Algorithms for Branching Markov Decision Processes

Kousha Etessami

University of Edinburgh

Based on joint works with:

Alistair Stewart & Mihalis Yannakakis

U. of Edinburgh Columbia Uni.

RP 2014

Oxford U.

September 22nd, 2014

Multi-type Branching Processes (Kolmogorov,1940s)

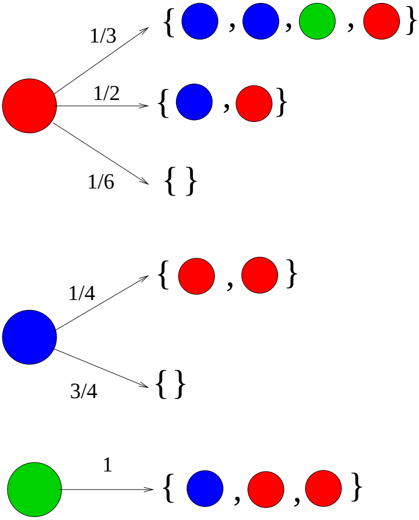

Multi-type Branching Processes (Kolmogorov,1940s)

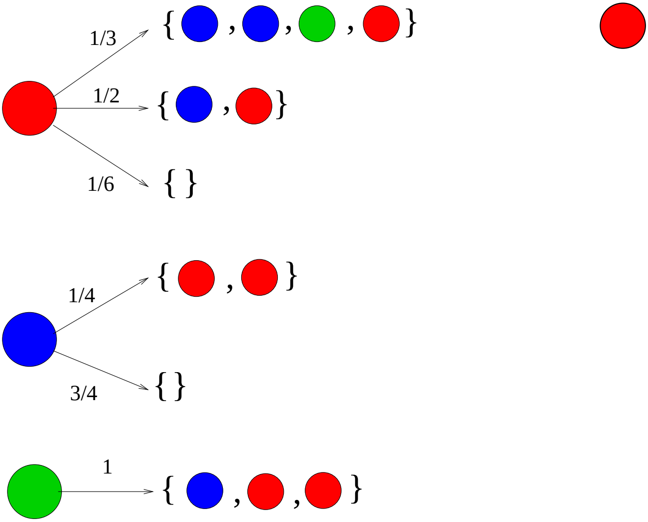

Multi-type Branching Processes (Kolmogorov,1940s)

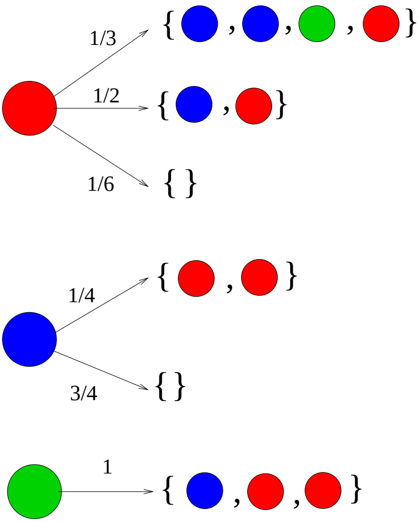

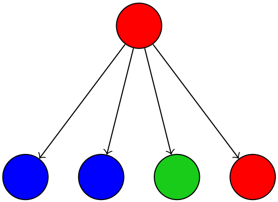

Multi-type Branching Processes (Kolmogorov,1940s)

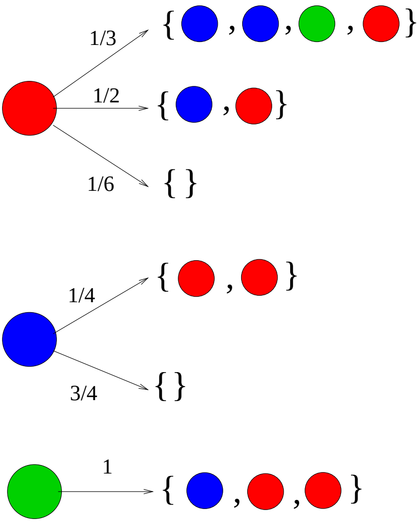

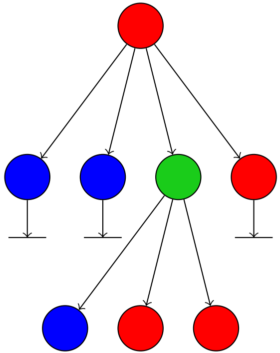

Multi-type Branching Processes (Kolmogorov,1940s)

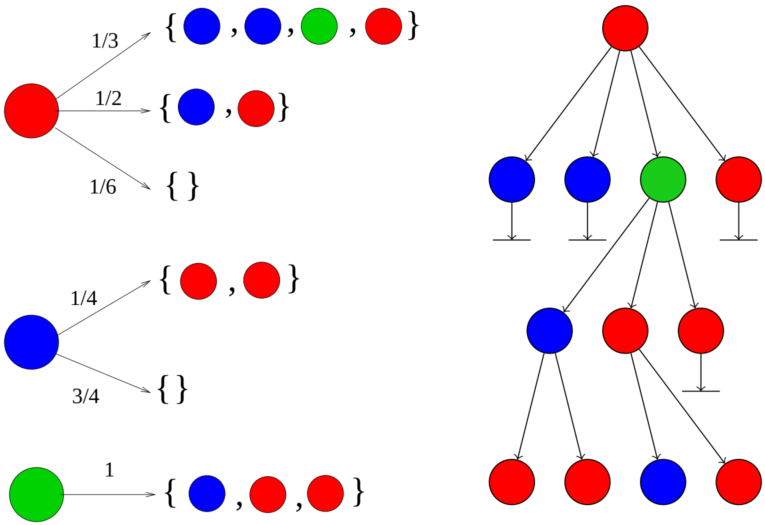

Multi-type Branching Processes (Kolmogorov,1940s)

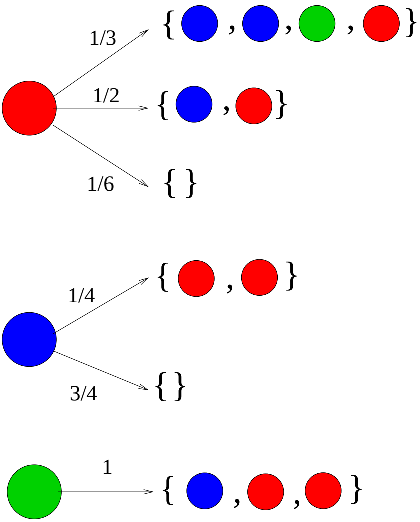

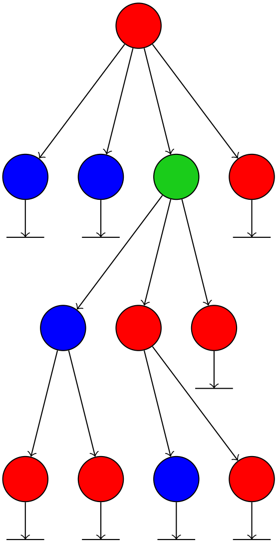

Multi-type Branching Processes (Kolmogorov 1940s)

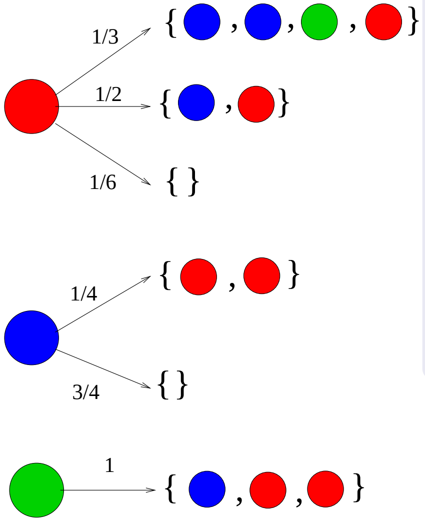

Question: What is the probability of eventual extinction, starting with one  ?

Multi-type Branching Processes (Kolmogorov 1940s)

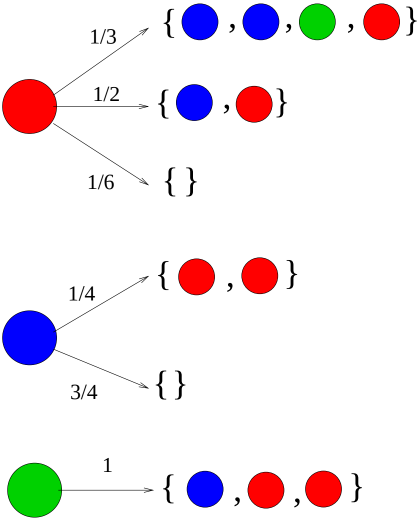

Question: What is the probability of eventual extinction, starting with one  ?

$$x_R \quad =$$

Multi-type Branching Processes (Kolmogorov 1940s)

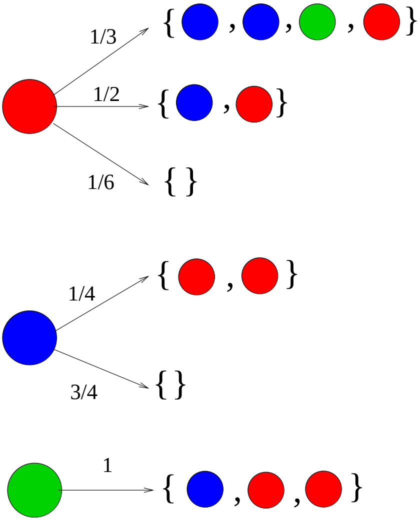

Question: What is the probability of eventual extinction, starting with one  ?

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

Multi-type Branching Processes (Kolmogorov 1940s)

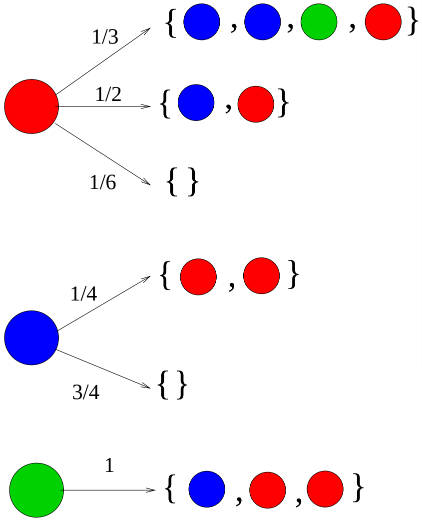

Question: What is the probability of eventual extinction, starting with one  ?

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

Multi-type Branching Processes (Kolmogorov 1940s)

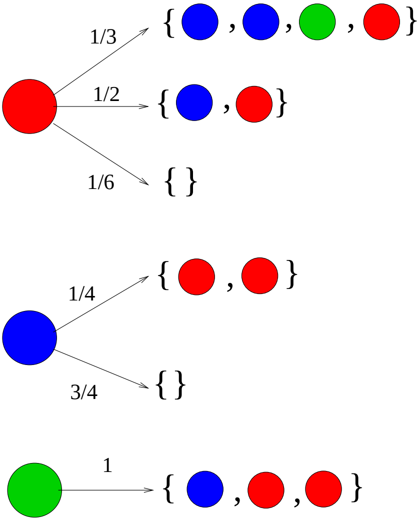

Question: What is the probability of eventual extinction, starting with one  ?

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

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

$$x_G = x_B x_R^2$$

We get nonlinear fixed point equations:

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

Multi-type Branching Processes (Kolmogorov 1940s)

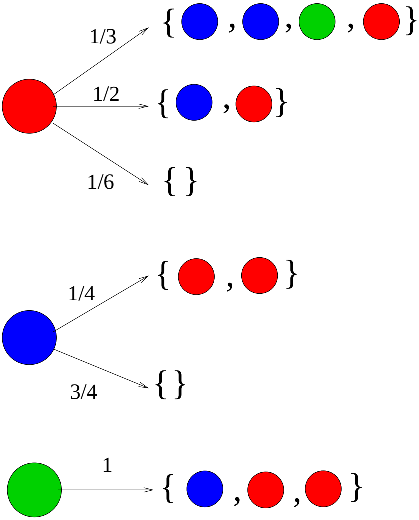

Question: What is the probability of eventual extinction, starting with one  ?

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

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

$$x_G = x_B x_R^2$$

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

Multi-type Branching Processes (Kolmogorov 1940s)

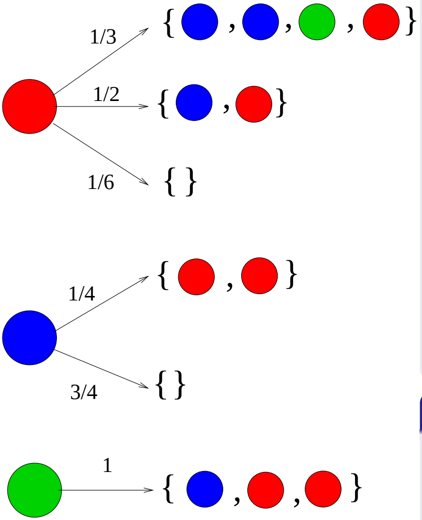

Question: What is the probability of eventual extinction, starting with one  ?

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

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 $\bar{\mathbf{x}} = P(\bar{\mathbf{x}})$.
$q_R^* = 0.276$; $q_B^* = 0.769$; $q_G^* = 0.059$.

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

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

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

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

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

Every multi-type Branching Process (BP) with n types, and every SCFG with n nonterminals, corresponds to a PPS, and vice-versa.

Basic properties of PPSs, $\mathbf{x} = P(\mathbf{x})$

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

Proposition

- *A PPS, $x = P(x)$ has a least fixed point, $q^* \in [0,1]^n$.*
  (q^* can be irrational.)
- $q^* = \lim_{k\to\infty} P^k(\mathbf{0})$.
- *q^* is vector of extinction/termination probabilities for the BP (SCFG).*

Question

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

First considered by Kolmogorov & Sevastyanov (1940s).

Newton's method

Newton's method

Seeking a solution to $F(\mathbf{x}) = 0$, we start at a guess $\mathbf{x}^{(0)}$, 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 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,'09], we gave no upper bounds on # of iterations needed for PPSs or MPSs.
- We proved hardness results (PosSLP-hardness) for obtaining any nontrivial approximation of the LFP of MPSs for recursive Markov chains.

What is Newton's worst case behavior for PPSs?

[Esparza,Kiefer,Luttenberger,'07,'10] studied Newton's method on MPSs further:

- Gave bad examples of PPSs, $\mathbf{x} = P(\mathbf{x})$, where $q^* = 1$, requiring exponentially many iterations, as a function of the encoding size $|P|$ of the equations, to converge to within additive error $< 1/2$.
- For strongly-connected equation systems they gave an exponential upper bound in $|P|$.
- But they gave no upper bounds for arbitrary PPSs or MPSs in terms of $|P|$.

What is Newton's worst case behavior for PPSs?

[Esparza,Kiefer,Luttenberger,'07,'10] studied Newton's method on MPSs further:

- Gave bad examples of PPSs, $\mathbf{x} = P(\mathbf{x})$, where $q^* = 1$, requiring exponentially many iterations, as a function of the encoding size $|P|$ of the equations, to converge to within additive error $< 1/2$.
- For strongly-connected equation systems they gave an exponential upper bound in $|P|$.
- But they gave no upper bounds for arbitrary PPSs or MPSs in terms of $|P|$.

Recently in [Stewart-E.-Yannakakis'13], we gave a matching exponential upper bound in $|P|$ for arbitrary PPSs and MPSs.

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.

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 ([ESY'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

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

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 ([ESY'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...

Branching Markov Decision Processes

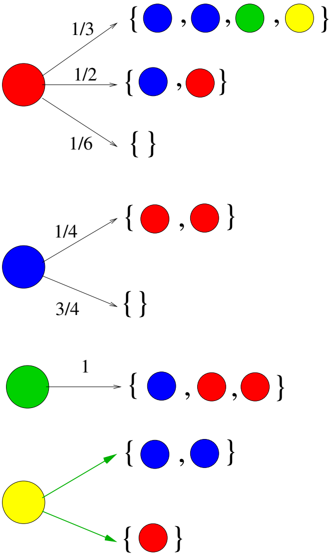

Branching Markov Decision Processes

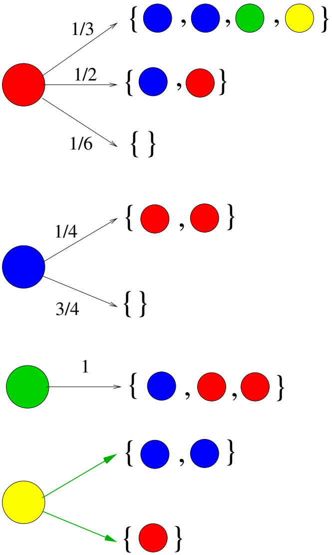

Branching Markov Decision Processes

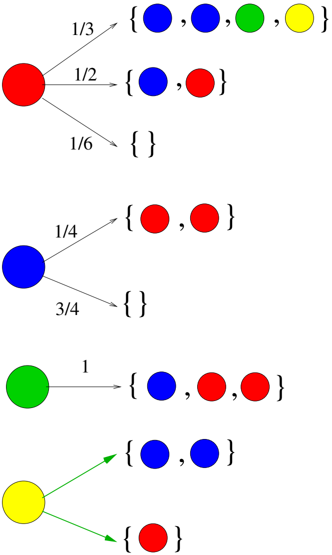

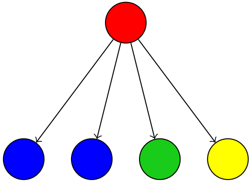

Branching Markov Decision Processes

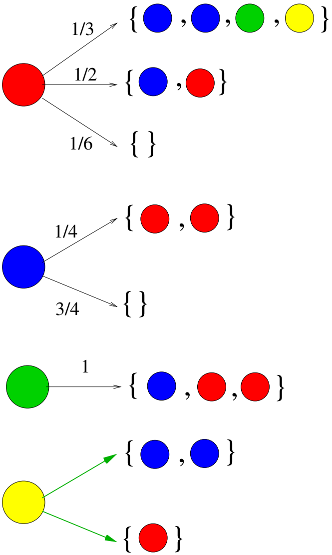

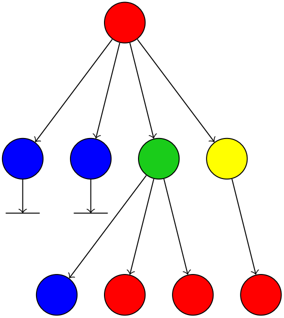

Branching Markov Decision Processes

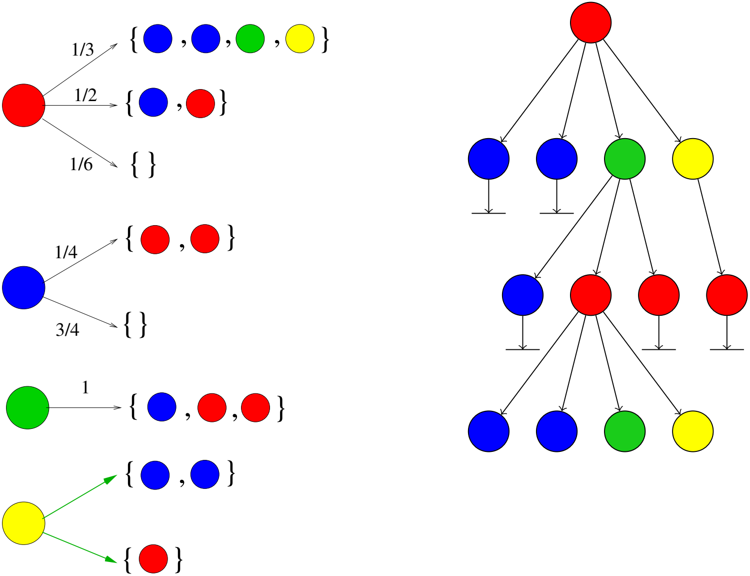

Branching Markov Decision Processes

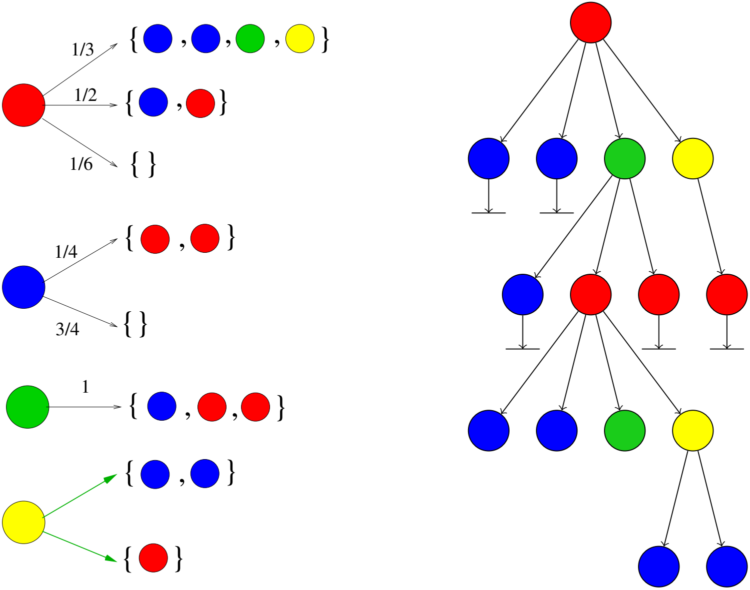

Branching Markov Decision Processes

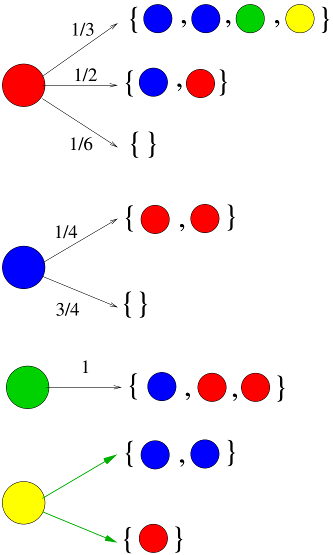

Question

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

Branching Markov Decision Processes

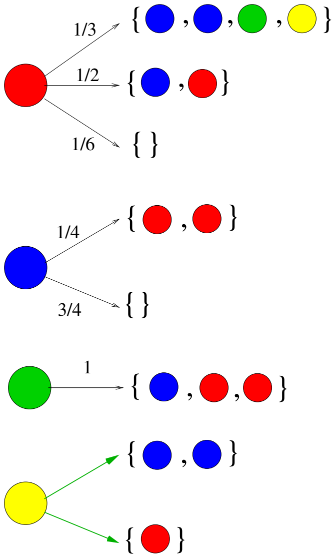

Question

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

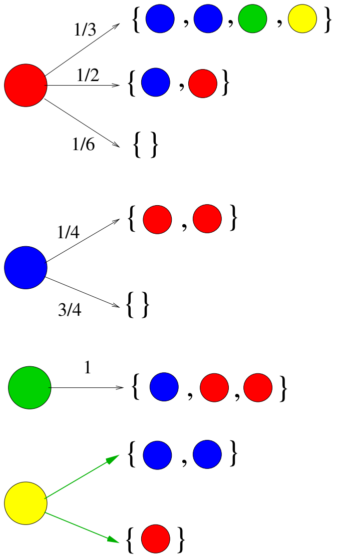

Question

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

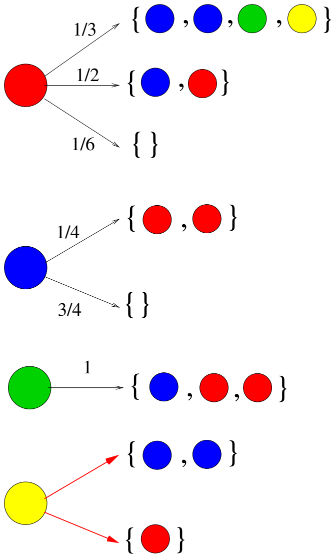

Question

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

Maximum Probabilistic Polynomial Systems of Equations

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.

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, $q^* \in [0,1]^n$.*
- $q^* = \lim_{k \to \infty} P^k(\mathbf{0})$.
- *q^* is vector of optimal extinction probabilities for the BMDP.*

Question

Can we compute the probabilities q^* efficiently (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.

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 a linear (first-order Taylor) approximation of $P(x)$.

Generalised Newton's method

Linearisation

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 linearisation, $P^y(x)$, 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

Linearisation

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 linearisation, $P^y(x)$, 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

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 $q_i^* = 1$. ($q_i^* = 1$ decidable in P-time using LP [E.-Yannakakis'06]: reduces 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 $q_i^* = 1$.
   ($q_i^* = 1$ decidable in P-time using LP [E.-Yannakakis'06]: reduces 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 $q_i^* = 1$.
   ($q_i^* = 1$ decidable in P-time using LP [E.-Yannakakis'06]: reduces 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 [ESY'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,

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

Branching Simple Stochastic Games

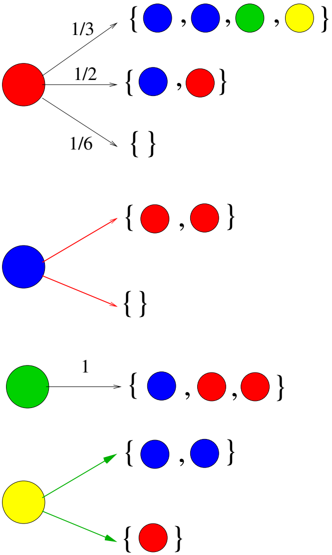

Branching Simple Stochastic Games

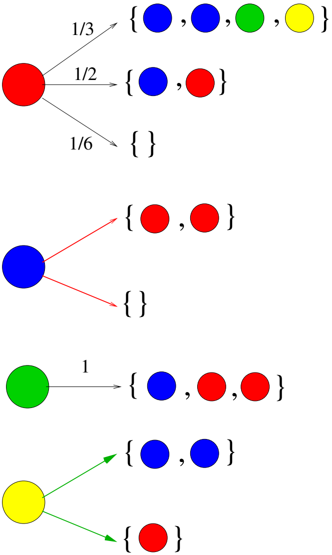

Types belonging to min:

Types belonging to max:

Branching Simple Stochastic Games

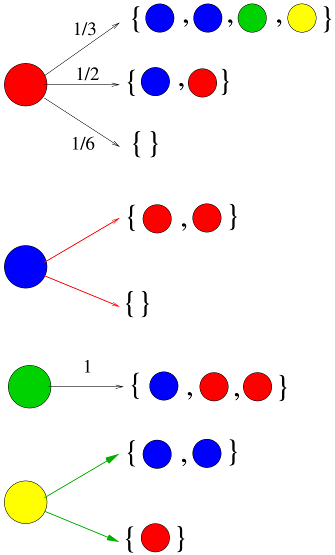

Question

What is the value of extinction, starting with one ?

Branching Simple Stochastic Games

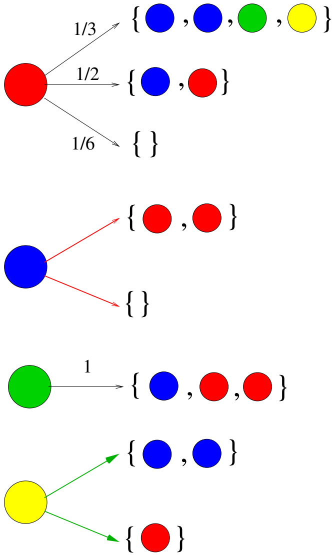

Question

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

Fact [E.-Yannakakis'05]

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

Qualitative and Quantitative problems for BSSGs

Theorem ([E.-Yannakakis'05])

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

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

Theorem ([E.-Yannakakis'06])

Given a BSSG, deciding if the extinction value is $q_i^* = 1$ is in **NP** $\cap$ **coNP**, *& is at least as hard as computing the exact value for a finite-state SSG.*

Theorem ([ESY'12])

Given a BSSG, and given $\epsilon > 0$, we can compute a vector $v \in [0,1]^n$, such that $\|v - q^*\|_\infty \leq \epsilon$, in **FNP** *(and in* **PLS***).*

Optimal Reachability problem for BMDPs

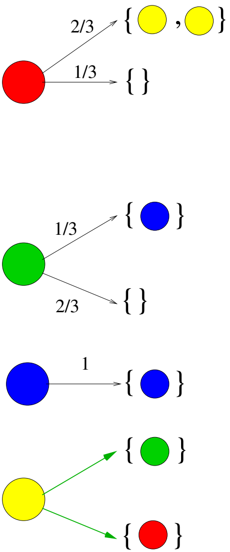

Optimal Reachability problem for BMDPs

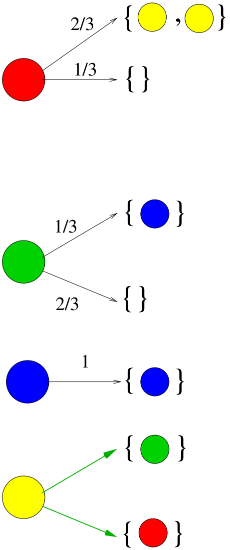

Question

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

Optimal Reachability problem for BMDPs

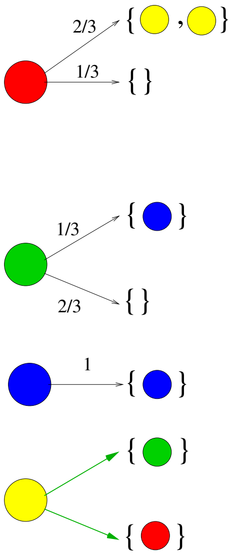

Same Question (rephrased)

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

Optimal Reachability problem for BMDPs

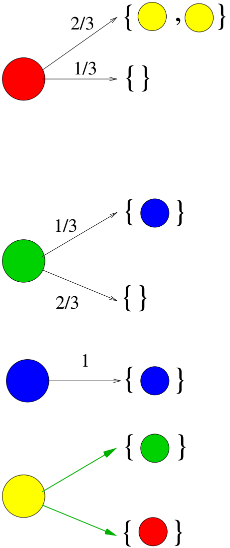

Same Question (rephrased)

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

$$y_R \quad =$$

Optimal Reachability problem for BMDPs

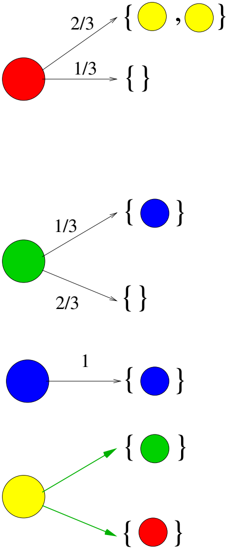

Same Question (rephrased)

What is the minimum 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. [ESY'14]

The maximum 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

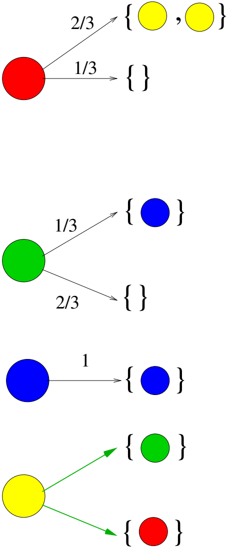

Question

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. [ESY'14]

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, 2014])

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 [ESY'14]

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.

Conclusion

- We have established P-time algorithms for a number of fundamental qualitative and quantative (approximate) analysis problems for Multi-type Branching Processes and for Branching MDPs, including for:
  - optimal extinction probabilities
  - optimal reachability probabilities
  - optimal expected progeny size ([E.-Wojtczak-Yannakakis'08])
- These algorithms also yield FNP (and PLS) complexity bounds for Branching Simple Stochastic Games with the same objectives.
- Many open questions still remain for these and related infinite-state recursive stochastic models and stochastic games.

Papers

- K. Etessami, A. Stewart, and M. Yannakakis.
  Polynomial time algorithms for multi-type branching processes and stochastic context-free grammars.
  Proceedings of *STOC'12*, pp. 579-588, 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'12*, pp. 314-326, 2012. Full version: arXiv:1202.4798
- K. Etessami, A. Stewart, and M. Yannakakis.
  Greatest Fixed Points of Probabilistic Min/Max Polynomial Bellman Equations, and Optimal Reachability for Branching MDPs.
  Forthcoming, 2014.

Other related papers accessible on my web page.