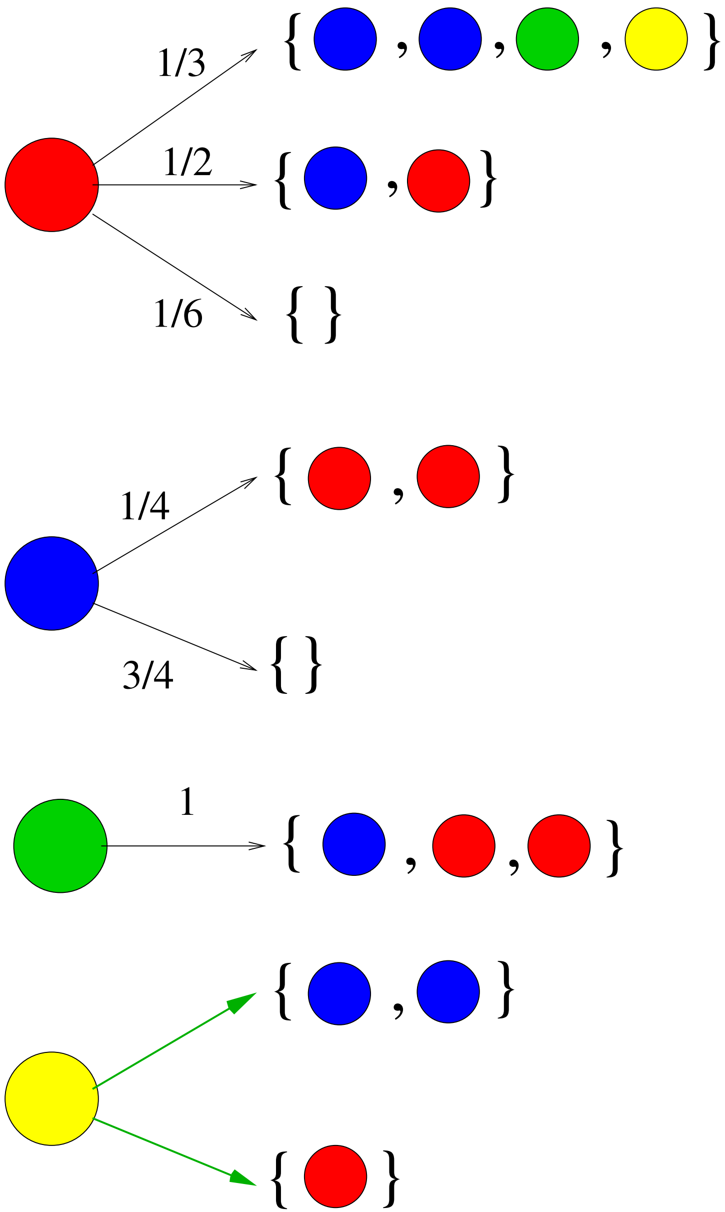

Question

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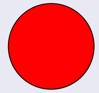 ?

Branching Markov Decision Processes

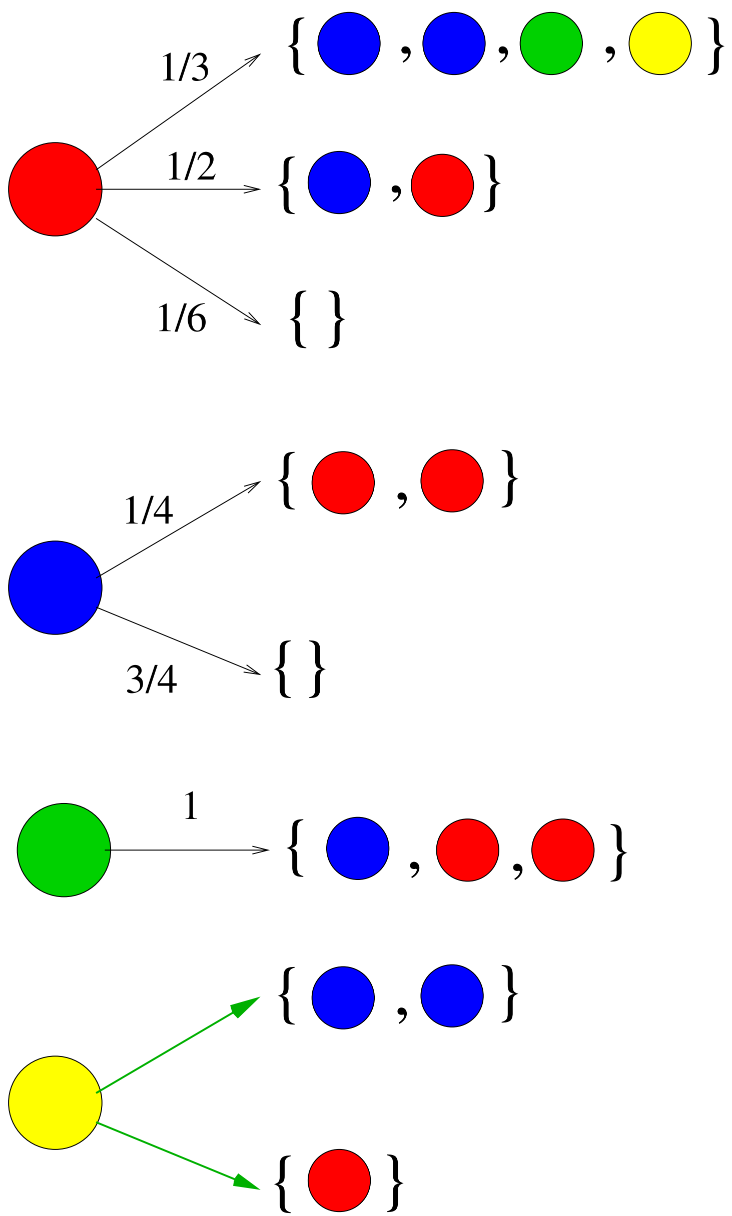

Question

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

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

Branching Markov Decision Processes

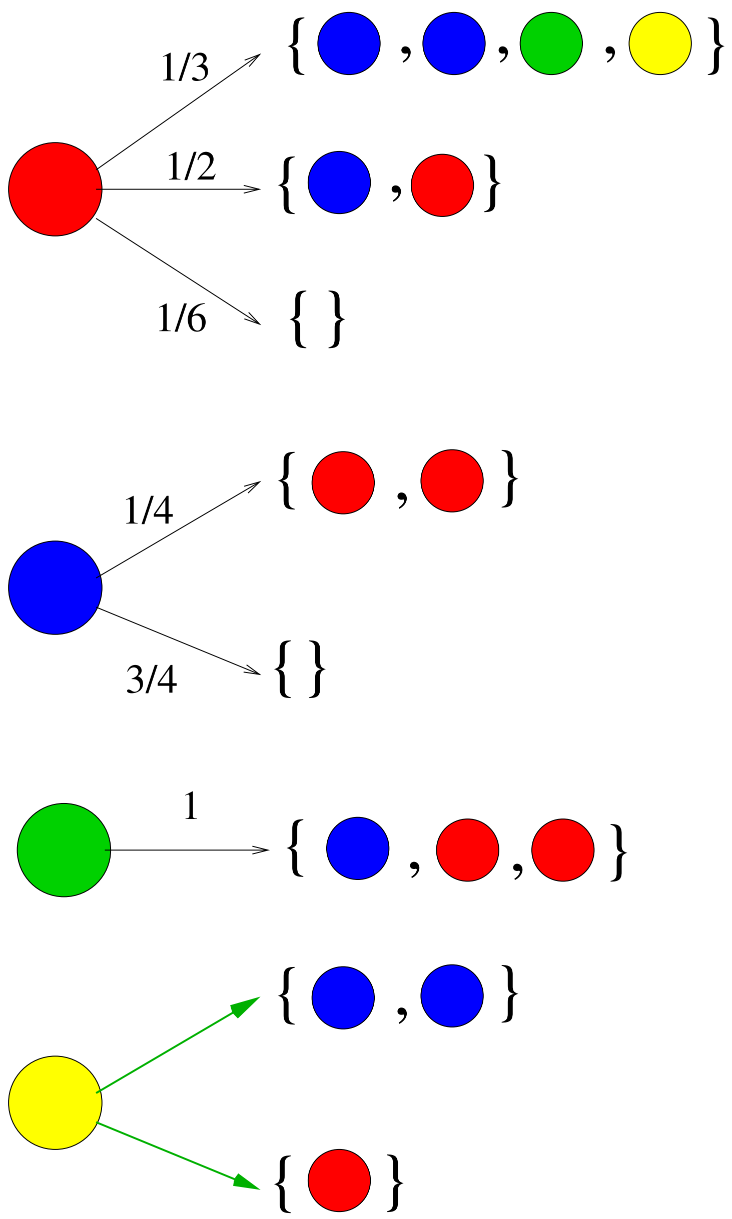

Question

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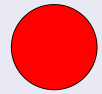 ?

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

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

Theorem [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

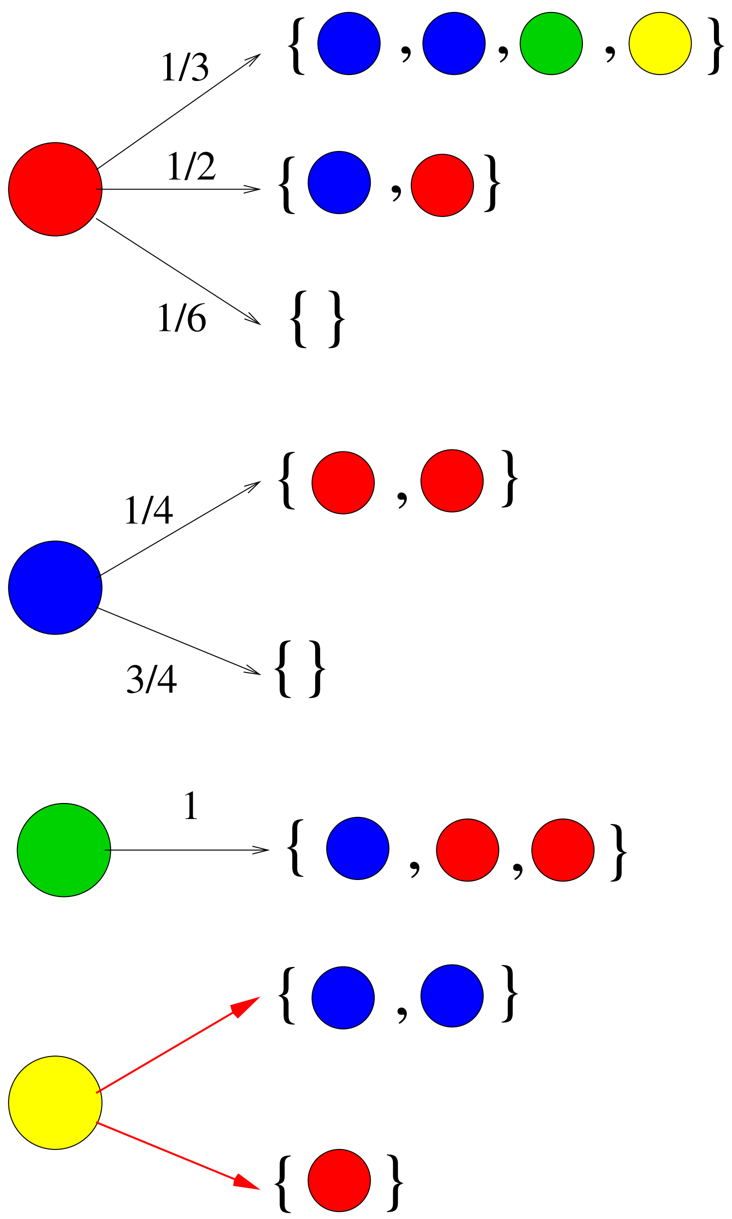

Question

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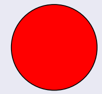 ?

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

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

Theorem [E.-Yannakakis'05]

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

Branching Simple Stochastic Games

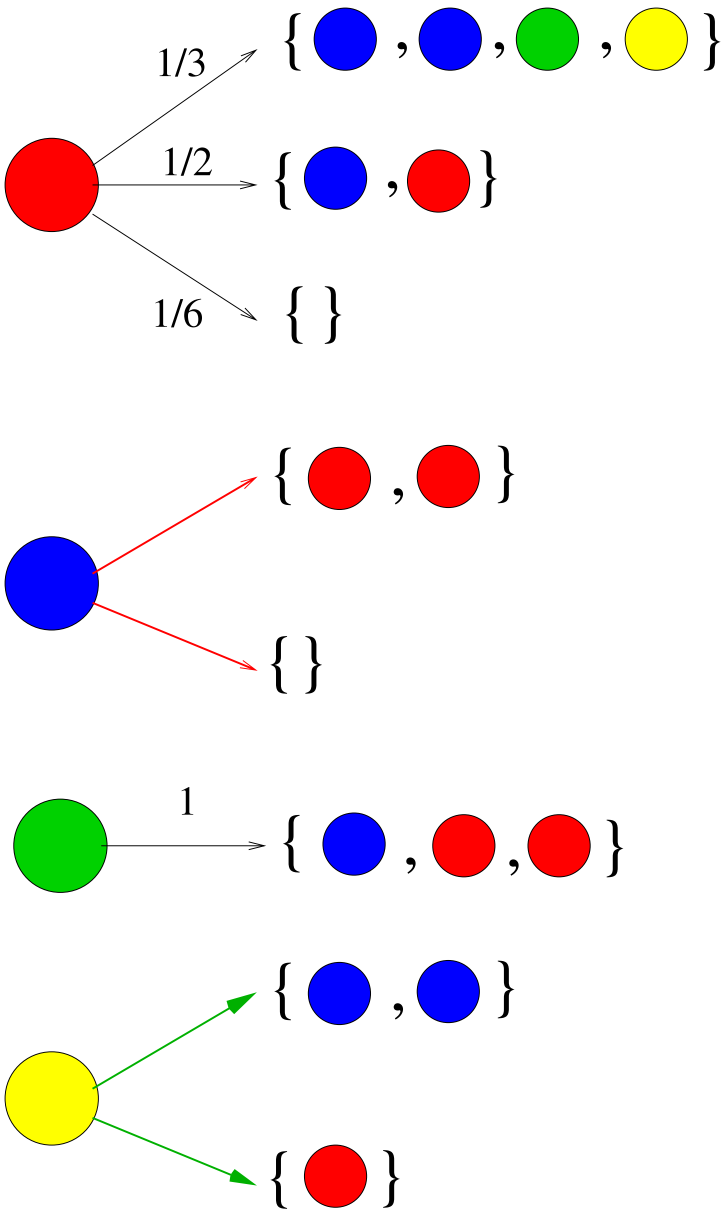

Question

What is the value of extinction, starting with one 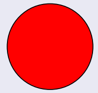 ?

Branching Simple Stochastic Games

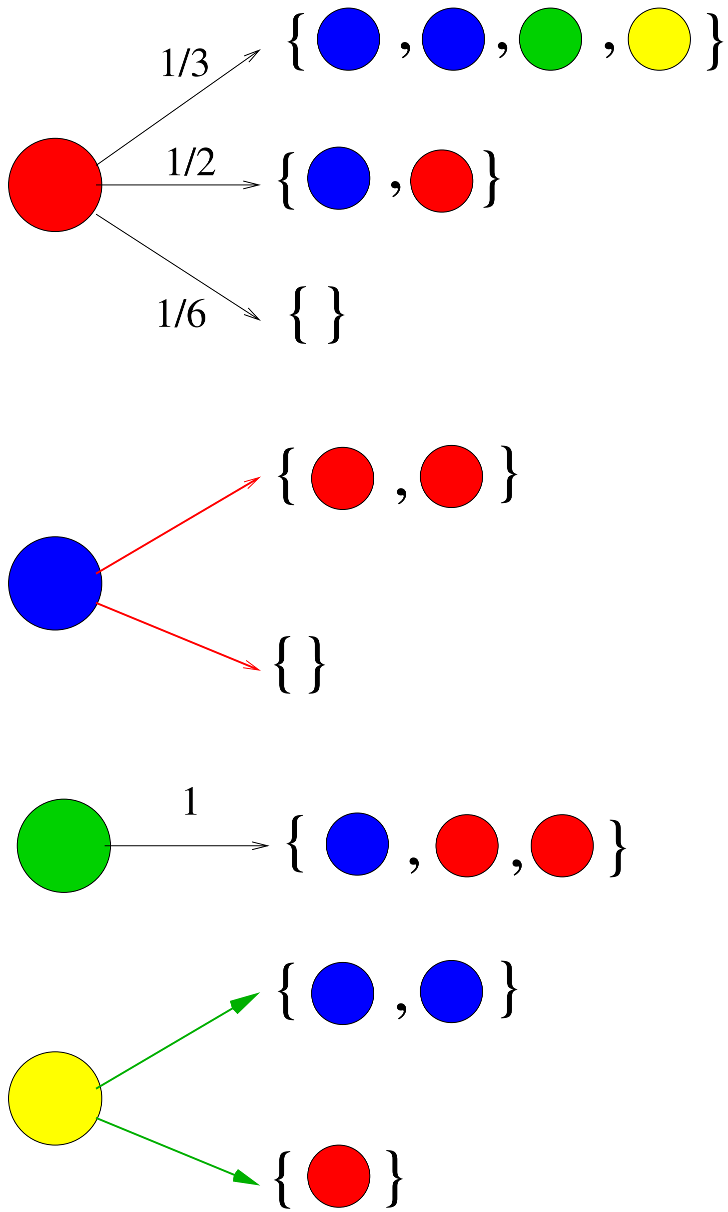

Question

What is the value of extinction, starting with one 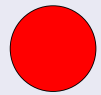 ?

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

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

Theorem [E.-Yannakakis'05]

The extinction values are the LFP, $\mathbf{q}^* \in [0,1]^3$ of $\bar{\mathbf{x}} = P(\bar{\mathbf{x}})$.

Branching Concurrent Stochastic Games

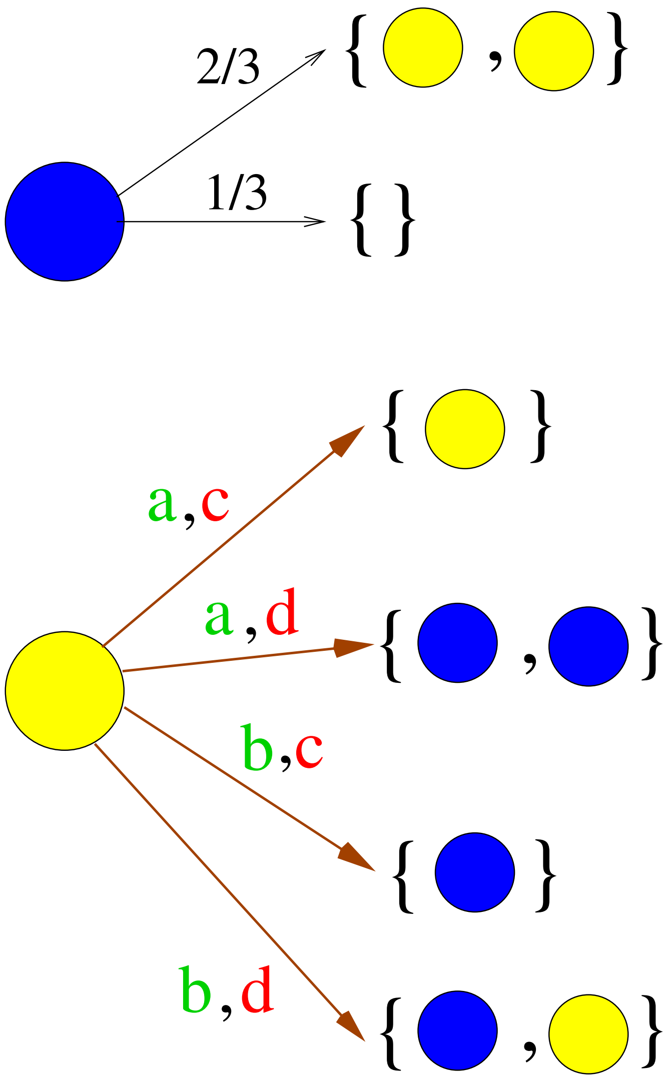

Branching Concurrent Stochastic Games

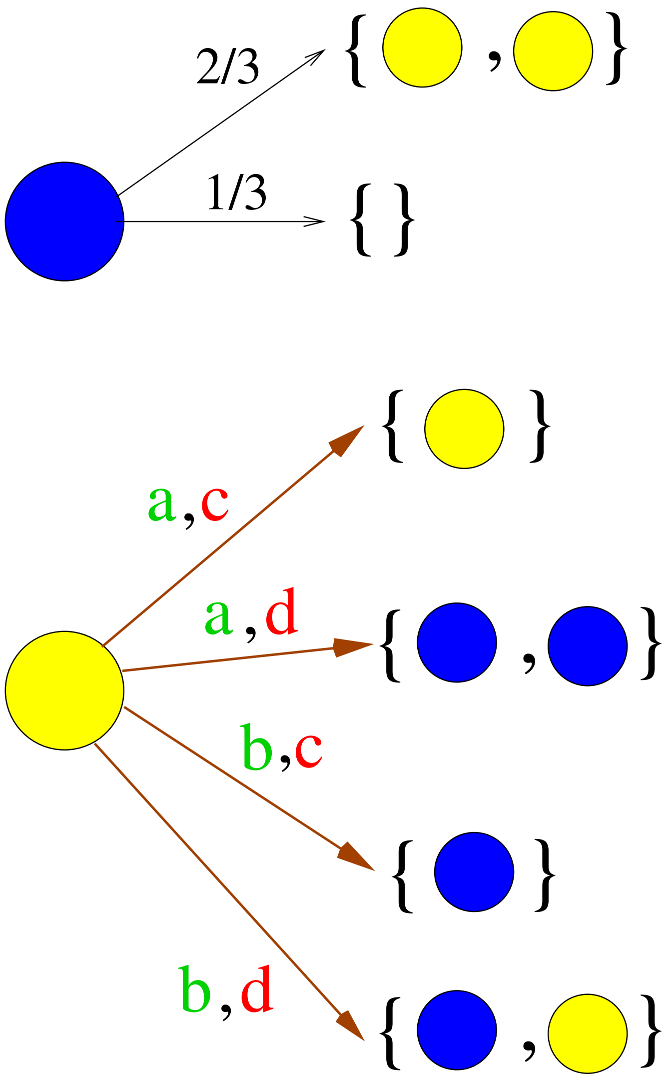

Question

What is the value of extinction, starting with one 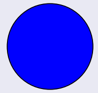 ?

Branching Concurrent Stochastic Games

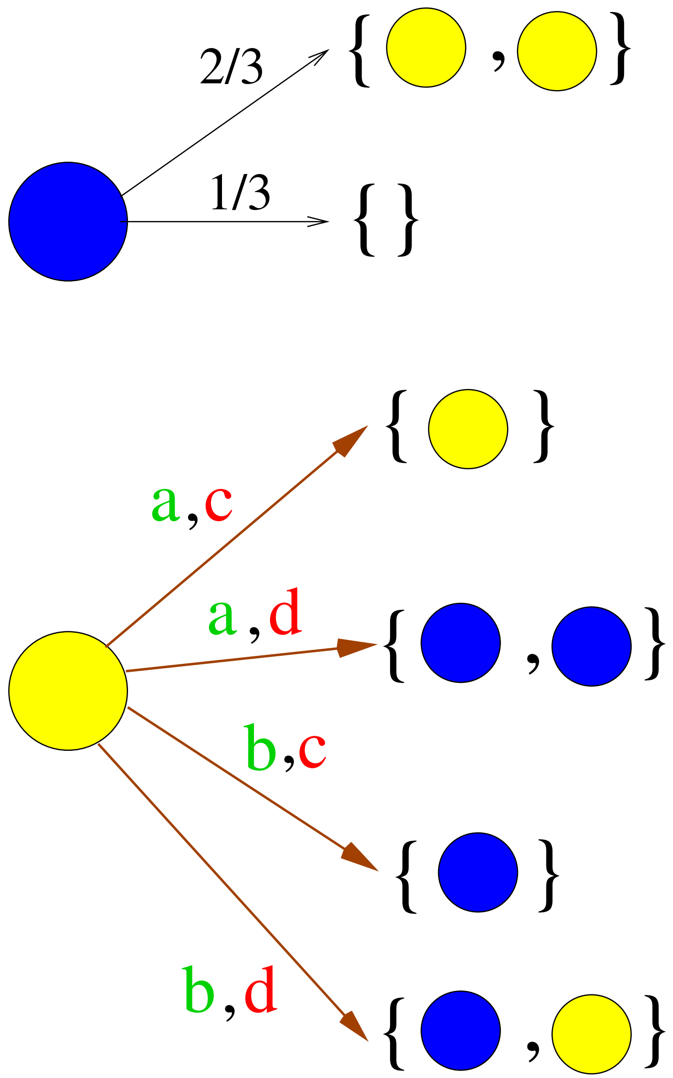

Question

What is the value of extinction, starting with one (blue) ?

$$x_B = \frac{2}{3}x_Y^2 + \frac{1}{3}$$

$$x_Y =$$

Branching Concurrent Stochastic Games

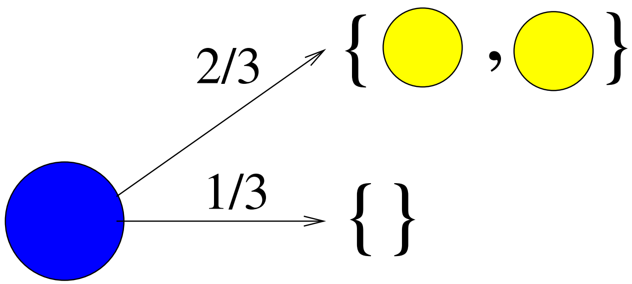

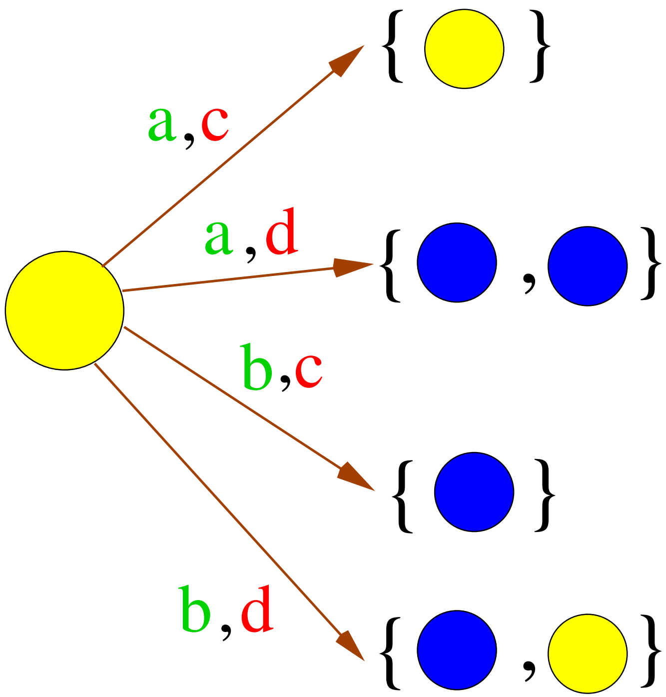

Question

What is the value of extinction, starting with one 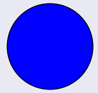 ?

$$x_B = \frac{2}{3}x_Y^2 + \frac{1}{3}$$

$$x_Y = \textbf{Val}\left(\begin{bmatrix} x_Y & x_B^2 \\ x_B & x_B x_Y \end{bmatrix}\right)$$

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

Theorem [E.-Yannakakis'06]

The extinction values are the LFP, $\mathbf{q}^* \in [0,1]^2$ of $\bar{\mathbf{x}} = P(\bar{\mathbf{x}})$.

BMDPs again: this time optimizing exptected tree size

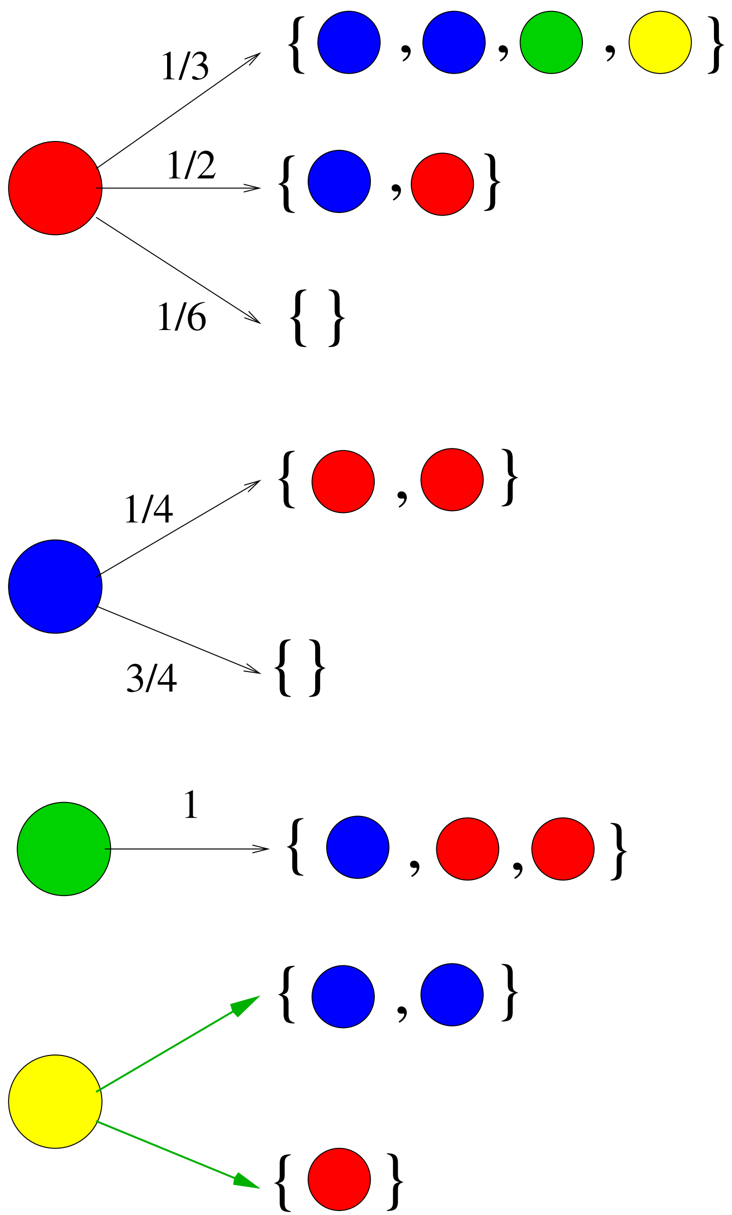

BMDPs again: this time optimizing expected tree size

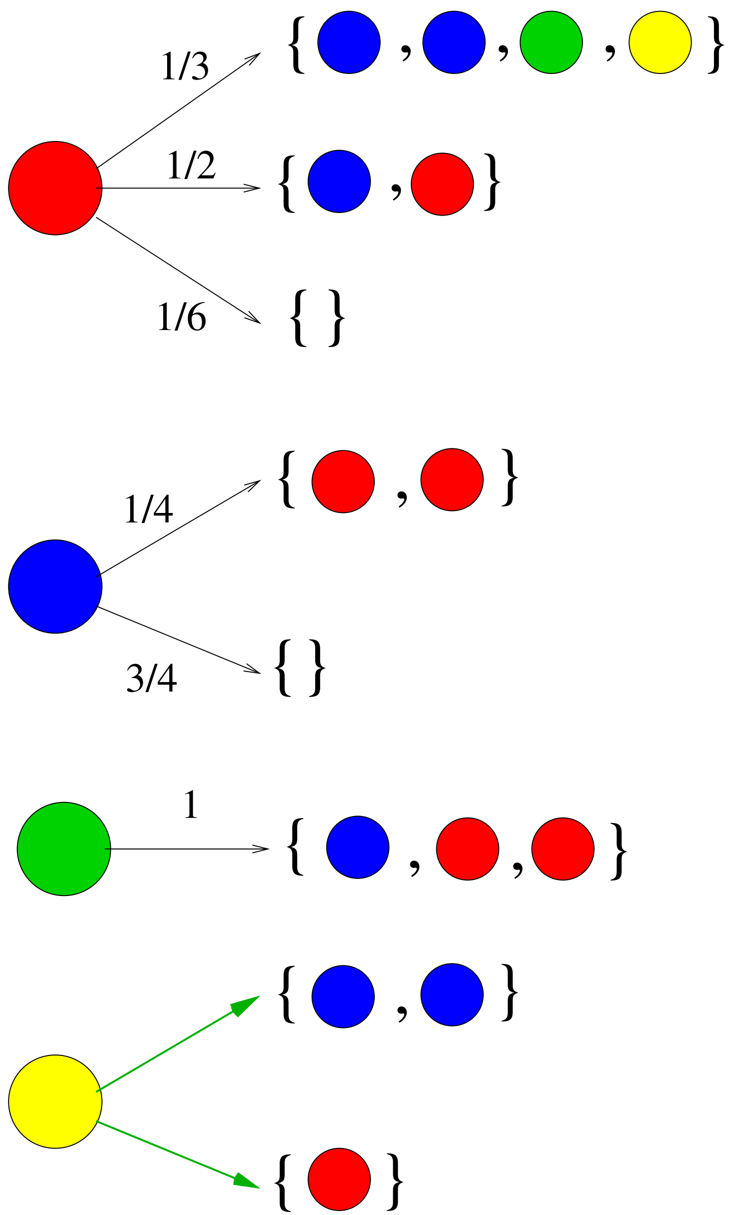

Question

What is the maximum expected size of the tree, starting with one ?

BMDPs again: this time optimizing expected tree size

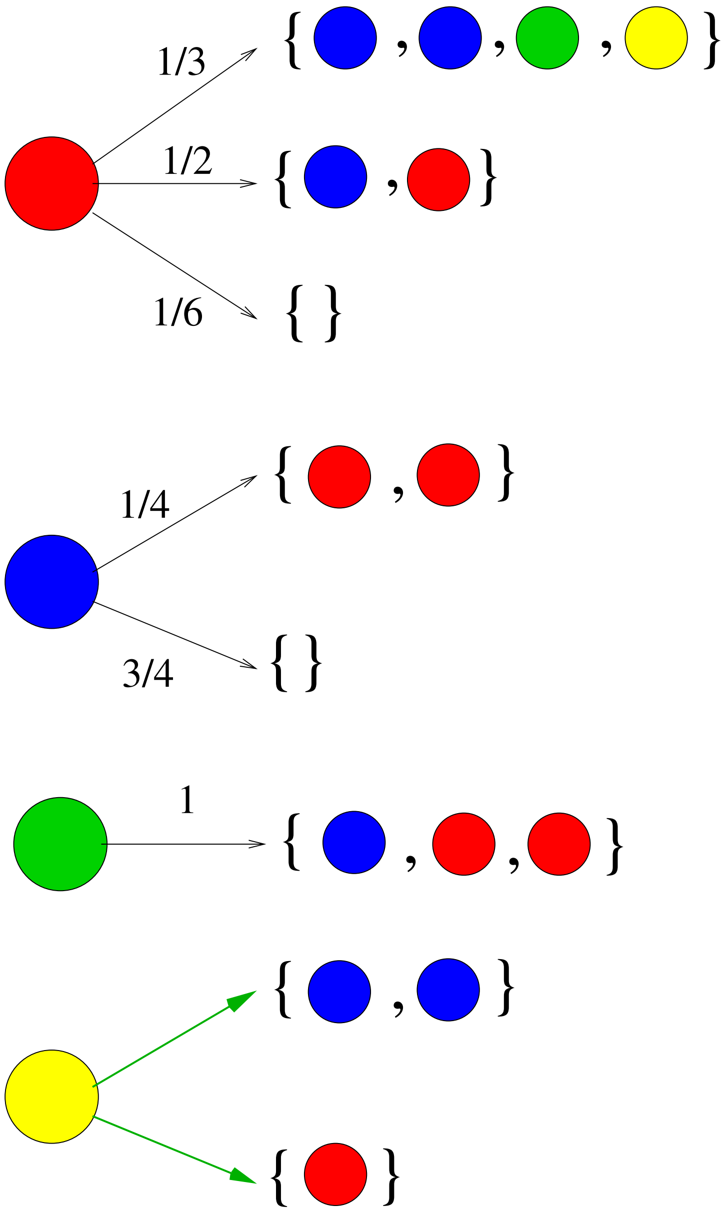

Question

What is the maximum expected size of the tree, starting with one 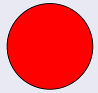 ?

$$x_R \;=$$

BMDPs again: this time optimizing expected tree size

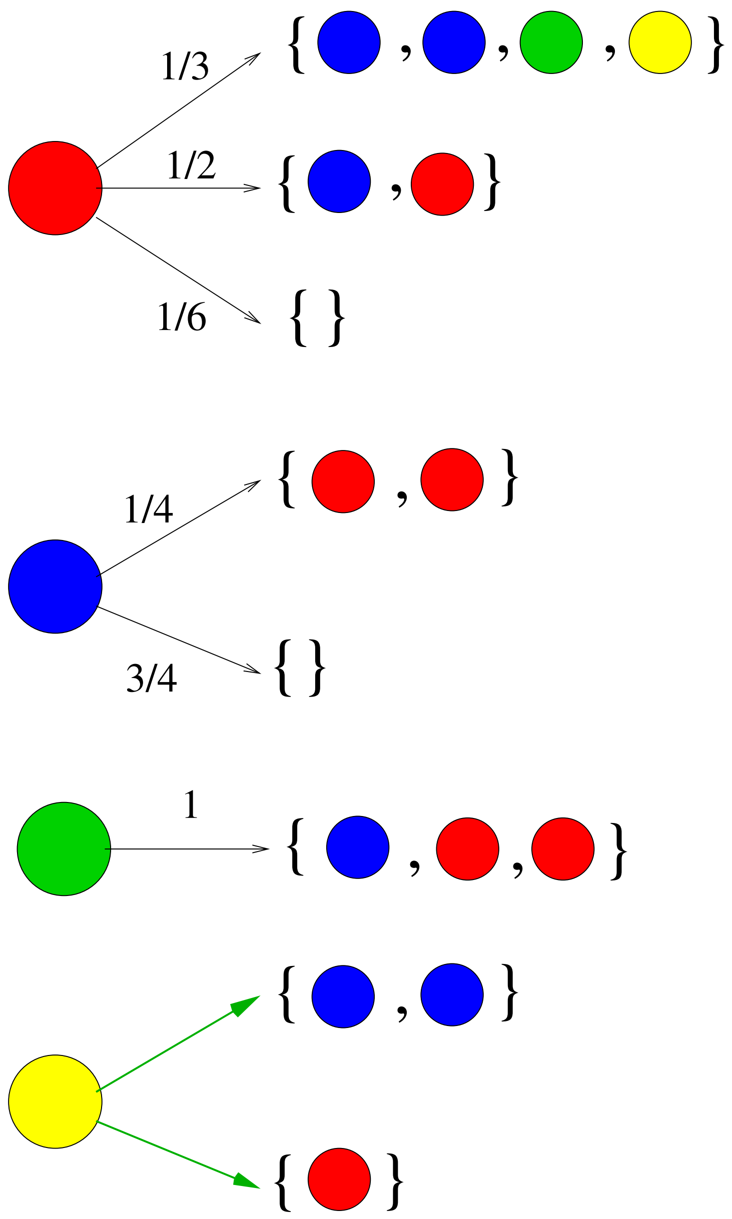

Question

What is the maximum expected size of the tree, starting with one 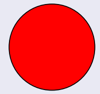 ?

$$x_R = 1 + \frac{1}{3}(2x_B + x_G + x_Y) + \frac{1}{2}(x_B + x_R)$$

BMDPs again: this time optimizing exptected tree size

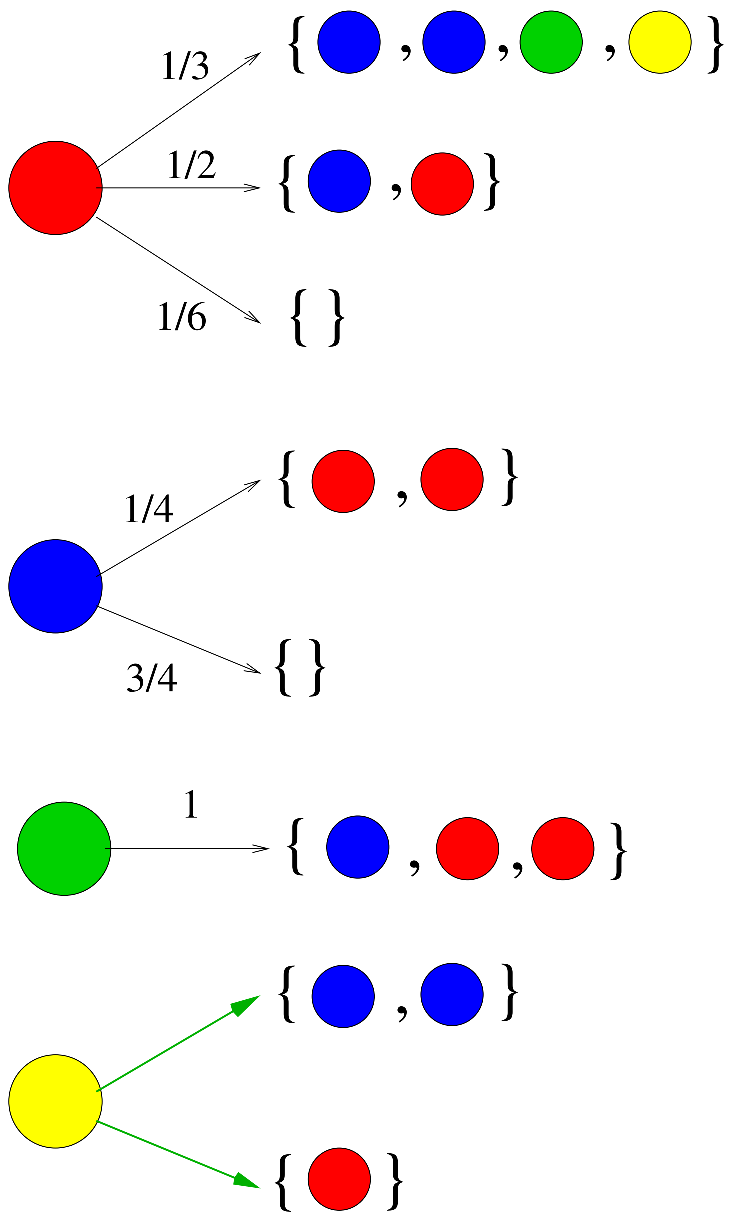

Question

What is the maximum expected size of the tree, starting with one ?

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

We get max/min-linear fixed point equations, $\bar{\mathbf{x}} = P(\bar{\mathbf{x}})$.

Prop [E.-Wojtczak-Yannakakis'09]

The maximum expected tree sizes are the LFP, $\mathbf{r}^* \in [0, +\infty]^3$, of $\bar{\mathbf{x}} = P(\bar{\mathbf{x}})$.

$$\frac{1}{3}x_B^2 x_G x_R + \frac{1}{2}x_B x_R + \frac{1}{6}$$

is a Probabilistic Polynomial: the coefficients are positive and sum to ≤ 1.

A Probabilistic Polynomial System (PPS) of equations, is a system of n equations in n variables, written

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

where each right-hand-side, $P_i(\mathbf{x})$, is a probabilistic polynomial.

A Maximum 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:

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

Minimum 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.
We use max-minPPS to refer to combined max and min PPS equations.

$$5x_B^2 x_G x_R + 2x_B x_R + \frac{1}{6}$$

is a Monotone Polynomial: the coefficients are positive.

A Monotone Polynomial System (MPS), is a system of n equations

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

in n variables where each $P_i(x)$ is a monotone polynomial.

We similiarly define max/minMPSs.

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

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

$$x \leq y \;\Rightarrow\; P(x) \leq P(y)$$

Proposition.

- *[Tarski'55] Every max-minPPS, $x = P(x)$ has a least fixed point, $q^* \in [0,1]^n$, and a greatest fixed point, $g^* \in [0,1]^n$.*
- $q^* = \lim_{k\to\infty} P^k(\mathbf{0})$ *and* $g^* = \lim_{k\to\infty} P^k(\mathbf{1})$.
- *[E.-Yannakakis'05,'06]: q^* is the vector of optimal extinction probabilities (values) for the BMDP (BSSG/BCSG).*

Key Question

Can we compute the probabilities q^* efficiently (in P-time for BMDPs)?

Fact: value iteration is too slow (double-exponentially slow) in worst cases.

Basic properties of (max-min)MPSs

For a max-minMPS, $\mathbf{x} = P(\mathbf{x})$,

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

Proposition

- *[Tarski'55] Every max-minMPS $x = P(x)$ has a LFP, $q^* \in [0, \infty]^n$, and a GFP, $t^* \in [0, \infty]^n$.*

  (We call a (max-min)MPS feasible if LFP $q^* \in [0, \infty)^n$.)

- *For a (max-min)MPS, q^* is the partition function of the corresponding (max-min) Weighted Context-Free Grammar.*

P-time approximation for BMDPs and max/minPPSs

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

Given a 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 new Generalized Newton's Method that uses linear programming in each iteration.

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

Moreover, we can compute an ϵ-optimal static strategy for maximizing or minimizing extinction probabilities for a BMDP, B, in time polynomial in $|B|$ and $\log(1/\epsilon)$.

Newton's method

Newton's method

Seeking a solution to differentiable $F(\mathbf{x}) = \mathbf{0}$, we start at a guess $\mathbf{x}^{(0)} \in \mathbb{R}^n$, and iterate:

$$\mathbf{x}^{(k+1)} := \mathbf{x}^{(k)} - (F'(\mathbf{x}^{(k)}))^{-1} F(\mathbf{x}^{(k)})$$

Here $F'(\mathbf{x})$, is the **Jacobian matrix**:

$$F'(\mathbf{x}) = \begin{bmatrix} \frac{\partial F_1}{\partial x_1} \cdots \frac{\partial F_1}{\partial x_n} \\ \vdots \vdots \vdots \\ \frac{\partial F_n}{\partial x_1} \cdots \frac{\partial F_n}{\partial x_n} \end{bmatrix}$$

For PPSs, $F(x) \equiv (P(x) - x)$, and Newton iteration looks like this:

$$\mathbf{x}^{(k+1)} := \mathbf{x}^{(k)} + (I - P'(\mathbf{x}^{(k)}))^{-1}(P(\mathbf{x}^{(k)}) - \mathbf{x}^{(k)})$$

where $P'(\mathbf{x})$ is the Jacobian of $P(\mathbf{x})$.

Newton's method on PPSs and MPSs

We can easily decompose $\mathbf{x} = P(\mathbf{x})$ into its strongly connected components (SCCs), based on variable dependencies, and eliminate "0" variables.

Theorem [E.-Yannakakis'05]

Decomposed Newton's method, starting at $x^{(0)} = \mathbf{0}$ converges monotonically to the LFP $\mathbf{q}^*$ for any feasible MPS.

But...

- In [E.-Yannakakis'05] we gave no upper bounds for Newton.
- [Esparza,Kiefer,Luttenberger'10] gave bad examples of PPSs, $\mathbf{x} = P(\mathbf{x})$, where $q^* = \mathbf{1}$, but requiring exponentially many Newton iterations, as a function of the encoding size $|P|$ of the equations, to converge to within additive error $< 1/2$.

P-time approximation for PPSs

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

Given a PPS, $\mathbf{x} = P(\mathbf{x})$, 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 both the encoding size $|P|$ of the equations and in j (the number of "bits of precision").

We use Newton's method..... but how?

Qualitative decision problems for PPSs are in P-time

Theorem ([Kolmogorov-Sevastyanov'47,Harris'63])

For certain classes of strongly-connected PPSs, $q_i^* = 1$ for all i iff the spectral radius $\varrho(P'(\mathbf{1}))$ for the moment matrix $P'(\mathbf{1})$ is ≤ 1, and otherwise $q_i^* < 1$ for all i.

Theorem ([E.-Yannakakis'05])

Given a PPS, $\mathbf{x} = P(\mathbf{x})$, deciding whether $q_i^* = 1$ is in P-time.

Qualitative decision problems for PPSs are in P-time

Theorem ([Kolmogorov-Sevastyanov'47,Harris'63])

For certain classes of strongly-connected PPSs, $q_i^* = 1$ for all i iff the spectral radius $\varrho(P'(\mathbf{1}))$ for the moment matrix $P'(\mathbf{1})$ is ≤ 1, and otherwise $q_i^* < 1$ for all i.

Theorem ([E.-Yannakakis'05])

Given a PPS, $\mathbf{x} = P(\mathbf{x})$, deciding whether $q_i^* = 1$ is in P-time.

(It is even in strongly-P-time ([Esparza-Gaiser-Kiefer'10]).)

Qualitative decision problems for PPSs are in P-time

Theorem ([Kolmogorov-Sevastyanov'47,Harris'63])

For certain classes of strongly-connected PPSs, $q_i^* = 1$ for all i iff the spectral radius $\varrho(P'(\mathbf{1}))$ for the moment matrix $P'(\mathbf{1})$ is ≤ 1, and otherwise $q_i^* < 1$ for all i.

Theorem ([E.-Yannakakis'05])

Given a PPS, $\mathbf{x} = P(\mathbf{x})$, deciding whether $q_i^* = 1$ is in P-time.

(It is even in strongly-P-time ([Esparza-Gaiser-Kiefer'10]).)

Deciding whether $q_i^* = 0$ is also easily in (strongly) P-time.

Algorithm for approximating the LFP q^* for PPSs

1. Find and remove all variables x_i such that $q_i^* = 0$ or $q_i^* = 1$.
2. On the resulting system of equations, run Newton's method starting from $\mathbf{0}$.

Algorithm for approximating the LFP q^* for PPSs

1. Find and remove all variables x_i such that $q_i^* = 0$ or $q_i^* = 1$.
2. On the resulting system of equations, run Newton's method starting from $\mathbf{0}$.

Theorem ([E.-Stewart-Yannakakis'12])

Given a PPS $\mathbf{x} = P(\mathbf{x})$ *with LFP* $\mathbf{0} < \mathbf{q}^* < \mathbf{1}$*, if we apply Newton starting at* $\mathbf{x}^{(0)} = \mathbf{0}$*, then*

$$\|\mathbf{q}^* - \mathbf{x}^{(4|P|+j)}\|_\infty \leq 2^{-j}$$

Algorithm with rounding

1. Find and remove all variables x_i such that $q_i^* = 0$ or $q_i^* = 1$.
2. On the resulting system of equations, run Newton's method starting from $\mathbf{0}$.
3. After each iteration, round down to a multiple of 2^{-h}

Theorem ([E.-Stewart-Yannakakis'12])

If, after each Newton iteration, we round down to a multiple of 2^{-h} where $h := 4|P| + j + 2$, then after h iterations $\|\mathbf{q}^* - \mathbf{x}^{(h)}\|_\infty \leq 2^{-j}$.

Thus, we obtain a P-time algorithm (in the standard Turing model) for computing q^* to any desired accuracy.

High level picture of proof

- For a PPS, $x = P(x)$, with LFP $\mathbf{0} < \mathbf{q}^* < \mathbf{1}$, $P'(q^*)$ is a non-negative square matrix with spectral radius $\varrho(P'(q^*)) < 1$.
- So, $(I - P'(q^*))$ is non-singular, and $(I - P'(q^*))^{-1} = \sum_{i=0}^{\infty}(P'(q^*))^i$.
- We can show the # of Newton iterations needed to get within $\epsilon > 0$ is

$$\approx \log \|(I - P'(q^*))^{-1}\|_\infty + \log \frac{1}{\epsilon}$$

- $\|(I - P'(q^*))^{-1}\|_\infty$ is inversely related to the distance $|1 - \varrho(P'(q^*))|$, which in turn is related to $\min_i(1 - q_i^*)$, which we can lower bound!
- Uses lots of Perron-Frobenius theory, among other things...

Towards Generalized Newton's Method: Newton iteration as a first-order (Taylor) approximation

An iteration of Newton's method on a PPS, applied on current vector $y \in \mathbb{R}^n$, solves the equation

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

where

$$P^{\mathbf{y}}(\mathbf{x}) \equiv P(\mathbf{y}) + P'(\mathbf{y})(\mathbf{x} - \mathbf{y})$$

is the linear (first-order Taylor) approximation of $P(x)$ at the point $\mathbf{y}$.

Generalized Newton's method

Linearization of max/minPPSs

Given a maxPPS

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

We define the linearization, $P^y(x)$, by:

$$(P^{\mathbf{y}}(\mathbf{x}))_i = \max\{p_{i,j}(\mathbf{y}) + \nabla p_{i,j}(\mathbf{y}).(\mathbf{x} - \mathbf{y}) : j = 1, \ldots, m_i\} \qquad i = 1, \ldots, n$$

Generalized Newton's method

Linearization of max/minPPSs

Given a maxPPS

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

We define the linearization, $P^y(x)$, by:

$$(P^{\mathbf{y}}(\mathbf{x}))_i = \max\{p_{i,j}(\mathbf{y}) + \nabla p_{i,j}(\mathbf{y}).(\mathbf{x} - \mathbf{y}) : j = 1, \ldots, m_i\} \qquad i = 1, \ldots, n$$

Generalised Newton's method: iteration applied at vector y

Solve $P^{\mathbf{y}}(\mathbf{x}) = \mathbf{x}$. Specifically:

For a maxPPS, minimize $\sum_i x_i$ subject to $P^{\mathbf{y}}(\mathbf{x}) \leq \mathbf{x}$;

For a minPPS, maximize $\sum_i x_i$ subject to $P^{\mathbf{y}}(\mathbf{x}) \geq \mathbf{x}$;

These can both be phrased as linear programming problems. Their optimal solution solves $P^{\mathbf{y}}(\mathbf{x}) = \mathbf{x}$, and yields one GNM iteration.

Algorithm for max/minPPSs

1) Find and remove all variables x_i such that $q_i^* = 0$ or $q_i^* = 1$. Checking $q_i^* = 0$ is again easy. Checking $q^* = 1$ is harder:

- **Theorem** ([E.-Yannakakis'06]) Checking $q_i^* = 1$ is decidable in P-time using linear programming.
  Reduces to spectral radius optimization for non-negative square matrices: given k choices for each row of a $n \times n$ matrix $M \geq \mathbf{0}$, can we choose the rows to make $\varrho(M) > 1$? Solvable by LP.

Algorithm for max/minPPSs

1) Find and remove all variables x_i such that $q_i^* = 0$ or $q_i^* = 1$. Checking $q_i^* = 0$ is again easy. Checking $q^* = 1$ is harder:

- **Theorem** ([E.-Yannakakis'06]) Checking $q_i^* = 1$ is decidable in P-time using linear programming.
  Reduces to spectral radius optimization for non-negative square matrices: given k choices for each row of a $n \times n$ matrix $M \geq \mathbf{0}$, can we choose the rows to make $\varrho(M) > 1$? Solvable by LP.

2) On the resulting equations, run Generalized Newton's Method, starting from $\mathbf{0}$. After each iteration, round down to a multiple of 2^{-h}.
Each iteration of GNM can be computed in P-time by solving an LP.

Algorithm for max/minPPSs

1) Find and remove all variables x_i such that $q_i^* = 0$ or $q_i^* = 1$. Checking $q_i^* = 0$ is again easy. Checking $q^* = 1$ is harder:

- **Theorem** ([E.-Yannakakis'06]) Checking $q_i^* = 1$ is decidable in P-time using linear programming.
  Reduces to spectral radius optimization for non-negative square matrices: given k choices for each row of a $n \times n$ matrix $M \geq \mathbf{0}$, can we choose the rows to make $\varrho(M) > 1$? Solvable by LP.

2) On the resulting equations, run Generalized Newton's Method, starting from $\mathbf{0}$. After each iteration, round down to a multiple of 2^{-h}.
Each iteration of GNM can be computed in P-time by solving an LP.

Theorem [E.-Stewart-Yannakakis'12]**:** Given a max/minPPS $\mathbf{x} = P(\mathbf{x})$ with LFP $\mathbf{0} < \mathbf{q}^* < \mathbf{1}$, if we apply rounded GNM starting at $\mathbf{x}^{(0)} = \mathbf{0}$, using $h := 4|P| + j + 1$ bits of precision, then $\|\mathbf{q}^* - \mathbf{x}^{(4|P|+j+1)}\|_\infty \leq 2^{-j}$.

Qualitative & quantitative extinction for BSSGs

Theorem ([E.-Yannakakis'06])

Given a BSSG, deciding whether the extinction value is $q_i^* = 1$ is in **NP** $\cap$ **coNP**.

And is at least as hard as computing the value of a finite-state SSG.

Theorem ([E.-Stewart-Yannakakis'12])

Given a BSSG extinction game, and given $\epsilon > 0$, we can compute a vector $v \in [0,1]^n$, such that $\|v - q^*\|_\infty \leq \epsilon$, and we can compute ϵ-optimal static strategies in **FNP**
(and in **PLS**, *using an approximate strategy improvement method).*

Optimal Reachability problem for BMDPs

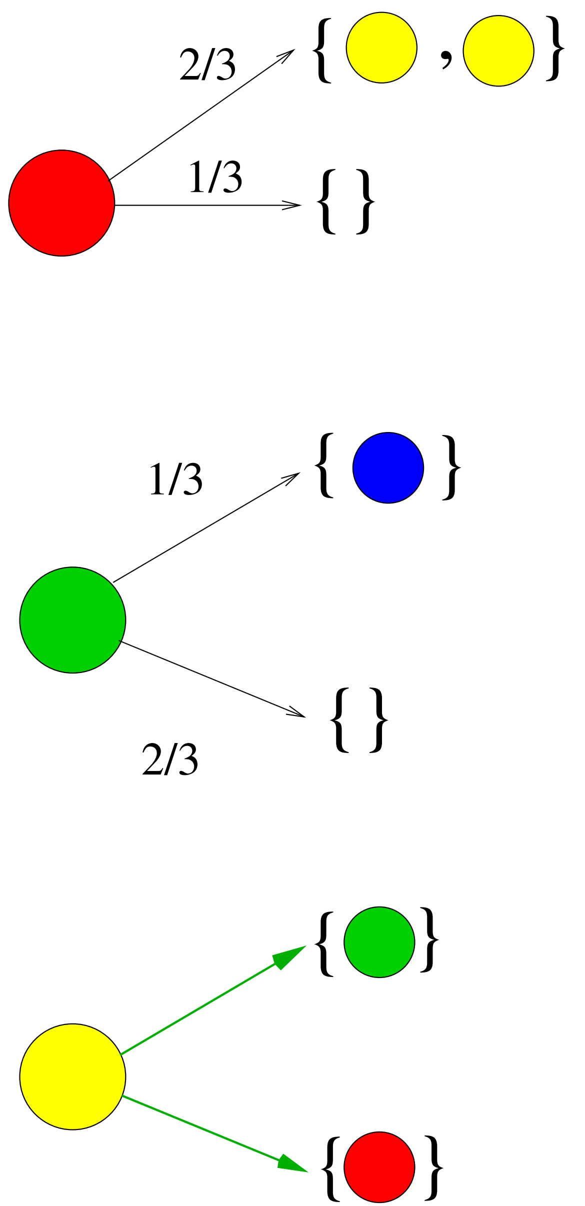

Optimal Reachability problem for BMDPs

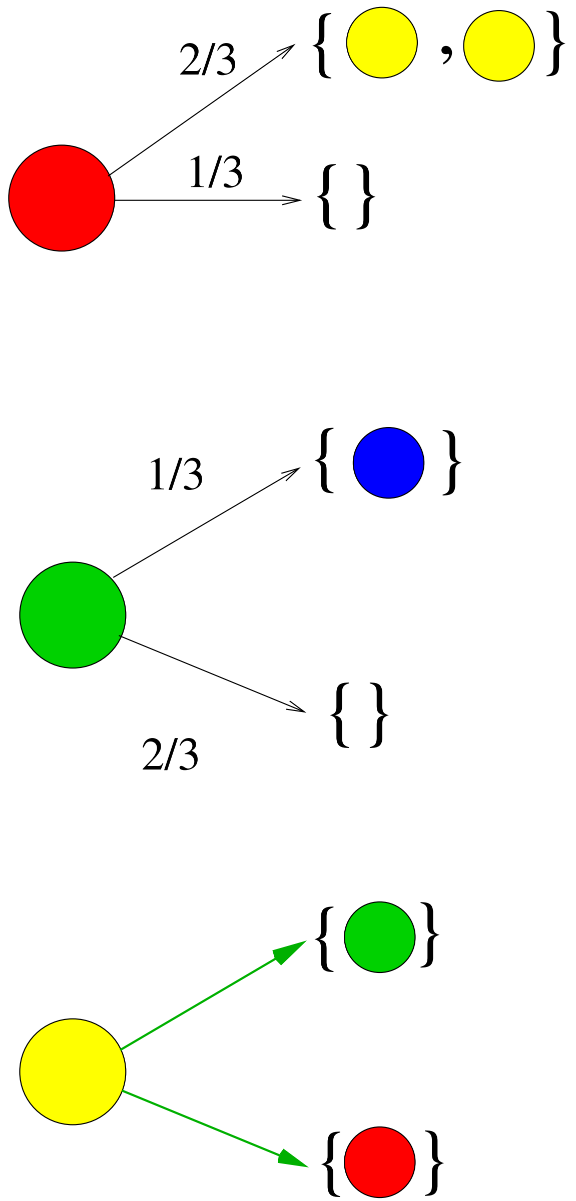

Question

What is the supremum probability of reaching (blue), starting with one (red)?

Optimal Reachability problem for BMDPs

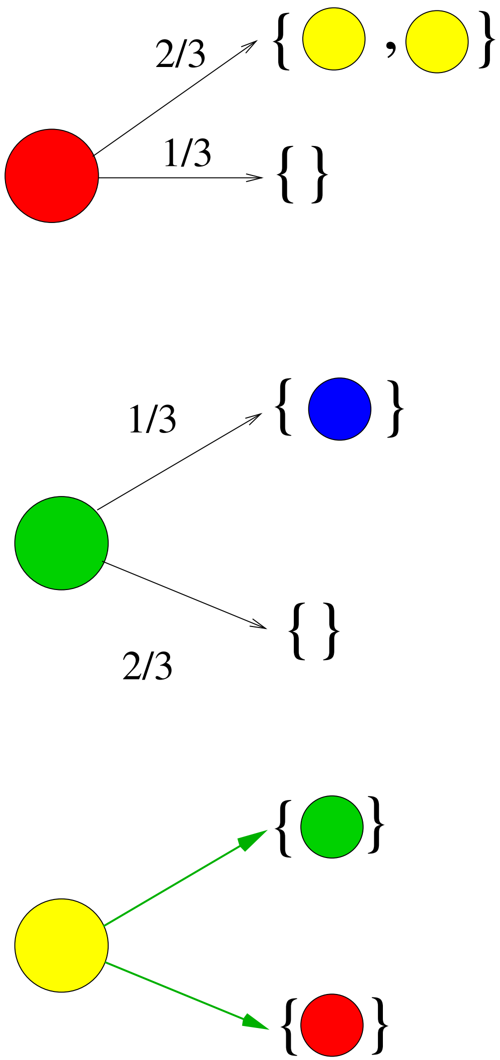

Same Question (rephrased)

What is the infimum probability of not reaching 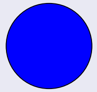, starting with one 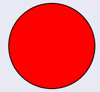 ?

Optimal Reachability problem for BMDPs

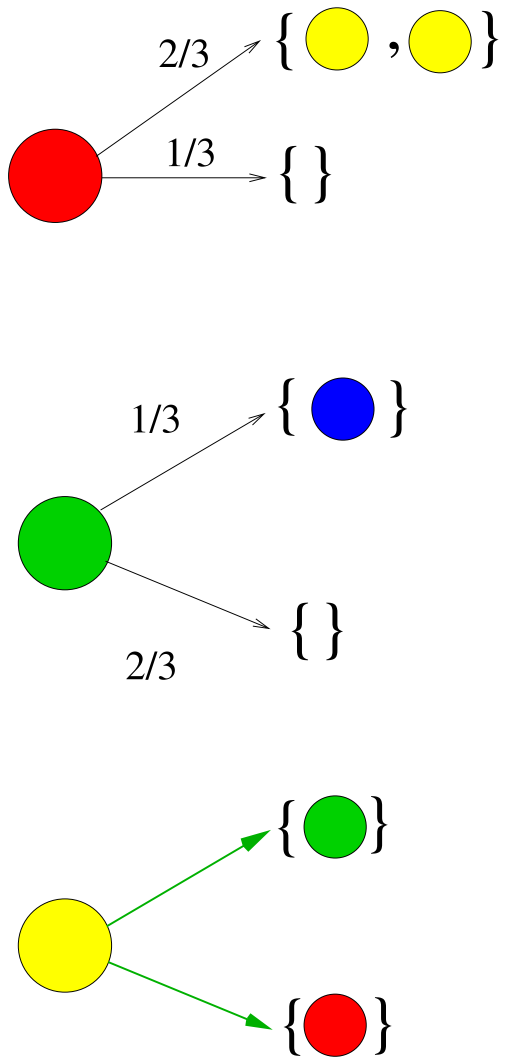

Same Question (rephrased)

What is the infimum probability of not reaching 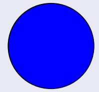, starting with one 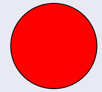 ?

$$y_R \;=$$

Optimal Reachability problem for BMDPs

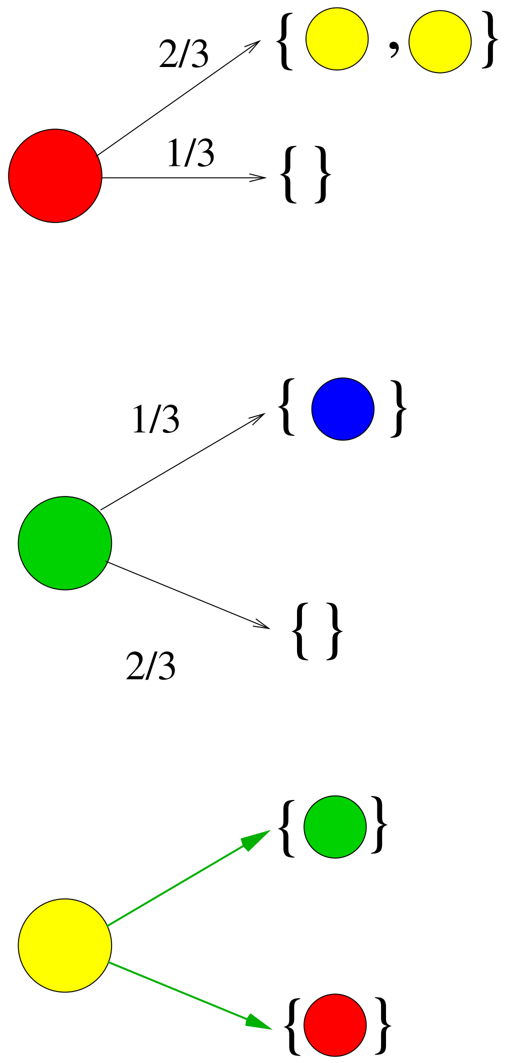

Same Question (rephrased)

What is the infimum probability of not reaching , starting with one ?

$$y_R = \frac{2}{3} y_Y y_Y + \frac{1}{3}$$

$$y_G = \frac{2}{3}$$

$$y_Y = \min\{y_G, y_R\}$$

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

Thm. [E.-Stewart-Yannakakis'15]

The supremum reachability probabilities are $\mathbf{1} - \mathbf{g}^*$, where $\mathbf{g}^* \in [0,1]^3$ is the Greatest Fixed Point, of $\bar{\mathbf{y}} = Q(\bar{\mathbf{y}})$.

Optimal Reachability problem for BMDPs

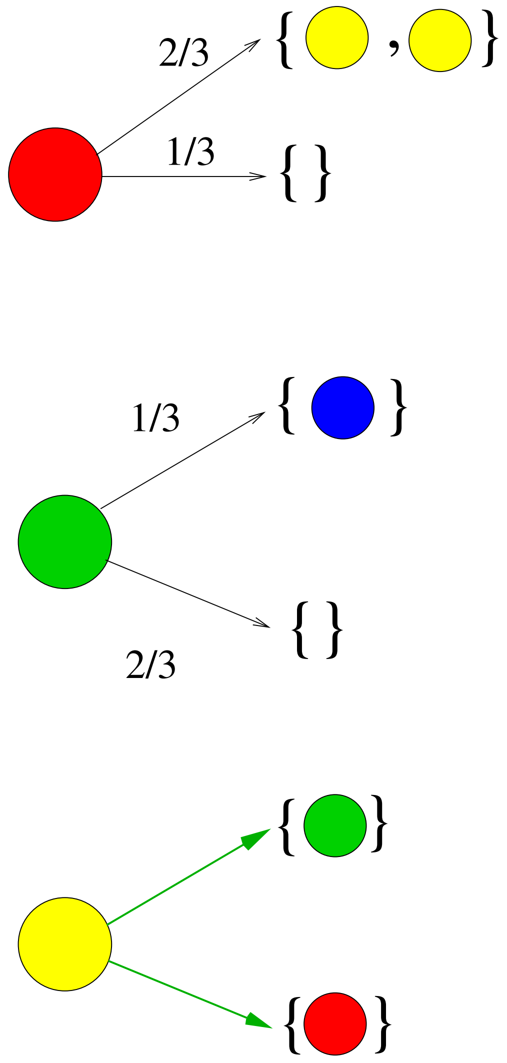

Question

What is the maximum probability of not reaching 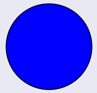, starting with one 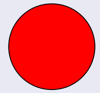 ?

$$y_R = \frac{2}{3} y_Y y_Y + \frac{1}{3}$$

$$y_G = \frac{2}{3}$$

$$y_Y = \max\{y_G, y_R\}$$

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

Thm. [E.-Stewart-Yannakakis'15]

The minimum reachability probabilities are $\mathbf{1} - \mathbf{g}^*$, where $\mathbf{g}^* \in [0,1]^3$ is the Greatest Fixed Point of $\bar{\mathbf{y}} = Q(\bar{\mathbf{y}})$.

P-time approximation of optimal reachability probability for BMDPs

Theorem ([E.-Stewart-Yannakakis, 2015])

Given a max/minPPS, $\mathbf{y} = Q(\mathbf{y})$, with GFP $\mathbf{g}^* \in [0,1]^n$, we can compute a rational vector $\mathbf{v} \in [0,1]^n$ such that

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

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

We again establish this via the Generalized Newton's Method, but with a subtly different preprocessing step, which results in convergence to the GFP $\mathbf{g}^*$, instead of the LFP $\mathbf{q}^*$.

Qualitative/ quantitative reachability problems for BSSGs

Theorem [E.-Stewart-Yannakakis'15]

- The value of a BSSG reachability game is captured by the GFP of a max-minPPS.

Qualitative/ quantitative reachability problems for BSSGs

Theorem [E.-Stewart-Yannakakis'15]

- The value of a BSSG reachability game is captured by the GFP of a max-minPPS.
- We can approximate the value, and compute ϵ-optimal stratgies, for a BSSG reachability game in FNP.
  (For BMDPs, we can compute ϵ-optimal strategies in P-time.)

Qualitative/ quantitative reachability problems for BSSGs

Theorem [E.-Stewart-Yannakakis'15]

- The value of a BSSG reachability game is captured by the GFP of a max-minPPS.
- We can approximate the value, and compute ϵ-optimal stratgies, for a BSSG reachability game in FNP.
  (For BMDPs, we can compute ϵ-optimal strategies in P-time.)
- For BSSG reachability games, limit-sure $=$ almost-sure, and we can decide all qualitative questions in P-time.

  (**Note:** This contrasts sharply with BSSG extinction games.)

one-counter Markov chains (discrete-time QBDs)

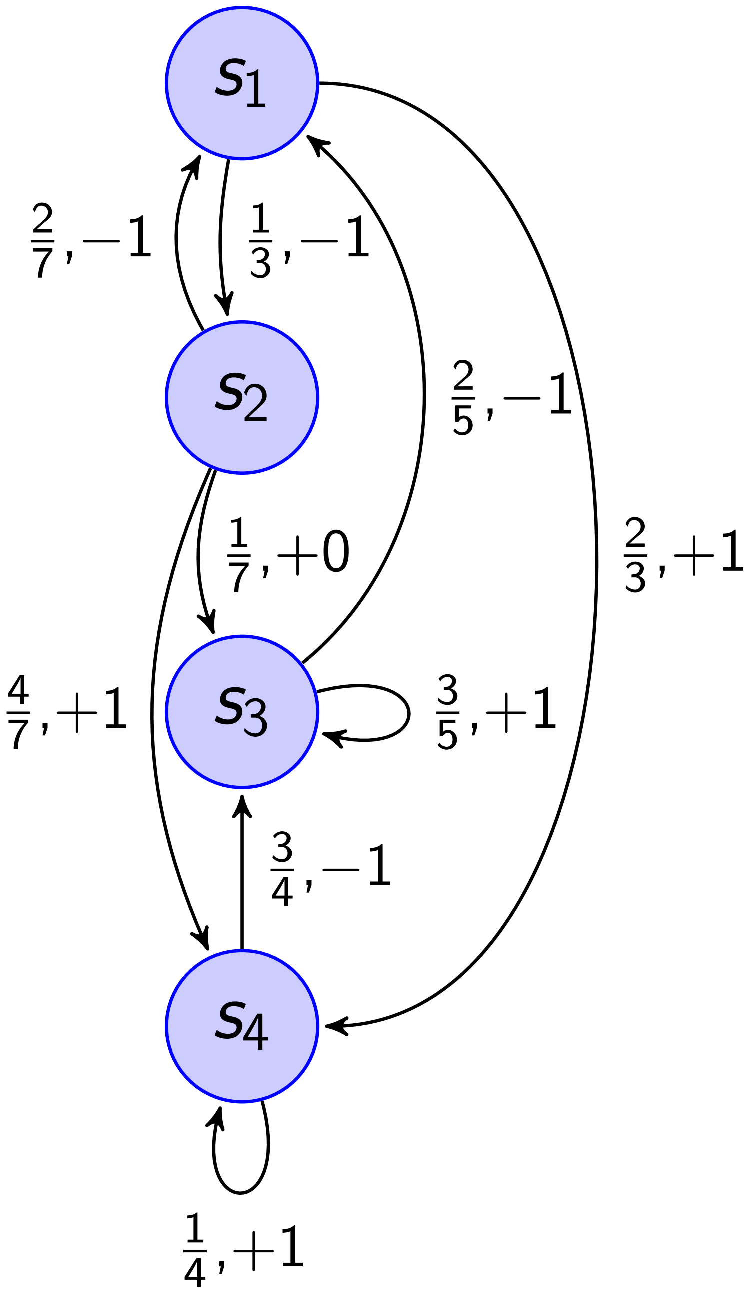

Question: What is the probability of terminating (reaching counter value $= 0$ for the first time) in state $\mathbf{s}_2$, if we start with counter value $= 1$ in state $\mathbf{s}_1$?

one-counter Markov chains (discrete-time QBDs)

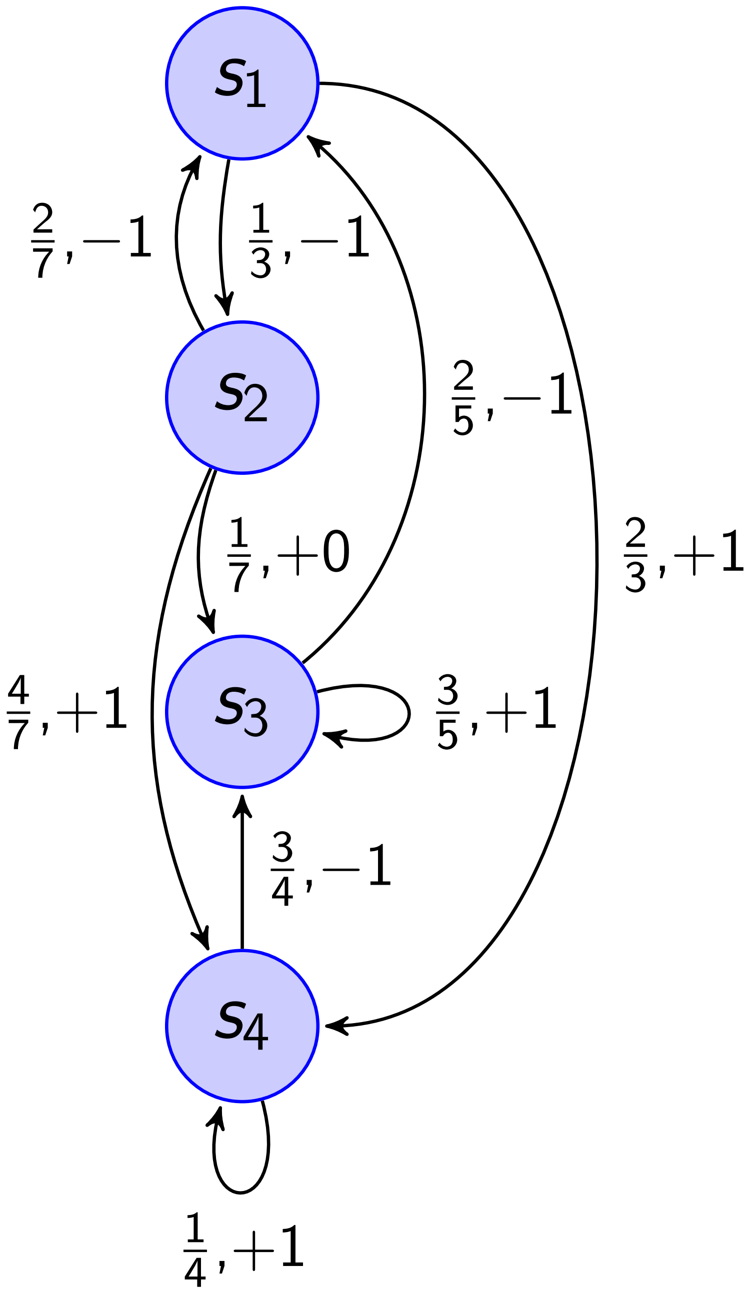

Question: What is the probability of terminating (reaching counter value $= 0$ for the first time) in state $\mathbf{s}_2$, if we start with counter value $= 1$ in state $\mathbf{s}_1$?

$$x_{1,2} \quad =$$

one-counter Markov chains (discrete-time QBDs)

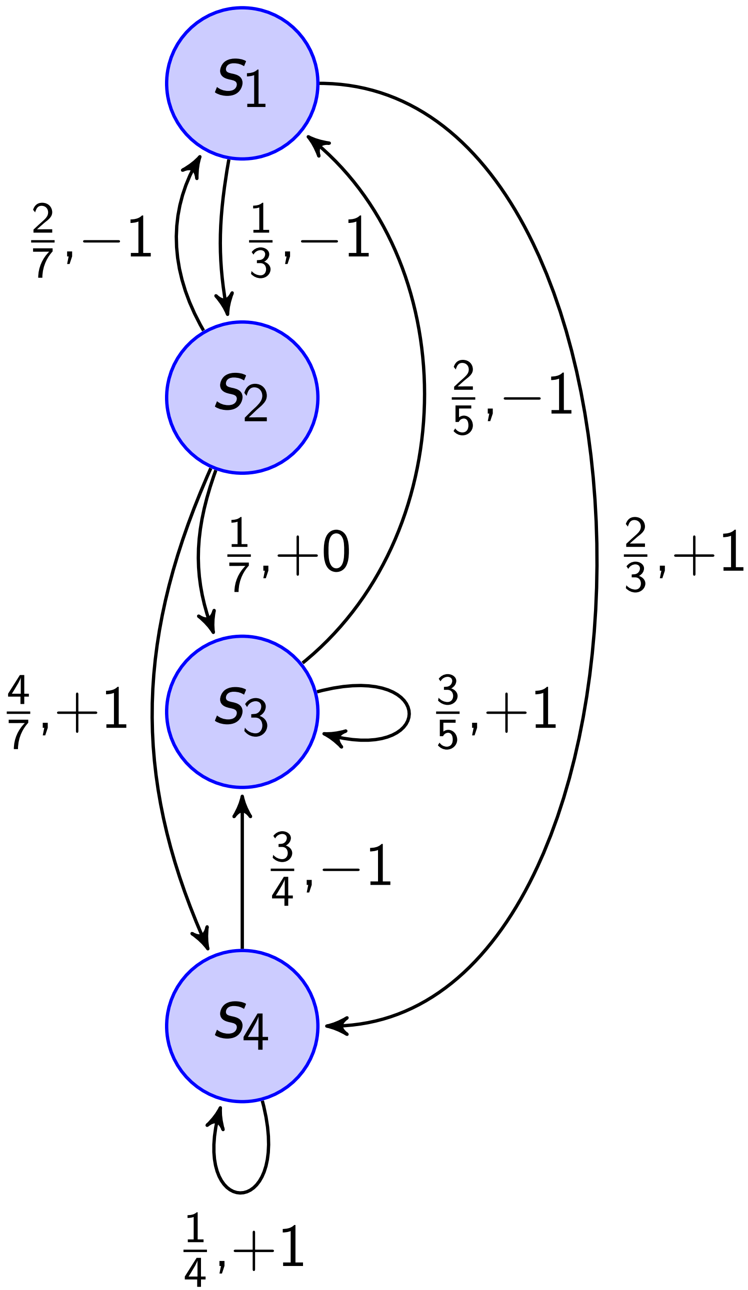

Question: What is the probability of terminating (reaching counter value $= 0$ for the first time) in state $\mathbf{s}_2$, if we start with counter value $= 1$ in state $\mathbf{s}_1$?

$$
\begin{aligned}
x_{1,2} &= \frac{1}{3} + \frac{2}{3} \sum_j x_{4,j} x_{j,2} \\
x_{4,3} &= \frac{3}{4} + \frac{1}{4} \sum_j x_{4,k} x_{k,2} \\
\cdots &= \cdots
\end{aligned}
$$

one-counter Markov chains (discrete-time QBDs)

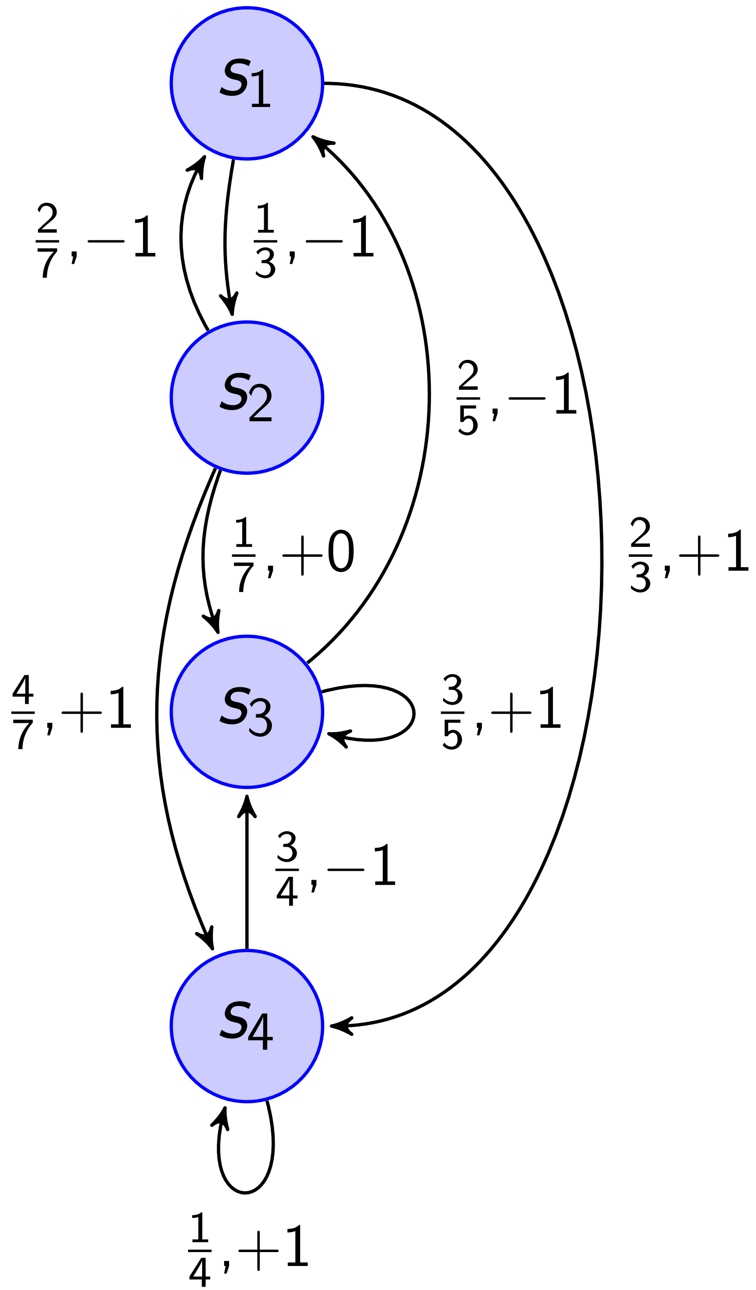

Question: What is the probability of terminating (reaching counter value $= 0$ for the first time) in state $\mathbf{s}_2$, if we start with counter value $= 1$ in state $\mathbf{s}_1$?

$$
\begin{aligned}
x_{1,2} &= \frac{1}{3} + \frac{2}{3} \sum_j x_{4,j} x_{j,2} \\
x_{4,3} &= \frac{3}{4} + \frac{1}{4} \sum_j x_{4,k} x_{k,2} \\
\cdots &= \cdots
\end{aligned}
$$

one-counter Markov chains (discrete-time QBDs)

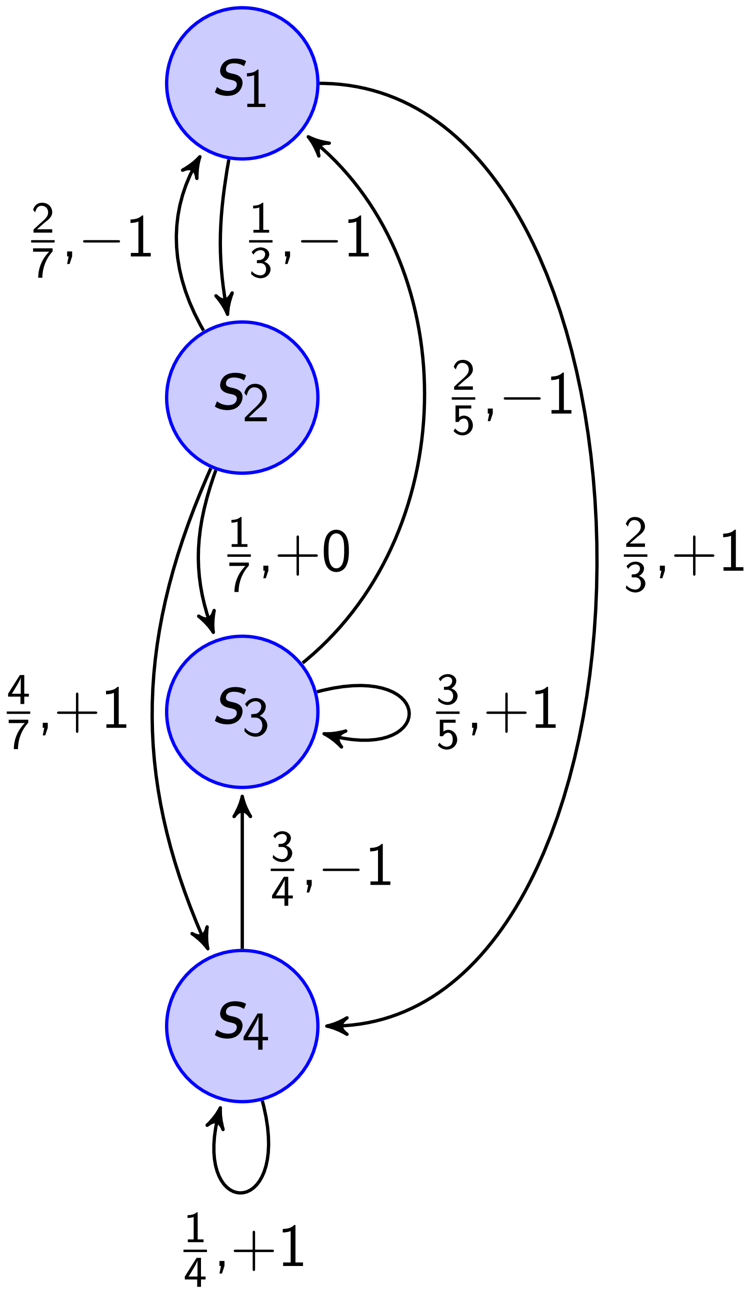

Question: What is the probability of terminating (reaching counter value $= 0$ for the first time) in state $\mathbf{s}_2$, if we start with counter value $= 1$ in state $\mathbf{s}_1$?

$$
\begin{aligned}
x_{1,2} &= \frac{1}{3} + \frac{2}{3} \sum_j x_{4,j} x_{j,2} \\
x_{4,3} &= \frac{3}{4} + \frac{1}{4} \sum_j x_{4,k} x_{k,2} \\
\cdots &= \cdots
\end{aligned}
$$

Fact (cf., [Neuts,1970s])

The termination probabilities are the LFP, $\mathbf{q}^* \in [0,1]^{4\times 4}$.

Theorem [E.-Wojtczak-Yannakakis'08], [Stewart-E.-Yannakakis'13]

The termination probabilities of a QBD, Q, can be computed to desired accuracy $\epsilon > 0$ in time polynomial in both the encoding size $|Q|$ and $\log(1/\epsilon)$ (in the standard Turing model of computation).

- Proof analyzes Newton's method on the very particular feasible MPSs arising for 1-counter Markov Chains (QBDs).
- [Stewart-E.-Yannakakis,'13] gives upper bounds for Newton's method on arbitrary feasible MPSs. Result for QBPs follows as a special case. (Worst-case bound, arising already for the feasible MPSs of Recursive Markov Chains, is exponential.)
- [Esparza-Kiefer-Luttenberger'10] earlier gave exponential upper bounds on Newton iterations for "strongly-connected"-MPSs.

one-counter Markov Decision Processes

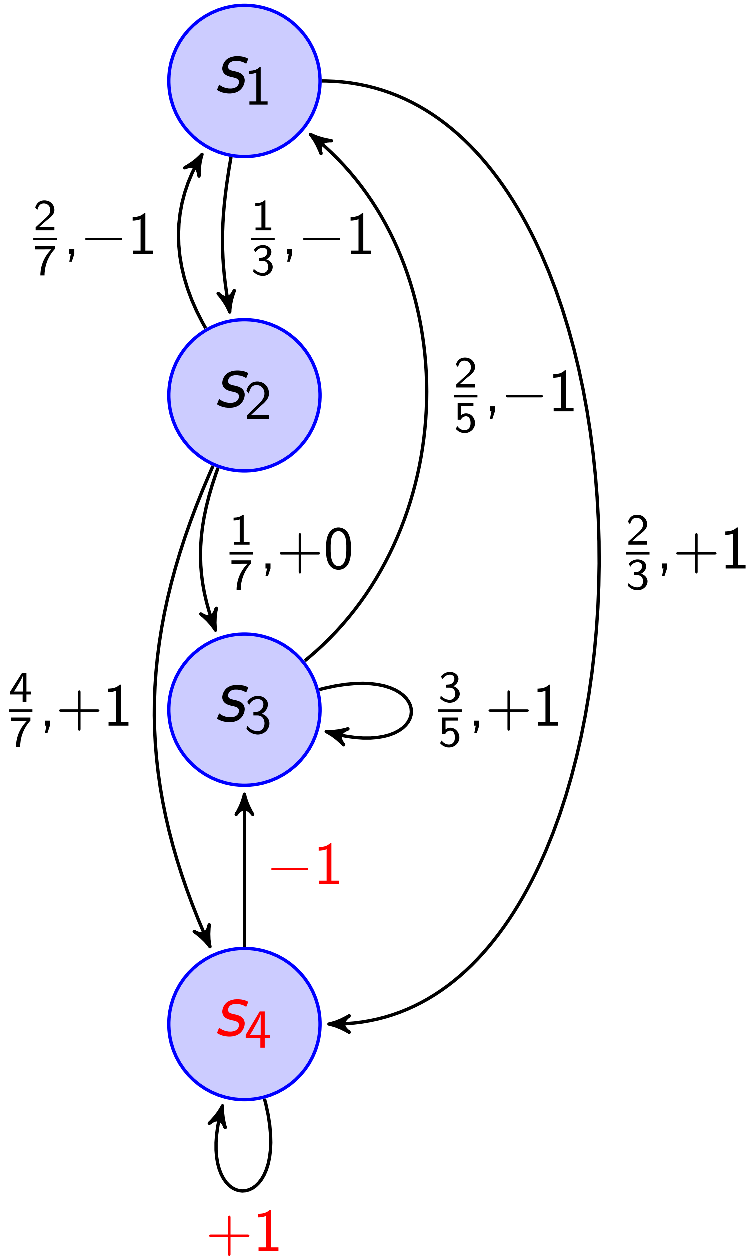

Question: What is the optimal (supremum or infimum) probability of termination (reaching counter value $= 0$) in any state, starting with counter value $= 1$ in state $\mathbf{s}_1$?

one-counter Markov Decision Processes

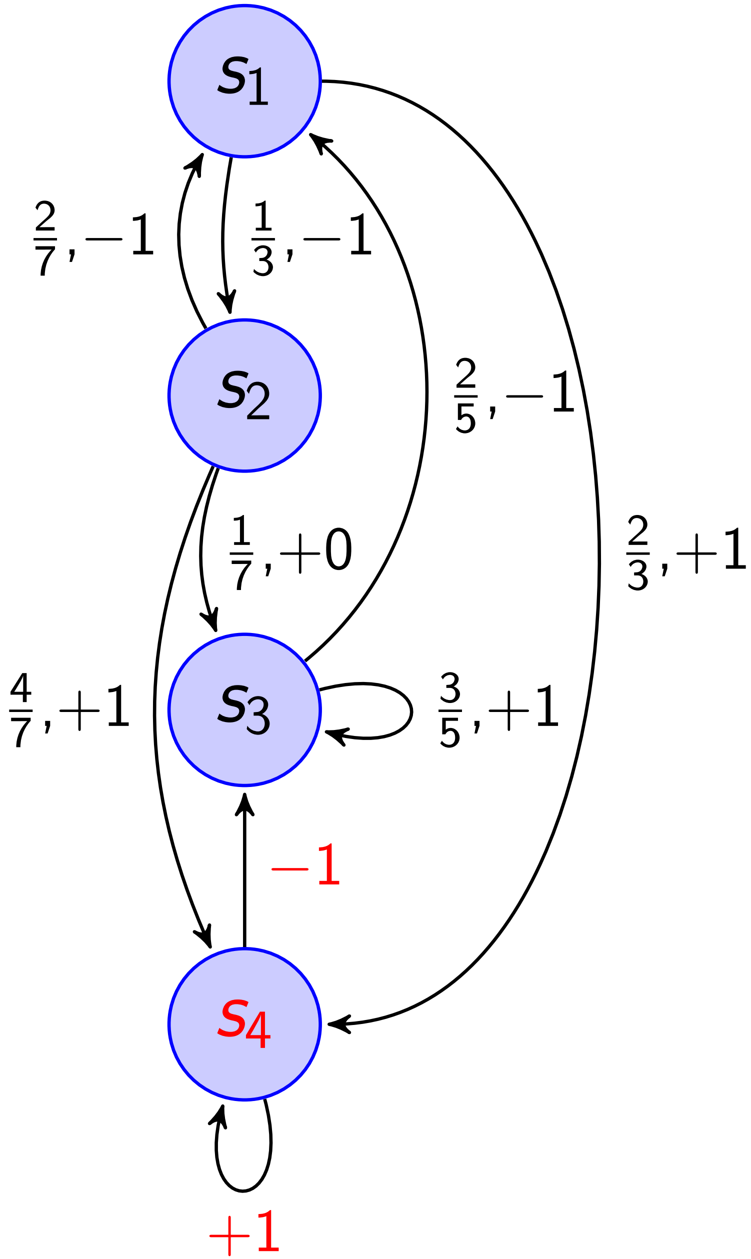

Question: What is the optimal (supremum or infimum) probability of termination (reaching counter value $= 0$) in any state, starting with counter value $= 1$ in state $\mathbf{s}_1$?

Unfortunately, we do not know any max/min-MPS equations that capture these optimal probabilities.

one-counter Markov Decision Processes

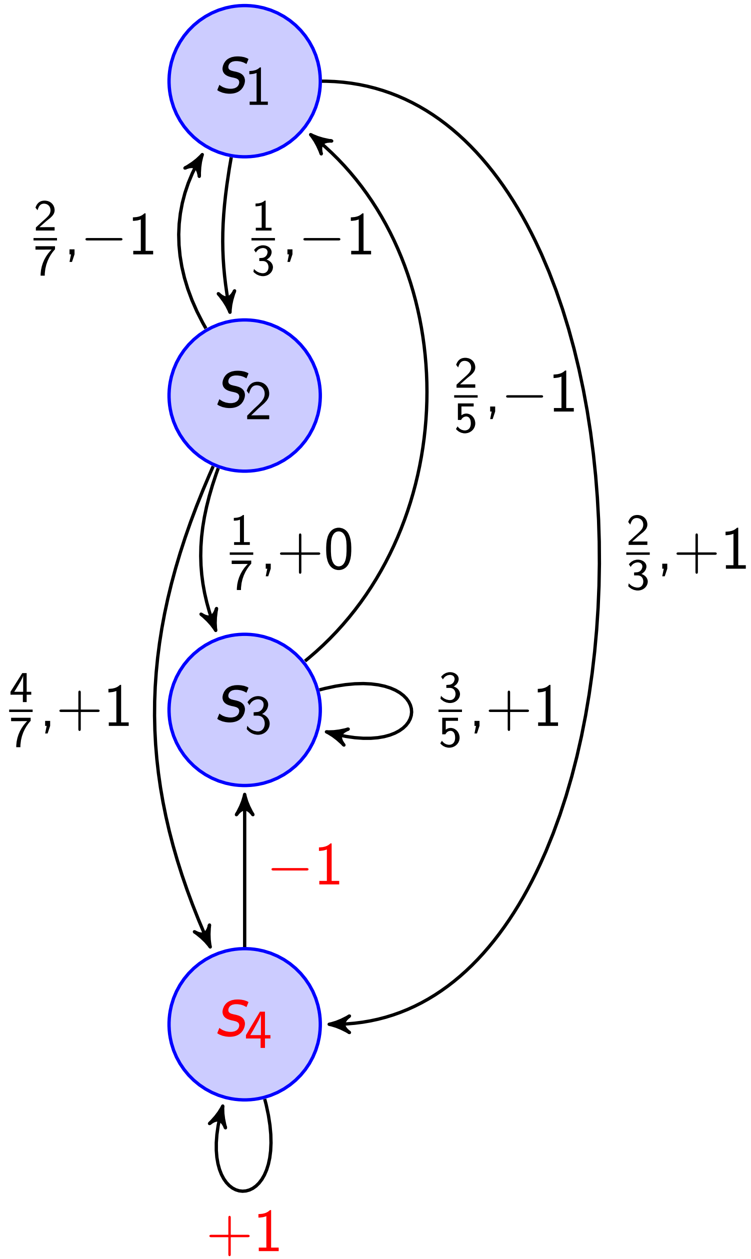

Question: What is the optimal (supremum or infimum) probability of termination (reaching counter value $= 0$) in any state, starting with counter value $= 1$ in state $\mathbf{s}_1$?

Unfortunately, we do not know any max/min-MPS equations that capture these optimal probabilities.

But we do have algorithms to compute them.....

Theorem [Brazdil-Brózek-E.-Kucera,2011]

Given a OC-MDP, M, we can compute the optimal (supremum/infimum) termination probability to accuracy $\epsilon > 0$ in time polynomial in $\log(1/\epsilon)$, and (unfortunately) exponential in $|M|$.

Algorithm involves solving exponentially large finite-state (mean-payoff) MDPs. Proof uses an intriguing martingale derived from LPs associated with optimizing mean-payoff MDPs, and the Azuma inequality.

Theorem [Brazdil-Brózek-E.-Kucera-Wojtzak,2010]

Given a OC-MDP, we can decide almost-sure = limit-sure termination in any state in P-time.

Proof uses LPs, and limit theorems for sums of i.i.d. random variables.

Theorem [Brazdil-Brózek-E.-Kucera-Wojtzak,2010]

Given a OC-MDP, deciding almost-sure termination in a specific state is **PSPACE-hard**, and in **EXPTIME**.

Recursive Markov Chains ($\approx$ pPDSs $\approx$ tree-like-QBDs)

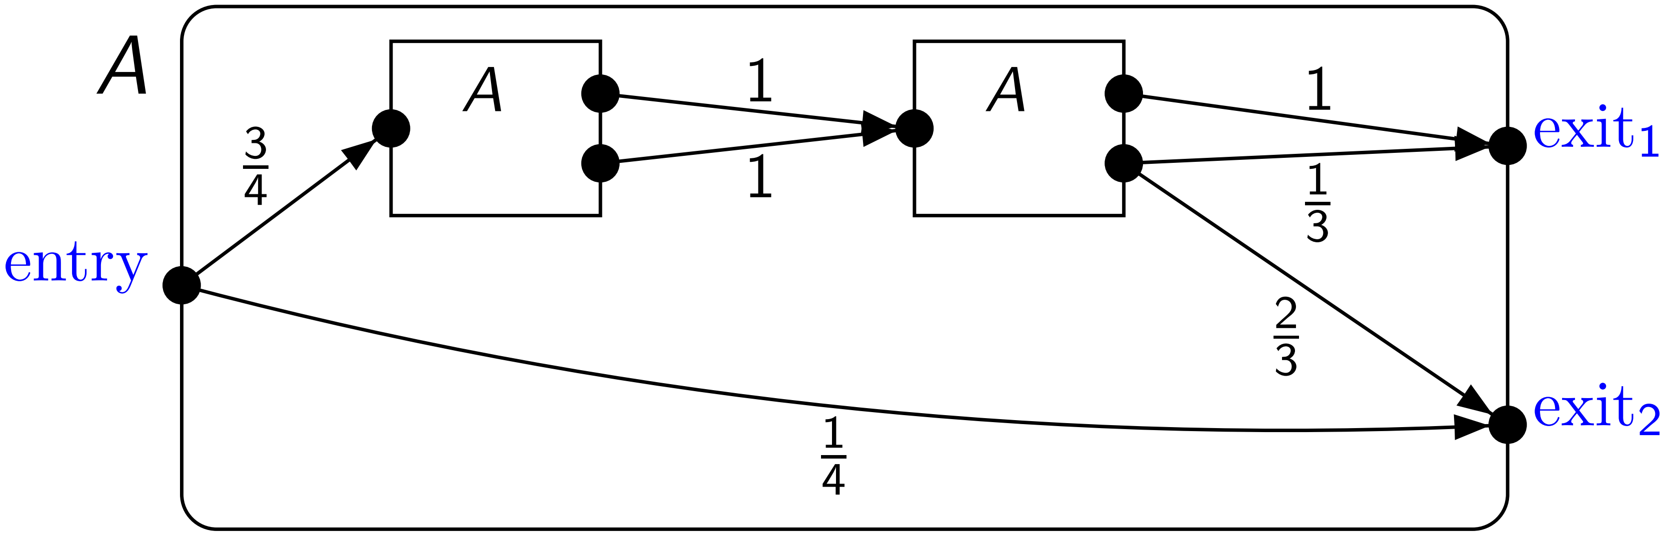

What is the probability of terminating at exit_2, starting at entry?

$$x_2 \quad =$$

Recursive Markov Chains ($\approx$ pPDSs $\approx$ tree-like-QBDs)

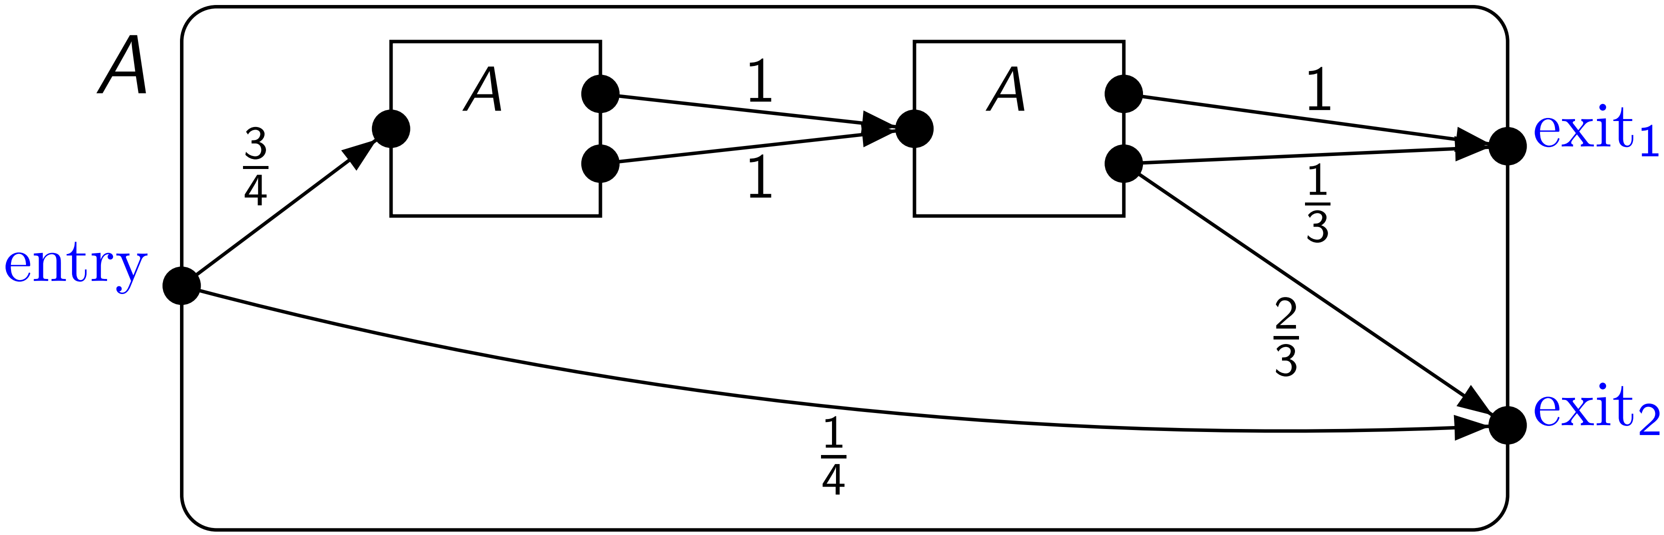

What is the probability of terminating at exit_2, starting at entry?

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

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

Fact: ([E.-Yannakakis'05]) The Least Fixed Point, $q^* \in [0,1]^n$, gives the termination probabilities.

approximation for Recursive Markov chains is "hard"

Theorem [E.-Yannakakis'05,'09]

Any non-trivial approximation of the termination probabilities q^* of an RMC (with 2 or more exits) is SqrtSum-hard and PosSLP-hard.

In fact, deciding whether (a.) $q_1^* = 1$ or (b.) $q_1^* < \epsilon$, given the promise that one of the two is the case, is PosSLP-hard.

(Thus, even approximation in **NP** would yield a major breakthrough on the complexity of the BSS model and exact numerical computation; and P-time approximation is very unlikely.)

Note: this is despite the fact that Newton's method converges monotonically, starting from $\mathbf{0}$, to the LFP q^*, for all feasible MPSs.

Theorem [E.-Yannakakis'05]

For Recursive Markov Decision Processes (with ≥ 10 exits), any non-trivial apporoximation of the optimal termination probabilities is not computable at all!

Model checking

Algorithms & complexity of many model-checking questions have also been addressed, for these infinite-state MCs, MDPs, and SSGs, often by building on termination/reachability analysis.

But still may open questions remain. For example:

Quantitative CTL model checking of BMDPs:
Given BMDP, M, start color c, and CTL formula φ over the color alphabet, can we compute/approximate:
$\sup_{\sigma \in \mathrm{S}trategy} Pr(\mathrm{Tree}_c^{\sigma}(M) \models \varphi)$.
(We only know approximation computability for fragments of CTL.)

Many embarrassing open questions

- The complexity, or even decidability, of optimizing the expected tree depth for a given BMDP. Optimizing expected tree size is in P-time ([E.-Wojtczak-Yannakakis'08]).
- The complexity, or even decidability, of optimizing reachability probability in 1-exit RMDPs (equivalently, BPA-MDPs).
  Even deciding limit-sure reachability for 1-exit RMDPs is open, although almost-sure reachability was shown decidable in P-time by [Brazdil-Brózek-Forejt-Kucera,2006].
- The previous question is a special case of optimizing termination probability in 2-exit RMDPs (which is also wide open).

More open questions

- The complexity, or even decidability, of limit-sure termination in a specific state, for a given OC-MDP.
(We know almost-sure termination is in EXPTIME & PSPACE-hard.)
- Can we approximate the optimal probability of termination in any state for a given OC-MDP in P-time? (We only know EXPTIME upper bounds.)
- [Esparza-Kiefer-Luttenberger'2010] ("Newtonian program analysis") studied analogs of Newton's method applied to MPSs for other (ω-continuous) semi-rings, beyond $[0, 1]$ or $[0, \infty]$.
Question: Can some version of Generalized Newton's Method for max/minPPSs be adapted to other semi-rings?

Some of my own related papers

- K. Etessami and M. Yannakakis. Recursive Markov chains, stochastic grammars, and monotone systems of nonlinear equations. Journal of the ACM, 56(1), 2009.
- K. Etessami and M. Yannakakis. Recursive Markov decision processes and recursive stochastic games. Journal of the ACM, 62(2), 2015.
- A. Stewart, K. Etessami, and M. Yannakakis. Upper bounds for Newton's method on monotone polynomial systems, and P-time model checking of probabilistic one-counter automata. Journal of the ACM, 64(4), 2015.
- K. Etessami, A. Stewart, and M. Yannakakis. Polynomial time algorithms for multi-type branching processes and stochastic context-free grammars. Proceedings of *STOC*, 2012. Full version: arXiv:1201.2374
- K. Etessami, A. Stewart, and M. Yannakakis. Polynomial time algorithms for Branching Markov Decision Processes and Probabilistic Min/Max Polynomial Bellman Equations. Proceedings of ICALP, 2012. Full version: arXiv:1202.4798
- K. Etessami, D. Wojtczak, and M. Yannakakis. Quasi-Birth-Death Processes, Tree-like QBDs, Probabilistic 1-Counter Automata, and Pushdown Systems. QEST'08, and Performance Evaluation, 67(9):837-857, 2010.
- T. Brazdil, V. Brozek, K. Etessami, & A. Kucera. Approximating the termination value of one-counter MDPs and stochastic games, ICALP'11 and Information and Computation, 222(2):121-138, 2013.

Other related papers accessible from my web page.