Algorithms for Branching MDPs and Branching stochastic games

Kousha Etessami

University of Edinburgh

Based on joint works with:

Alistair Stewart & Mihalis Yannakakis

U. of Edinburgh (now USC) Columbia Uni.

Cassting Workshop (ETAPS'16)
April 2016

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

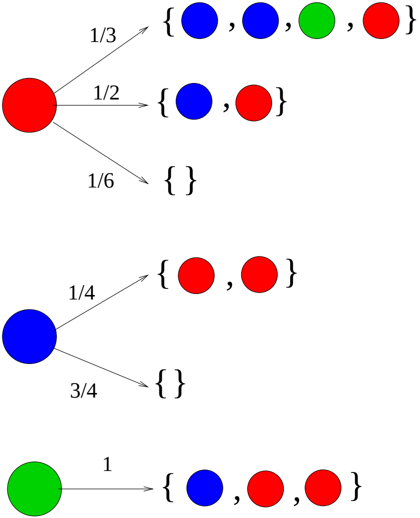

Multi-type Branching Processes (BPs) (Kolmogorov,1940s)

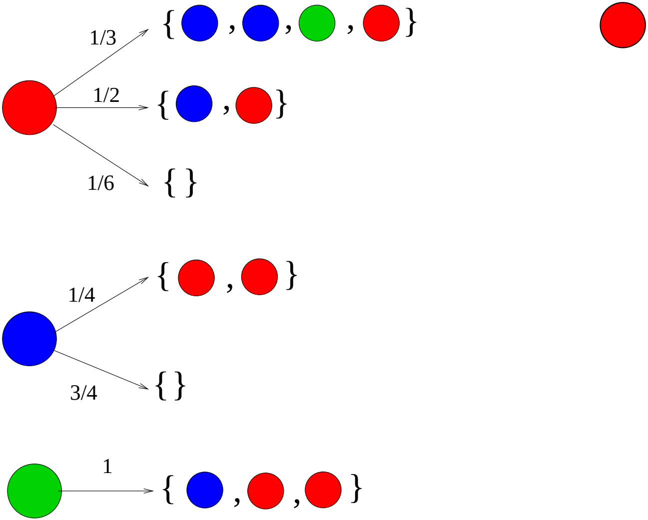

Multi-type Branching Processes (BPs) (Kolmogorov,1940s)

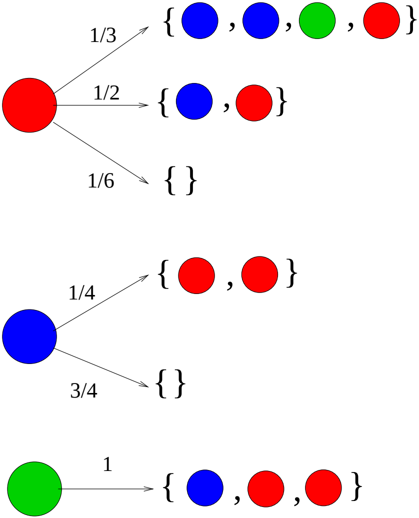

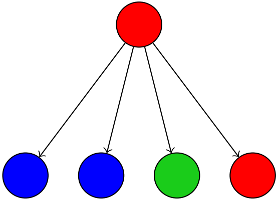

Multi-type Branching Processes (BPs) (Kolmogorov,1940s)

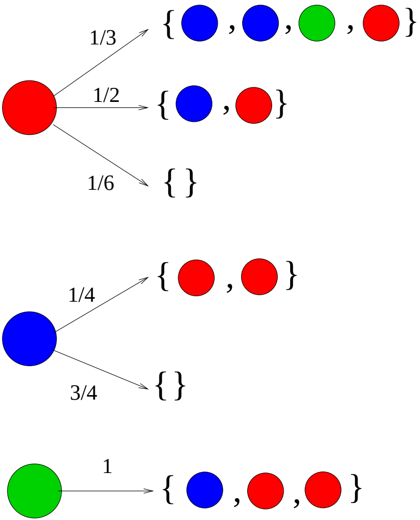

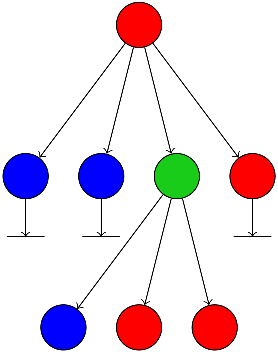

Multi-type Branching Processes (BPs) (Kolmogorov,1940s)

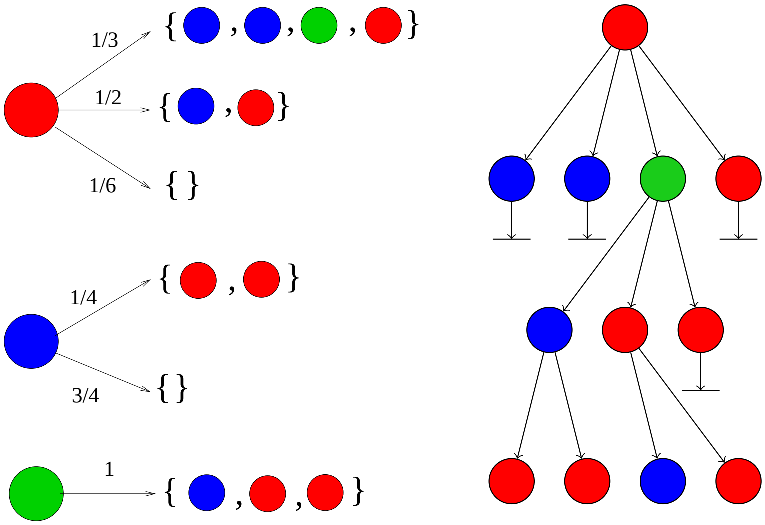

Multi-type Branching Processes (BPs) (Kolmogorov,1940s)

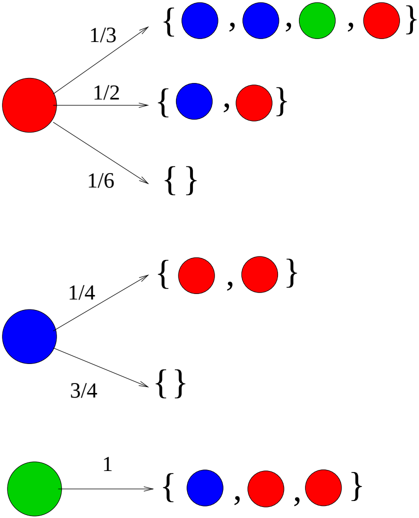

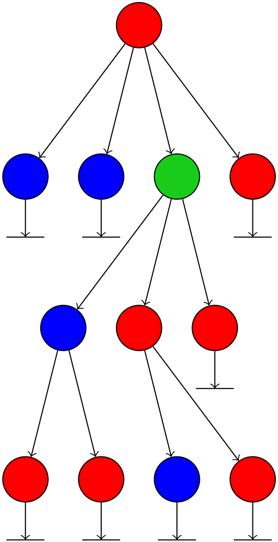

BPs are classic, fundamental, stochastic processes, studied for decades in probability theory, with many applications, eg.: population biology, nuclear chain reactions, cancer tumor models, . . .

BPs are also "intimately related" to:

- probabilistic BPPs (pBPPs)
- probablistic BPAs
- stochastic (probabilistic) Context-Free Grammars (SCFGs).
- 1-exit Recursive Markov Chains (1-RMCs)
- stateless probabilistic Pushdown Systems (stateless pPDS).

Nevertheless, even basic algorithm questions about BPs remained open until recently.

Branching Markov Decision Processes

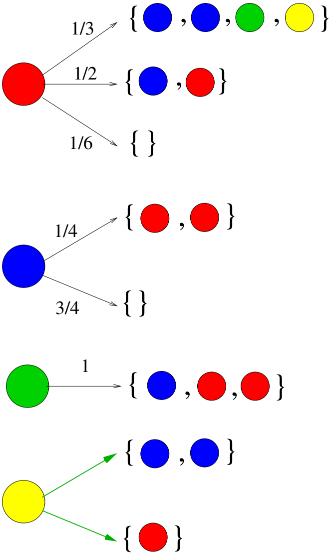

Branching Markov Decision Processes

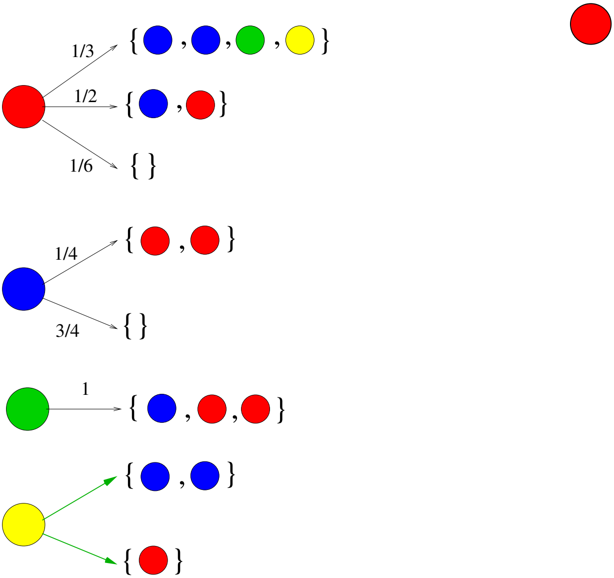

Branching Markov Decision Processes

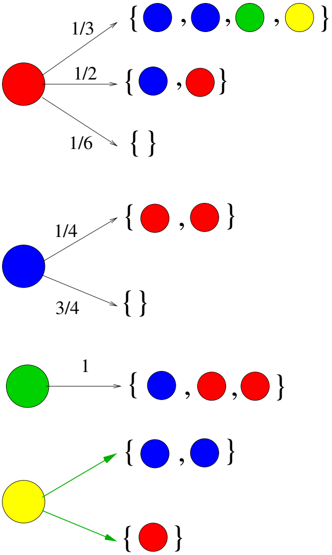

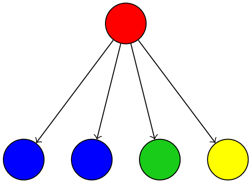

Branching Markov Decision Processes

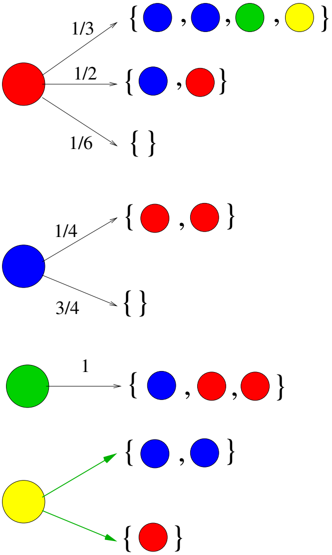

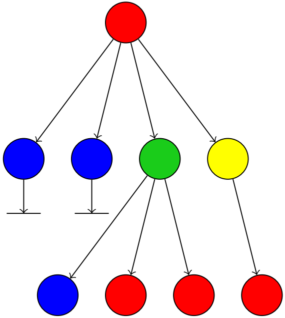

Branching Markov Decision Processes

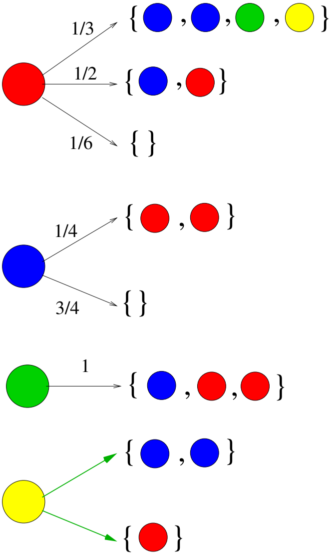

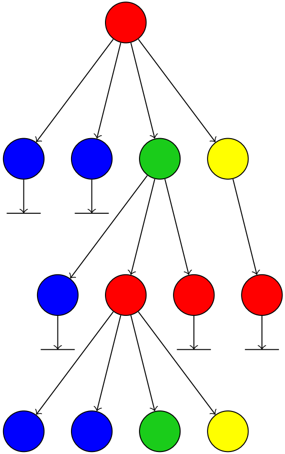

Branching Markov Decision Processes

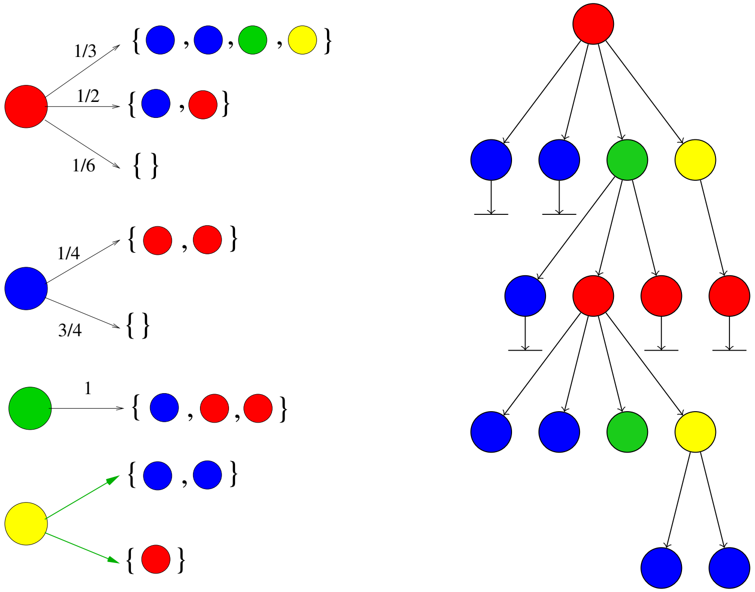

Branching Simple Stochastic Games

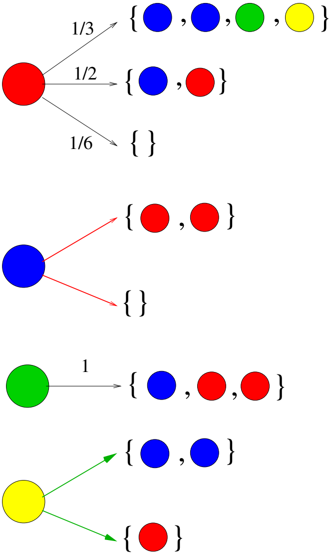

Branching Simple Stochastic Games

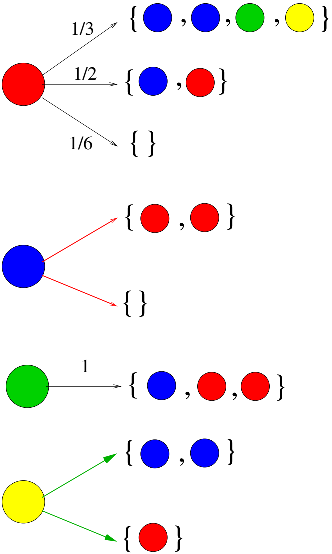

Types belonging to min:

Types belonging to max:

Multi-type Branching Processes (Kolmogorov,1940s)

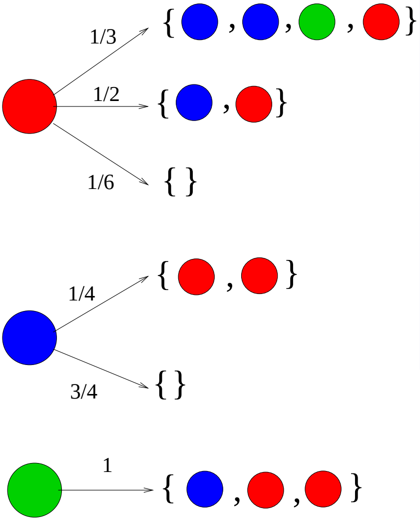

Question: What is the probability of eventual extinction, starting with one  ?

Multi-type Branching Processes (Kolmogorov,1940s)

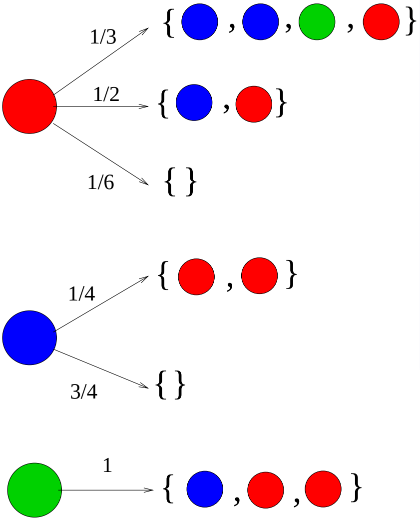

Question: What is the probability of eventual extinction, starting with one  ?

$$x_R \;=$$

Multi-type Branching Processes (Kolmogorov,1940s)

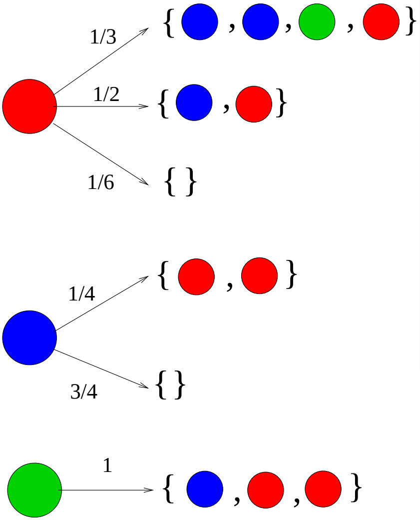

Question: What is the probability of eventual extinction, starting with one  ?

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

Multi-type Branching Processes (Kolmogorov,1940s)

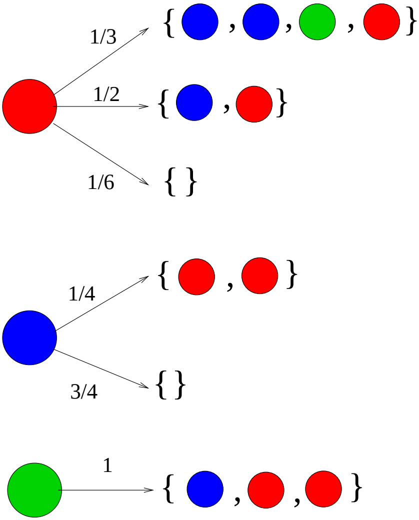

Question: What is the probability of eventual extinction, starting with one  ?

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

Multi-type Branching Processes (Kolmogorov,1940s)

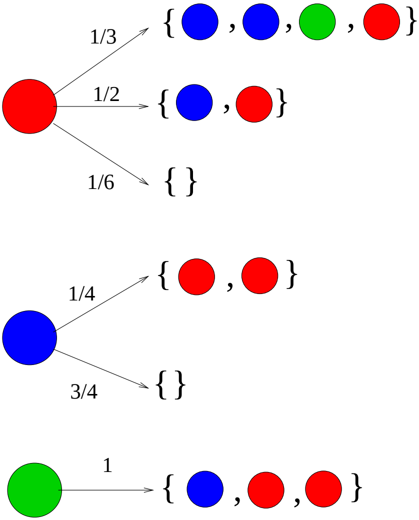

Question: What is the probability of eventual extinction, starting with one  ?

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

We get nonlinear fixed point equations:

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

Multi-type Branching Processes (Kolmogorov,1940s)

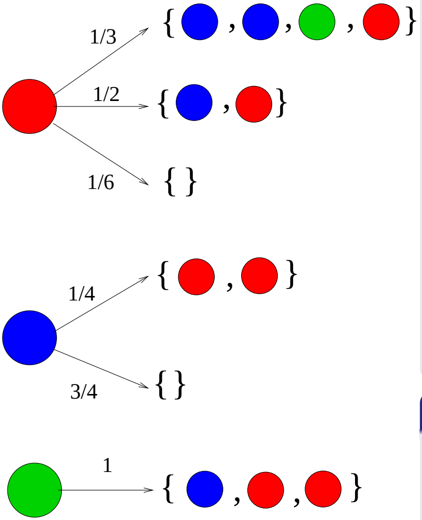

Question: What is the probability of eventual extinction, starting with one  ?

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

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

Multi-type Branching Processes (Kolmogorov,1940s)

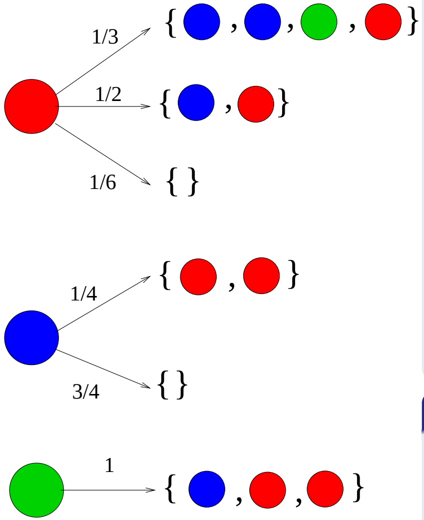

Question: What is the probability of eventual extinction, starting with one  ?

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

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

$q_R^* = 0.276$; $q_B^* = 0.769$; $q_G^* = 0.059$.

Branching Markov Decision Processes

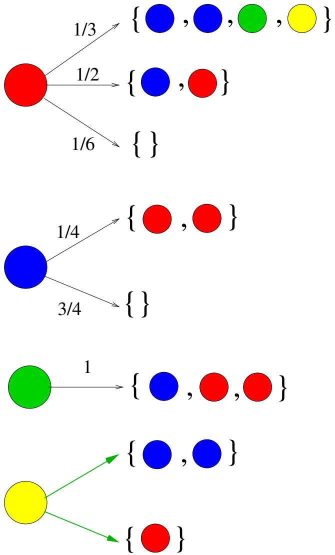

Branching Markov Decision Processes

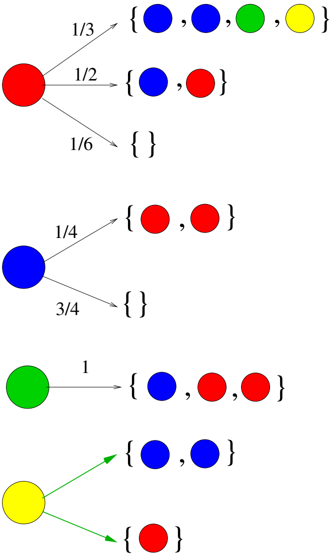

Question

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

Branching Markov Decision Processes

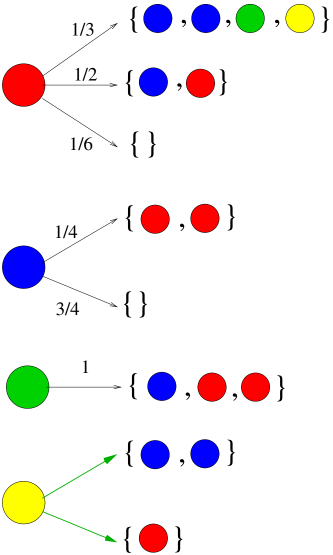

Question

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

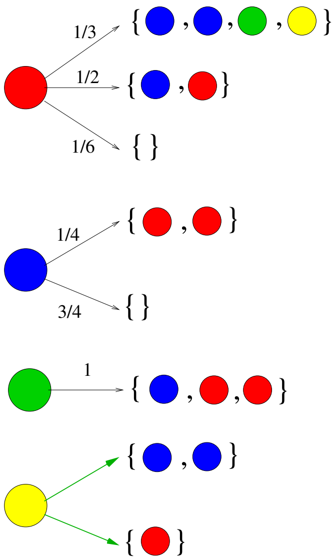

Question

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

Theorem [E.-Yannakakis'05]

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

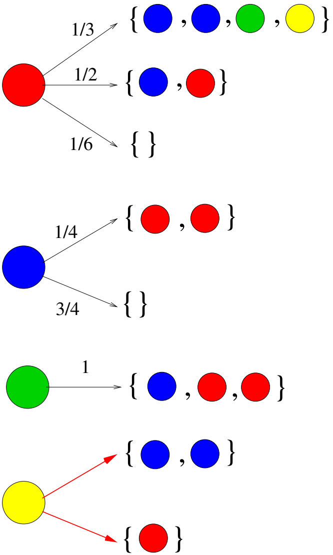

Question

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

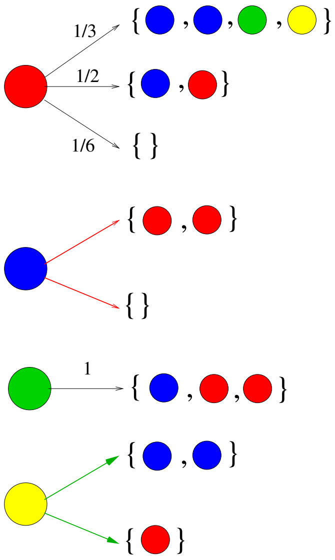

Question

What is the value of extinction, starting with one  ?

Branching Simple Stochastic Games

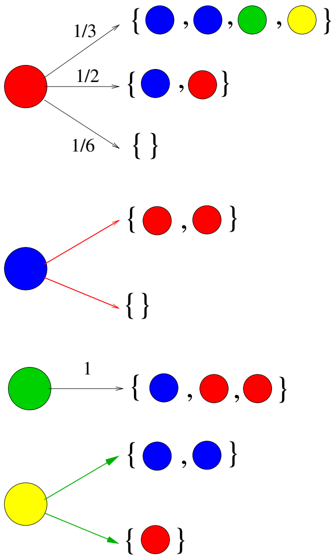

Question

What is the value of extinction, starting with one  ?

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

$$x_B = \min\{x_R^2, 1\}$$

$$x_G = x_B x_R^2$$

$$x_Y = \max\{x_B^2, x_R\}$$

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

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

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

A 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) are 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.

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

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

Proposition. [E.-Yannakakis'05]

- *Every max-minPPS, $x = P(x)$ has a least fixed point, $q^* \in [0,1]^n$.*
- $q^* = \lim_{k\rightarrow\infty} P^k(\mathbf{0})$.
- *q^* is the vector of optimal extinction probabilities (values) for the BMDP (the BSSG).*

Question

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

Static optimal strategies for BMDP/BSSG extinction

Theorem ([E.-Yannakakis'05])

For any BSSG extinction game, both players have static optimal strategies for maximizing (minimizing) extinction probability.

(However, computing an optimal strategy, even for BMDPs, is PosSLP-hard ([E.-Yannakakis'05,'09]); and of course, for BSSGs this is also as hard as solvings Condon's finite-state SSGs.)

A static strategy is one that, for every type belonging to a player, always chooses the same rule (i.e., it is deterministic, memoryless, and "context-oblivious".)

Question

Can we compute an ϵ-optimal strategy for the controller maximizing/minimizing extinction probability in a BMDP in P-time?

P-time approximation for BMDPs and max/minPPSs

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

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,ICALP'12])

Moreover, we can compute an ϵ-optimal static strategy for (maximizing/minimizing) extinction probability 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**:

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

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 on PPSs

We can 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 converges monotonically to the LFP $\mathbf{q}^*$ for PPSs, and for more general Monotone Polynomial Systems (MPSs).

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 $q^* = 1$, 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,STOC'12])

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 approximating q^*.

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, and (we show)

$$\text{(spectral radius of } P'(q^*)\text{)} \equiv \varrho(P'(q^*)) < 1$$

- 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

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

- $\|(I - P'(q^*))^{-1}\|_\infty$ is tied 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.

   Theorem ([E.-Yannakakis'06]) Checking $q_i^* = 1$ is decidable in P-time using linear programming.
   (Reduces, with some work, to a spectral radius optimization problem for non-negative square matrices.)

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.

   Theorem ([E.-Yannakakis'06]) Checking $q_i^* = 1$ is decidable in P-time using linear programming.
   (Reduces, with some work, to a spectral radius optimization problem for non-negative square matrices.)

2. On the resulting system of 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.

   Theorem ([E.-Yannakakis'06]) Checking $q_i^* = 1$ is decidable in P-time using linear programming.
   (Reduces, with some work, to a spectral radius optimization problem for non-negative square matrices.)

2. On the resulting system of 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

$$\|\mathbf{q}^* - \mathbf{x}^{(4|P|+j+1)}\|_\infty \leq 2^{-j}.$$

Thus, algorithm runs in time polynomial in $|P|$ and j.

Proof outline: some key lemmas

$(\mathbf{1} - \mathbf{q}^*)$ is the vector of pessimal survival probabilities.

Lemma

If $\mathbf{q}^* - \mathbf{x}^{(k)} \leq \lambda(\mathbf{1} - \mathbf{q}^*)$ *for some* $\lambda > 0$*, then* $\mathbf{q}^* - \mathbf{x}^{(k+1)} \leq \frac{\lambda}{2}(\mathbf{1} - \mathbf{q}^*)$.

Lemma

For any Max(Min) PPS with LFP $\mathbf{q}^*$*, such that* $\mathbf{0} < \mathbf{q}^* < \mathbf{1}$*, for any* i*,*

$$q_i^* \leq 1 - 2^{-4|P|}.$$

Qualitative and Quantitative extinction problems for BSSGs

Theorem ([E.-Yannakakis'06])

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

And, it is at least as hard as computing the exact value for 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***).*

Optimal Reachability problem for BMDPs

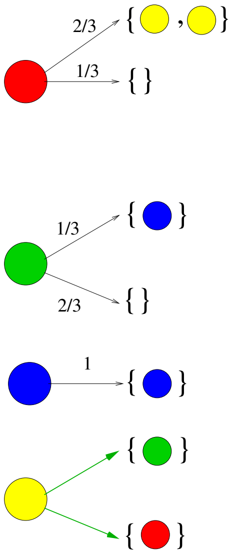

Optimal Reachability problem for BMDPs

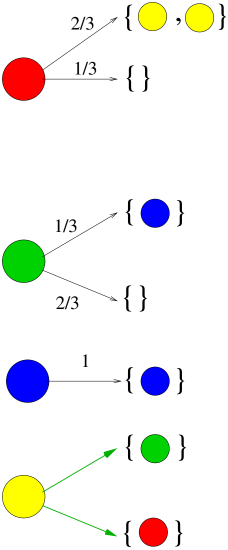

Question

What is the maximum (actually supremum) probability of reaching , starting with one  ?

Optimal Reachability problem for BMDPs

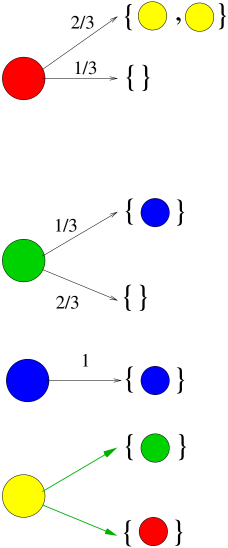

Same Question (rephrased)

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

Optimal Reachability problem for BMDPs

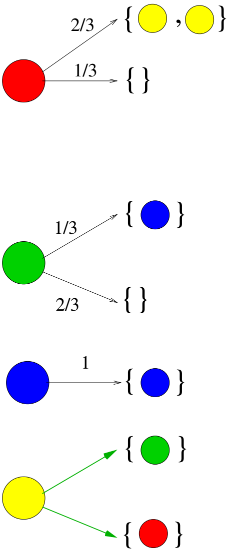

Same Question (rephrased)

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

$$y_R \quad =$$

Optimal Reachability problem for BMDPs

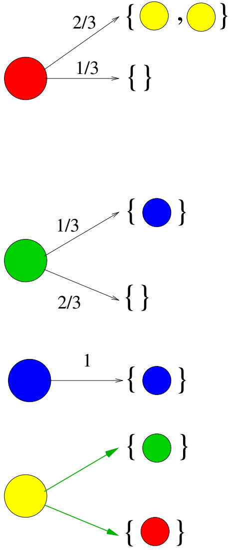

Same Question (rephrased)

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

$$
\begin{aligned}
y_R &= \frac{2}{3} y_Y y_Y + \frac{1}{3} \\
y_G &= \frac{2}{3} \\
y_Y &= \min\{y_G, y_R\}
\end{aligned}
$$

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

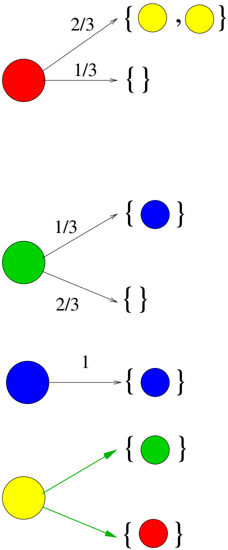

Question

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 Generalized Newton's Method.

Algorithm for GFP of max/minPPSs

1. Find and remove all variables x_i such that $g_i^* = 1$.
   (This can be done in P-time, by qualitative analysis of $\mathbf{y} = Q(\mathbf{y})$.)
2. Interestingly, we do not need to eliminate the variables x_i such that $g_i^* = 0$. (And we do not want to eliminate variables with $q_i^* = 0$.)
3. On the resulting system of equations, run Generalized Newton's Method, starting from $\mathbf{0}$. After each iteration, round down to a multiple of 2^{-h}.
4. Amazingly this works! Note the very subtle difference with the algorithm for approximating the LFP of the same max/minPPS.

Algorithm for GFP of max/minPPSs

1. Find and remove all variables x_i such that $g_i^* = 1$.
   (This can be done in P-time, by qualitative analysis of $\mathbf{y} = Q(\mathbf{y})$.)
2. Interestingly, we do not need to eliminate the variables x_i such that $g_i^* = 0$. (And we do not want to eliminate variables with $q_i^* = 0$.)
3. On the resulting system of equations, run Generalized Newton's Method, starting from $\mathbf{0}$. After each iteration, round down to a multiple of 2^{-h}.
4. Amazingly this works! Note the very subtle difference with the algorithm for approximating the LFP of the same max/minPPS.

Theorem [E.-Stewart-Yannakakis'ICALP15]

Given a max/minPPS $\mathbf{x} = P(\mathbf{x})$ with GFP $\mathbf{0} \leq \mathbf{g}^* < \mathbf{1}$, if we apply rounded GNM starting at $\mathbf{x}^{(0)} = \mathbf{0}$, using $h := 4|P| + j + 1$ bits of precision, then

$$\|\mathbf{g}^* - \mathbf{x}^{(4|P|+j+1)}\|_\infty \leq 2^{-j}.$$

Thus, algorithm runs in time polynomial in $|P|$ and j.

Qualitative & quantitative reachability for BSSGs

Theorem [E.-Stewart-Yannakakis'ICALP15]

- The value of a BSSG reachability game is captured by the GFP of max-minPPS.
- The player minimizing reachability probability has a static positional optimal strategy. But, already for BMDPs, the player maximizing it may have no optimal strategy at all, only ϵ-optimal (randomized-static, or deterministic-memoryful) strategies.
- We can approximate the value, and compute ϵ-optimal stratgies, for a BSSG reachabilty 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 an answer all qualitative questions in P-time for BSSG reachability games, including compute qualitative-optimal (not static) strategies.
  (**Note:** This contrasts sharply with qualitative extinction, which is as hard as computing the value of finite-state SSGs [E.-Yannakakis'05].)

Conclusion

We have established P-time algorithms for a number of fundamental quantitative and qualitative analysis problems for Branching MDPs (and related results for Branching SSGs), including for:

- optimal extinction probabilities
- optimal reachability probabilities
- optimal expected total progeny size and "weight" ([E.-Wojtczak-Yannakakis'08], which I didn't speak about.)

Many 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, compute: $\sup_{\sigma \in \mathrm{S}trategy} Pr(\mathrm{Tree}_c^{\sigma}(M) \models \varphi)$.
  (Our results only imply computability for fragments of CTL.)
- Multi-player branching stochastic games? We know nothing!

Papers

- 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 and M. Yannakakis. Recursive Markov decision processes and recursive stochastic games. Journal of the ACM, 62(2):169, 2015.
- 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, A. Stewart, and M. Yannakakis. Greatest Fixed Points of Probabilistic Min/Max Polynomial Bellman Equations, and Reachability for Branching Markov Decision Processes. Proceedings of ICALP, 2015. Full version: arXiv:1502.05533
- 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 Concurrent Stochastic Games. LMCS, 4(4), 2008.
- K. Etessami, D. Wojtczak, and M. Yannakakis. Recursive Stochastic Games with Positive Rewards. ICALP'08. Full version:

`http://homepages.inf.ed.ac.uk/kousha/`

Other related papers accessible from my web page.