Adding Recursion to Markov Chains, Markov Decision Processes, and Stochastic Games: Algorithms and Complexity

Kousha Etessami

LFCS, School of Informatics
University of Edinburgh

QEST'11 Tutorial
Aachen, September 2011

Overview

- It turns out that many important classes of countable infinite-state stochastic processes can be captured by adding a natural recursion feature to finite-state Markov chains.
  Adding recursion also provides a natural abstract model of probabilistic procedural programs.
- Equivalently, such models can be captured by probabilistic extensions to classic infinite-state automata-theoretic models like context-free grammars, pushdown automata, and one-counter automata.
- The algorithmic theory of these recursive stochastic models, and their extensions to Markov decision processes and stochastic games, has turned out to be an extremely rich subject.
- This is the theory I will survey during this tutorial.

What is a Recursive Graph?

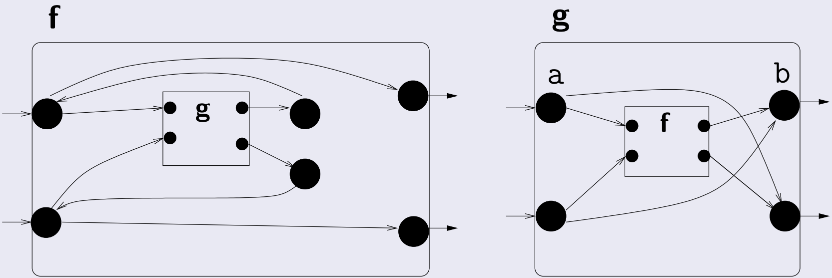

- Question: Is it possible to reach **b** from **a**?
- Such questions can be answered in P-time. Worst-case cubic time, and linear time if # of entries or # of exits of each component is bounded by a fixed constant. (Reducible to And/Or-graph reachability.)
- More generally, we can model check LTL and ω-regular properties of recursive state machines in the same time in the size of the model.
- Such models are studied heavily in program analysis and verification.

What is a Recursive Markov Chain?

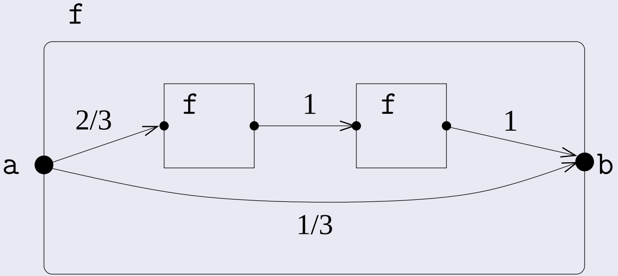

- Question: What is the probability of eventually reaching **b** from **a**?
- Is there an efficient algorithm for computing such probabilities?
- Recall standard algorithm for computing hitting probabilities for finite-state MCs: solve a linear system of equations.
- More general model checking question: what is the probability that a run of the RMC satisfies a given LTL or ω-regular property?

Let's calculate the termination probability for a RMC

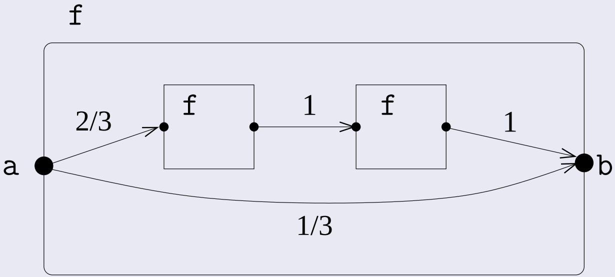

- Let x be the (unknown) probability

Let's calculate the termination probability for a RMC

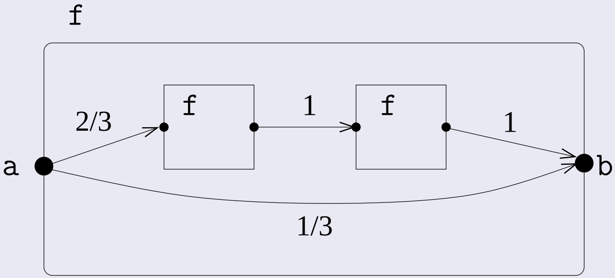

- Let x be the (unknown) probability
- An equation for x: $x = (2/3)x^2 + 1/3$

Let's calculate the termination probability for a RMC

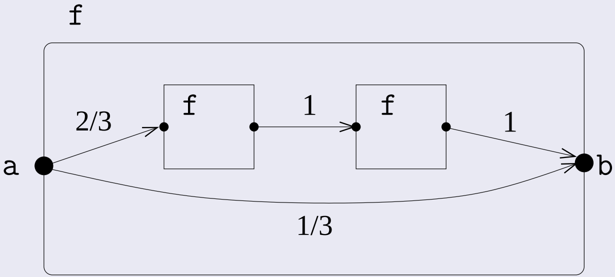

- Let x be the (unknown) probability
- An equation for x: $x = (2/3)x^2 + 1/3$
- This nonlinear equation has two solutions: $x = 1/2$ and $x = 1$.

Let's calculate the termination probability for a RMC

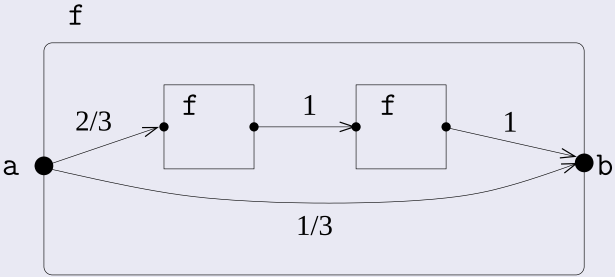

- Let x be the (unknown) probability
- An equation for x: $x = (2/3)x^2 + 1/3$
- This nonlinear equation has two solutions: $x = 1/2$ and $x = 1$.
- The least solution, call it the least fixed point (LFP), is $1/2$.
  Fact: The LFP, $1/2$, is the probability of termination.

A "structurally identical" example

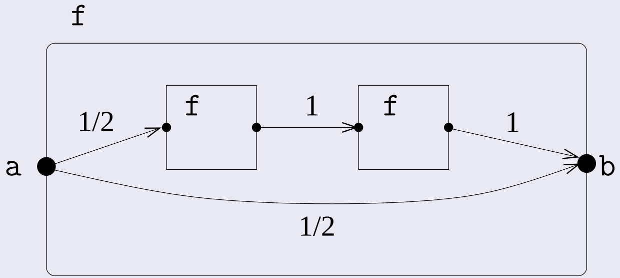

- Equation: $x = (1/2)x^2 + 1/2$.
- Two (degenerate) solutions: $x = 1$ and $x = 1$. LFP: $x^* = 1$.
- So, even for *structurally identical* RMCs, almost sure termination depends on actual probabilities.

A "structurally identical" example

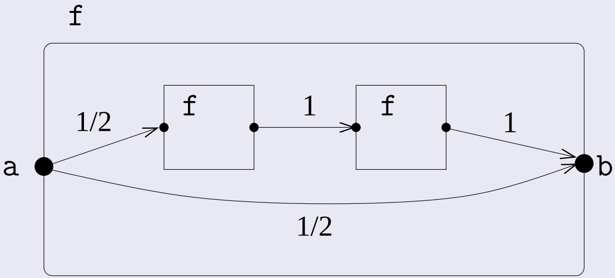

- Equation: $x = (1/2)x^2 + 1/2$.
- Two (degenerate) solutions: $x = 1$ and $x = 1$. LFP: $x^* = 1$.
- So, even for *structurally identical* RMCs, almost sure termination depends on actual probabilities.
- This does not happen for finite-state MCs.

another example: irrational probabilities

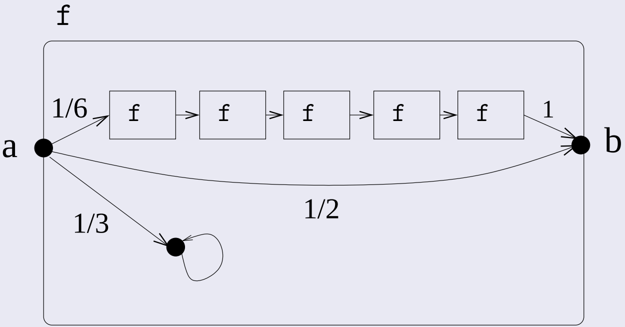

- Equation: $x = (1/6)x^5 + 1/2$.
- This equation is not "solvable by radicals" (LFP is $\sim$ 0.50550123 . . .).
- For finite-state MCs, hitting probabilities are "concise" rationals.

RMCs subsume many well-studied infinite-state stochastic models

Models subsumed by RMCs:

- Stochastic Context-Free Grammars (= 1-exit RMCs)
- Multi-Type Branching Processes (Kolmogorov 1940's)
  (extinction probability = 1-exit RMC termination probability)
- (discrete-time) Quasi-Birth-Death Processes
  (= 1-box RMCs = 1-counter probabilistic automata)
- Backbutton Processes (=special subclass of 1-exit RMCs)

Models equivalent to RMCs:

- probabilistic Pushdown Automata
- Tree-Like-QBDs

Many important computational problems for all these models boil down to computation of termination probabilities.

RMCs, more formally

An RMC, $A = \langle A_1, \ldots, A_k \rangle$ consists of **components** $A_1, \ldots, A_k$, with each A_i given by:

- A set N_i of **nodes**, and a set B_i of **boxes**.
  A mapping $Y_i : B_i \mapsto \{1, \ldots, k\}$ of each box to a component.
- A set $En_i \subseteq N_i$ of **entry nodes**, and a set $Ex_i \subseteq N_i$ of **exit nodes**.
- A **transition relation** δ_i, where each $(u, p_{u,v}, v) \in \delta_i$ has the form:
  - $u \in N_i$, or $u = (b, ex)$ where $b \in B_i$ and $ex \in Ex_{i_b}$.
  - $v \in N_i$, or $v = (b, en)$ where $b \in B_i$ and $en \in En_{i_b}$.
  - $p_{u,v} \in \mathbb{R}_{\geq 0}$,
    and $\sum_v p_{u,v} = 1$ or $= 0$.

How do we get a countable-state Markov chain from this?

The infinite-state Markov chain underlying a RMC

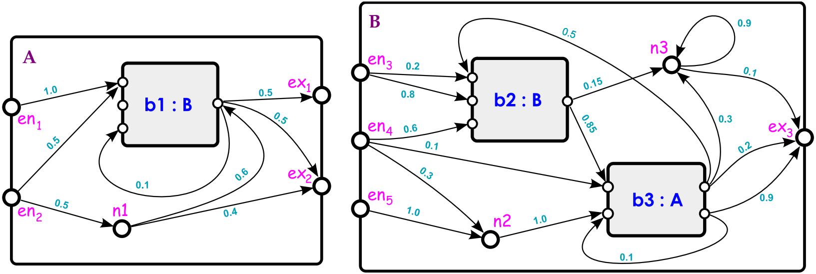

- "Expanding" the boxes of an RMC, forever, defines a countable MC.
- States of the MC have the form $s = \langle b_1 b_2 \ldots b_r, u \rangle$, where b_i's are boxes (the "call stack"), and u is a node.

Termination probabilities: What is the probability that starting at vertex u with empty call stack (i.e., $\langle \epsilon, u \rangle$), we eventually terminate at exit ex (i.e., reach state $\langle \epsilon, ex \rangle$)? Denote these (unknown) probabilities by $x_{u,ex}$.

The nonlinear system of equations for RMC termination

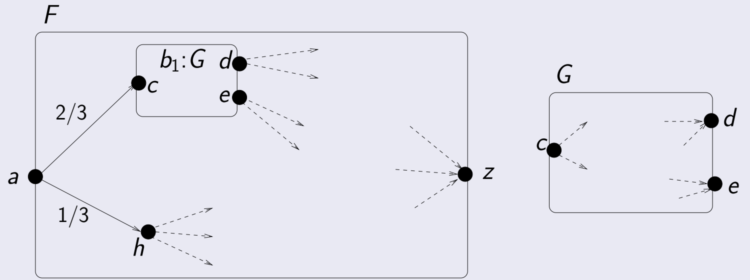

- $x_{z,z} = 1$
- $x_{a,z} = \frac{1}{3} x_{h,z} + \frac{2}{3} x_{(b_1,c),z}$
- $x_{(b_1,c),z} = x_{c,d} x_{(b_1,d),z} + x_{c,e} x_{(b_1,e),z}$

These "patterns" cover all cases. Yields a system of polynomial equations:

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

Basic fact about $\mathbf{x} = P(\mathbf{x})$

$P : \mathbb{R}^n \mapsto \mathbb{R}^n$ defines a monotone operator on $\mathbb{R}^n_{\geq 0}$.
It has a least fixed point, $x^* \in \mathbb{R}^n_{\geq 0}$.
(I.e., $x^* = P(x^*)$), and for any fixed point $y^* \in \mathbb{R}^n_{\geq 0}$, $x^* \leq y^*$.)

Theorem

- *x^* is the vector of termination probabilities for the RMC.*
- $x^* = \lim_{m \to \infty} \mathrm{P}^m(\mathbf{0})$

Question

Can we compute these probabilities efficiently?

Why not just do value iteration?
I.e., start with $x^0 := \mathbf{0}$, and let $x^{i+1} := P(x^i) = P^i(0)$, $i = 1, 2, 3, \ldots$

Value iteration can require exponentially many iterations

Question

How many iterations, m, is required for $x^m = P^m(0)$ to be within i bits of precision (i.e., to within additive error $1/2^i$) of the solution x^*?

Answer

In the worst case, at least exponentially many iterations in i, even for a fixed RMC:

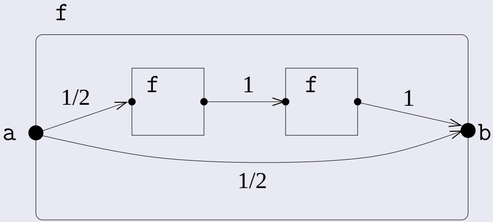

Equation: $x = (1/2)x^2 + 1/2$.

Fact: $x^* = 1$, but for all $m \leq 2^i$, $|1 - P^m(0)| \geq 1/2^i$.

More examples of pathologies: very small, and very large, probabilities

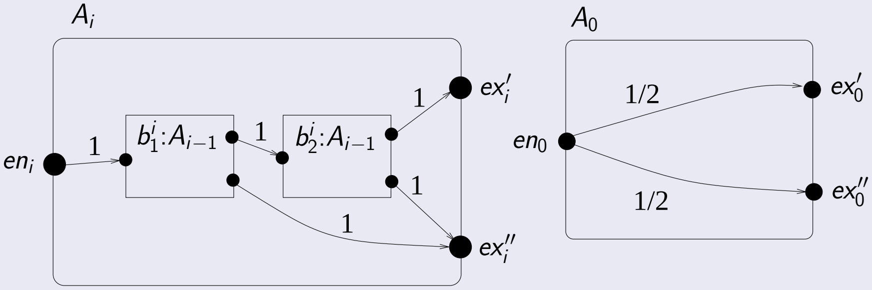

$i = 1, \ldots, n$

Fact: $x^*_{(en_n, ex_n')} = \frac{1}{2^{2^n}}$ and $x^*_{(en_n, ex_n'')} = 1 - \frac{1}{2^{2^n}}$.

some complexity upper bounds

existential theory of reals ($\exists\mathbb{R}$)

An $\exists\mathbb{R}$ sentence looks like: $\exists x_1, \ldots, x_k B(\bar{x})$, where $B(\bar{x})$ is a boolean combination of polynomial predicates: $F(\bar{x}) > 0$, $F(\bar{x}) = 0$, $F(\bar{x}) \geq 0$, etc.

Theorem

Given an RMC, and a probability p, we can decide whether $x^*_{u,ex} \geq p$ in PSPACE, and given i (in unary), we can approximate $x^*_{u,ex}$ to within i bits (i.e., within additive error $\geq 1/2^i$) in PSPACE.

Proof.

$\exists \mathbf{x}(\mathbf{x} = P(\mathbf{x}) \wedge \mathbf{a} \leq \mathbf{x} \leq \mathbf{b})$ can be expressed as a formula in $\exists\mathbb{R}$.
There are PSPACE decision procedures for $\exists\mathbb{R}$ ([Canny'89,Renegar'92]).
In order to approximate the probabilities to within i bits, we can do binary search, using i queries to $\exists\mathbb{R}$. $\square$

some "hard" problems

Sqrt-Sum

The square-root sum problem is the following decision problem:

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

It is solvable in PSPACE.

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

PosSLP

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

PosSLP captures everything one can do in polynomial time in the unit-cost arithmetic RAM model of computation.

[Allender et. al.'06] Gave a (Turing) reduction from Sqrt-Sum to PosSLP and showed both can be decided in the Counting Hierarchy: $P^{PP^{PP^{PP}}}$.

some bad news

Theorem ([E.-Yannakakis'05,'07])

Both Sqrt-Sum and PosSLP are P-time reducible to all of the following problems:

1. *Given a 1-exit RMC, and a rational p, decide whether* $x^*_{u,ex} \geq p$.
2. *Given a 2-exit RMC, decide whether* $x^*_{u,ex} = 1$.
3. *Given a 2-exit RMC, compute any non-trivial approximation of* $x^*_{u,ex}$. *More precisely:*
   - *For any fixed* $\epsilon > 0$, *given a 2-exit RMC in which either (a)* $x^*_{u,ex} = 1$, *or (b)* $x^*_{u,ex} \leq \epsilon$; *decide which of (a) or (b) is the case.*

- *Note: Sqrt-Sum and PosSLP are also reducible to approximating any (actual) Nash Equilibrium in a 3-player game. (See ([E.-Yannakakis,FOCS'07]).)*

toward some good news

Newton's method

Seeking a solution to $F(\mathbf{x}) = 0$, we start at initial guess vector $\mathbf{x}_0$, and compute the sequence, x_k , $k \to \infty$, where:

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

Here $F'(\mathbf{x})$, is the *Jacobian matrix*, of partial derivatives, given by

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

- Method not defined unless the matrices $F'(\mathbf{x}_k)$ are non-singular.
- Even when defined, it can diverge (even for univariate polynomials).
- But when it does converge, it is typically quite fast....

The good news for RMCs

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

Theorem (Decomposed Newton's method for RMCs [E.-Yannakakis'05])

Starting at $x_0 := \mathbf{0}$, and working "bottom-up" on the SSCs of the decomposition DAG of $x = P(x)$, Newton's method "monotonically converges" to the LFP, i.e., $\lim_{k \to \infty} \mathbf{x}_k \uparrow x^*$.

Implemented in PReMo (`http://groups.inf.ed.ac.uk/premo/`), tool developed by Dominik Wojtczak [Wojtczak-E.,'07].
Experiments on large benchmarks (from NLP) and very large random instances (up to 0.5M variables), yield good results.

Note

In [E.-Yannakakis'05] we actually showed decomposed Newton's method works more generally for all monotone systems of polynomial equations.

What is Newton's worst case behavior for RMCs?

- [Kiefer,Luttenberger,Esparza,'07] gave examples requiring exponentially many iterations (as a function of the encoding size of the equations) to converge to within additive error $< 1/2$.
- For strongly-connected equation systems [Esparza, Kiefer, Luttenberger, 2008] gave an exponential upper bound, as a function of the size of the system, and linear in the number of bits of precision required.
- We know no good general upper bounds on the number of iterations, *as a function of the system size*, for arbitrary RMCs.
- But we do know much more for special subclasses of RMCs....

A subclass of RMCs: Quasi-Birth-Death processes

Theorem [E.-Wojtczak-Yannakakis'08]

For discrete-time Quasi-Birth-Death Processes (= 1-box RMCs = probabilistic 1-counter automata (p1CA)), polynomially many iterations of Newton's method in the encoding size of the QBD, and in i, suffice to get termination probabilities (a.k.a. the G matrix) within additive error $1/2^i$.

Proof establishes interesting combinatorial properties for QBDs/p1CAs. Decomposition is a key to the analysis. Another key is results by [Esparza-Kiefer-Luttenberger'08] for their exponential upper bounds.

- Yields P-time in Blum-Shub-Smale model for approximating G.
- QBDs have been heavily studied in queueing theory and performance evaluation (basic model of a multi-phase queue) since the 1970s.
- Computing the G matrix (the termination probabilities) for QBDs is a key ingredient for many other quantitative analysis tasks, e.g., for computing steady-state probabilities, etc.

- Special "Matrix Analytic" numerical methods have been developed for many years for analyzing QBDs and related Structured Markov chains (e.g., M/G/1-Type). See, e.g., the books: [Neuts'81],[Latouche-Ramaswami'99],[Bini-Latouche-Meini'05]
- Among the key matrix analytic methods are logarithmic reduction and cyclic reduction. (Implemented in tools like SMCSolver [Bini-Meini-Steffe-Van Houdt'06].)
- These methods far outperform decomposed Newton's method on dense instances of QBDs, but decomposed Newton's method can outperform them on very sparse instances (see [E.-Wojtczak-Yannakakis'08-'10] for some comparisons).
- Tree-Like QBDs are a generalization of QBDs studied in the more recent structured-MC literature (see, e.g., [Bini-Latouche-Mini'05]). They are equivalent to RMCs (see [E.-Wojtczak-Yannakakis'08]).
- There are interesting open problems related to efficiently/practically combining matrix-analytic methods with decomposition methods for analysis of QBDs and TL-QBDs (= RMCs). See [EWY'08].

Another subclass of RMCs: multi-type Branching Processes, and Stochastic Context-Free Grammars

- 1-exit RMCs, where every component has one exit, are "equivalent" in precise senses to stochastic context-free grammars (SCFGs) and multi-type branching processes (MT-BPs).
- SCFGs are a fundamental model in statistical natural language processes, and are also used extensively in biological sequence analysis (RNA secondary structure analysis).
- MT-BPs are a classic and heavily studied class of stochastic processes ([Kolmogorov'1940s]), with many applications.
- Again, computing termination probabilities is a key ingredient for many important analysis problems for both SCFGs and MT-BPs. Termination probabilities known as the partition function for SCFGs. Also known as the extinction probabilities for MT-BPs.

Multi-type Branching Processes (Kolmogorov,1940s)

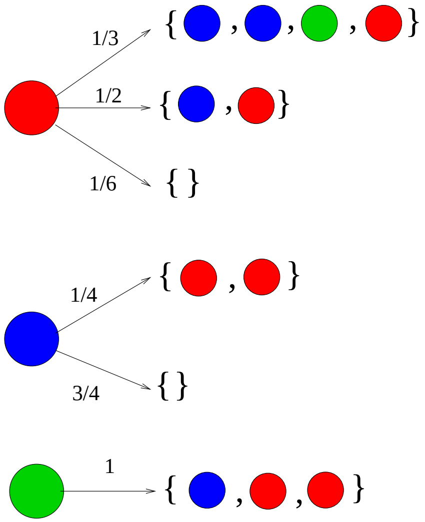

Multi-type Branching Processes (Kolmogorov,1940s)

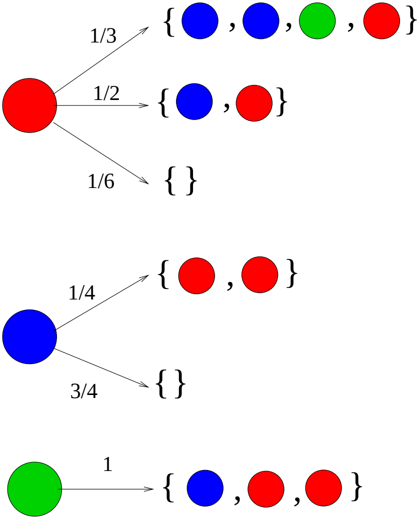

Question

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

Multi-type Branching Processes (Kolmogorov,1940s)

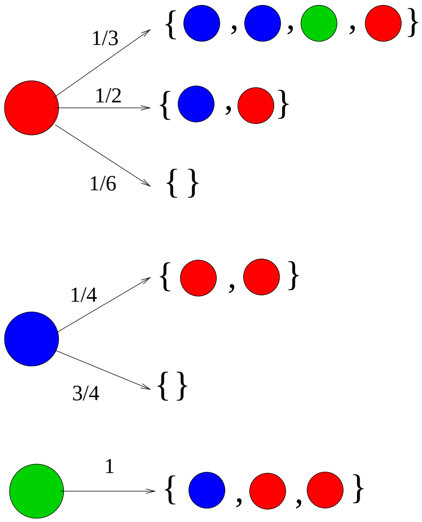

Question

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

$x_R \; =$

Multi-type Branching Processes (Kolmogorov,1940s)

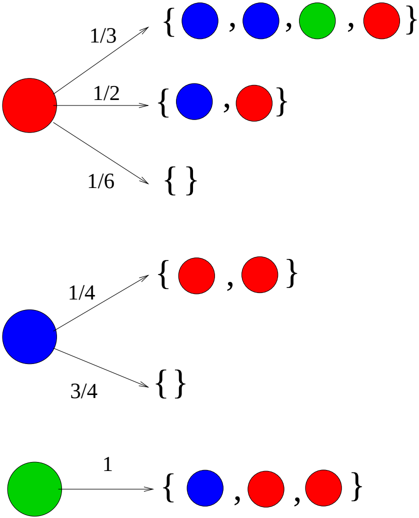

Question

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)

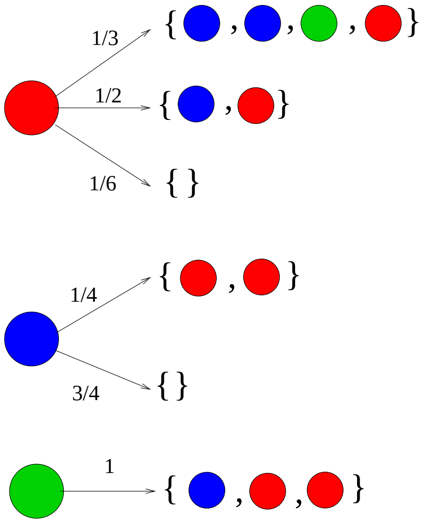

Question

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

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

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

$$x_G = x_B x_R^2$$

Multi-type Branching Processes (Kolmogorov,1940s)

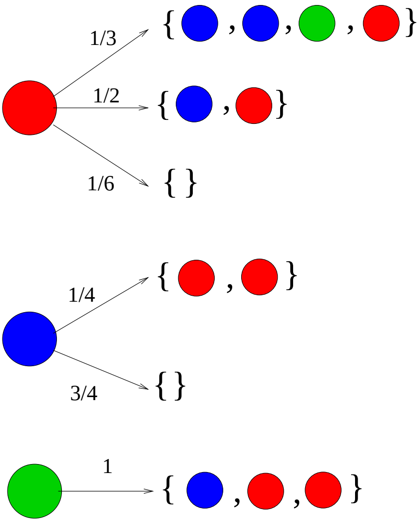

Question

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 fixed point equations, $\bar{\mathbf{x}} = F(\bar{\mathbf{x}})$.

Multi-type Branching Processes (Kolmogorov,1940s)

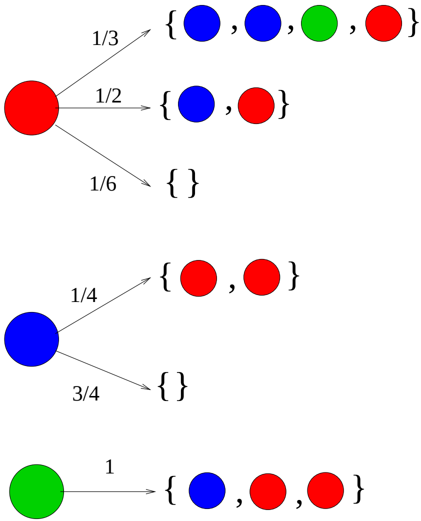

Question

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 fixed point equations, $\bar{\mathbf{x}} = F(\bar{\mathbf{x}})$.

Fact

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

Multi-type Branching Processes (Kolmogorov,1940s)

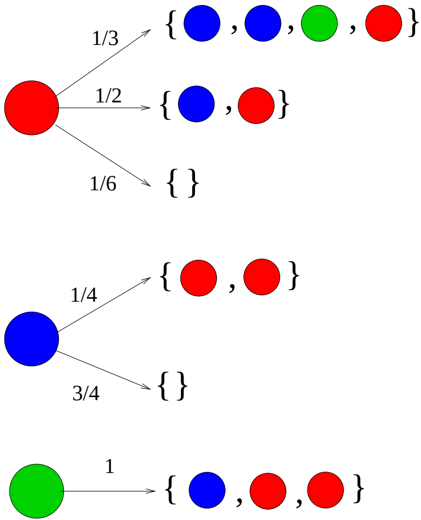

Question

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 fixed point equations, $\bar{\mathbf{x}} = F(\bar{\mathbf{x}})$.

Fact

The extinction probabilities are the least fixed point, $\mathbf{x}^* \in [0,1]^3$, of $\bar{\mathbf{x}} = F(\bar{\mathbf{x}})$.
$x_R^* = 0.276$; $x_B^* = 0.769$; $x_G^* = 0.059$.

Better algorithms for 1-exit RMCs, SCFGs and MT-BPs

Theorem ([E.-Yannakakis'05])

For 1-exit RMC, SCFGs, and MT-BPs, deciding whether the termination probability is $= 1$ is in polynomial time.

Proof: combines eigenvalue methods and graph-theoretic methods.
Key problem can be reduced to deciding whether certain moment matrices (Jacobian of $P(x)$ evaluated at the all 1 vector) have spectral radius > 1.
([Kolmogorov-Sevastyanov,'47,Harris'63])

Theorem [Fagin,Karlin,Kleinberg,Raghavan,Raj.,Rubinfeld,Sudan,'2000]

For Backbotton Processes ($=$ special subclass of 1-exit RMCs), x^* can be approximated in polynomial time.

Proof: semi-definite programming (constraints become convex).
Note: Not convex for more general 1-exit RMCs.

In [E.-Yannakakis'05 (JACM'09)] we also gave polynomial time algorithms for termination probabilities for various other special subclasses of RMCs:

- linearly-recursive RMCs.
- bounded RMCs, where the total number of entries and exits of all components is bounded by a fixed constant.

Model checking of RMCs

For an RMC, A, and an LTL formula or Büchi automaton, φ:
Let $P_A(\varphi)$ denote the probability that a run π of A satisfies property φ, i.e., that $\pi \in L(\varphi)$, where $L(\varphi)$ denotes ω-regular language defined by φ.

We are interested in the following model checking problems:

(1) *Qualitative model checking problems*:
Is $\mathrm{P}_A(\varphi) = 1$? Is $P_A(\varphi) = 0$?

(2) *Quantitative model checking problems*: for given $p \in [0, 1]$, is $P_A(\varphi) \geq p$?
Also approximate $P_A(\varphi)$ to within desired additive error $\epsilon > 0$.

[Esparza-Kucera-Mayr'04], [Brazdil-Kucera-Strazovsky'05] studied model checking of probabilistic pushdown systems (pPDSs) ($=$ RMCs).
[BKS'05] showed quantitative model checking for pPDS, A, and Büchi automaton, φ, is in EXPTIME in $|A|$, and 3-EXPTIME in $|\varphi|$.

Our results on ω-regular and LTL model checking of RMCs

Theorem [E.-Yannakakis'05,'05 (journal version, ToCL'11)]

Qualitative Model Checking of RMC A against *omega*-regular or LTL property φ:

	reachability	det. Büchi	Büchi/LTL
1-exit	P	P	P-time in $\|A\|$, EXPTIME in $\|\varphi\|$
Bounded	P	P	P-time in $\|A\|$, EXPTIME in $\|\varphi\|$
General	PSPACE	PSPACE	PSPACE in $\|A\|$, EXPTIME in $\|\varphi\|$

Quantitative Model Checking of RMC A against ω-regular or LTL property φ:

	reachability	det. Büchi	Büchi/LTL
1-exit	PSPACE	PSPACE	PSPACE in $\|A\|$, EXPSPACE in $\|\varphi\|$
Bounded	P	P-time in $\|A\|$	P-time in $\|A\|$
General	PSPACE	PSPACE	PSPACE in $\|A\|$, EXPSPACE in $\|\varphi\|$

Moreover, already for a fixed 1-exit RMC, qualitative model checking against an LTL formula, or a non-deterministic Büchi automaton, is EXPTIME-hard (thus EXPTIME-complete).

Brief hints of the many techniques involved

- A finite *conditioned summary chain*, $\mathcal{M}_A$ can be "built" using the termination probabilities x^*.
  This extends a summary graph construction for RSMs from [Alur-E.-Yannakakis'01] to the probabilistic setting.
- Many extensions of techniques from [Courcoubetis-Yannakakis'89] for model checking finite-state Markov chains.
- A certain kind of *unique fixed point theorem* for RMCs.
- Our upper bounds for Bounded RMCs involve monotone rational functions whose least fixed point characterizes desired probabilities.
- Note: Model checking LTL and Büchi automata properties of RMCs have same complexity in $|\varphi|$. Not surprising if you know [Courcoubetis-Yannakakis'89].
- Some of these quantitative algorithms are highly impractical because they use decision procedures for the existential theory of reals.

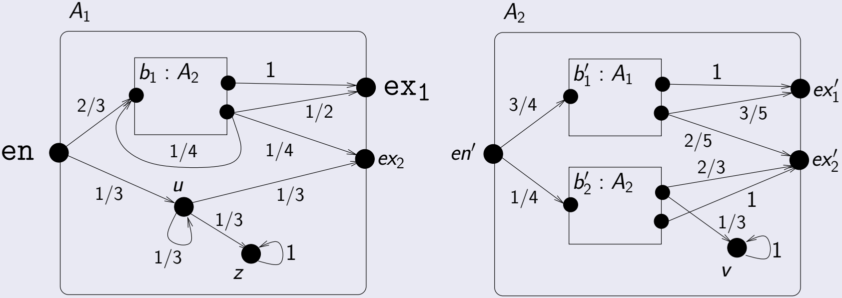

Let x^* be the LFP solution to $x = P(x)$ for this RMC.
For a vertex u in A_i, let $ne(u) = 1 - \sum_{ex \in Ex_i} x^*_{(u,ex)}$, be the probability of never exiting the component when starting at u.

The corresponding conditioned summary chain M_A

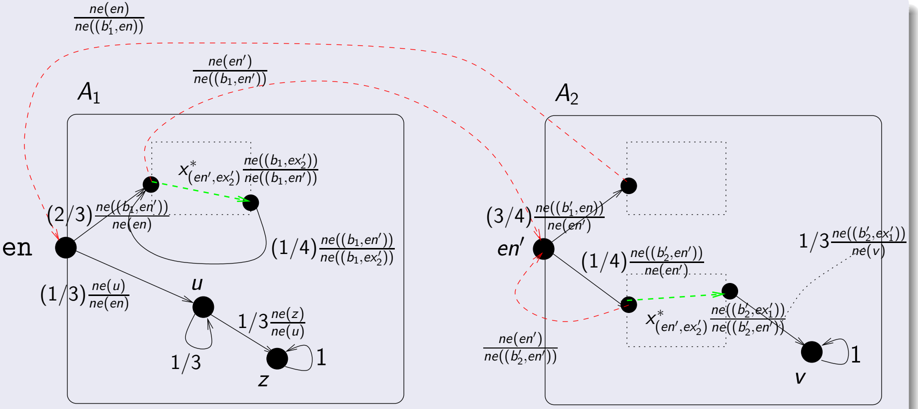

Each transition probability is now the conditional probability of making that transition, conditioned on never exiting that component. We get, e.g., $P_A(\Box\Diamond v) =$ probability of reaching bottom SCC containing v in M_A.

Some recent developments on model checking p1CAs = 1-box RMCs = discrete-time QBDs

- [Brazdil-Kiefer-Kucera'11] used the result of [E.-Wojtczak-Yannakakis'08] on P-time approximation of termination probabilities in the Blum-Shub-Smale model, to show that also approximating the model checking probability $P_A(\varphi)$ for a discrete-time QBD or p1CA, A, and for a deterministic Büchi or Rabin property, φ, can be done in P-time in the Blum-Shub-Smale model of computation.

  Basically involves showing transition probabilities of the conditioned summary chain can be approximated using the approximated termination probabilities. The key is bounding away the termination probabilities from 1, if they are below 1.
  [BKR'11] do this with a nice martingale construction.
- We do not yet know a P-time approximation algorithm in the standard Turing model of computation, even for approximating termination probabilities of QBDs and p1CAs.

On to Markov Decision Processes and Stochastic Games

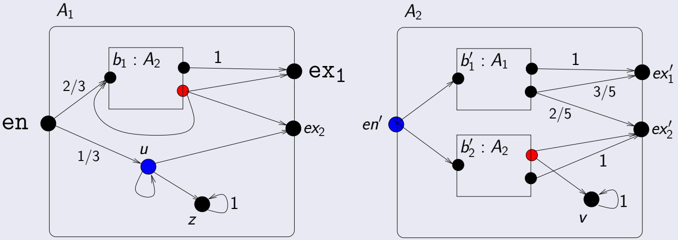

- Recursive MDPs (RMDPs): some nodes controlled, others probabilistic.
- Recursive Simple Stochastic Games (RSSGs): Two adversarial players, who control different nodes.
- Recursive Concurrent Stochastic Games (RCSGs): Two adversarial players who jointly control each state, and whose joint action determines a probability distribution on the next state.

infinite-state games defined by RMDPs/RSSGs/RCSGs

- Models define countable-state MDPs, SSGs, CSGs, in obvious way.
- Strategies for players define how they would choose to move at each node, possibly based on history.
  Fixing strategies for players induces a countable Markov chain.
- Many different objectives can be studied for RMDPs, RSSGs, RCSGs. Again, a key objective for many analyses is maximize/minimize termination probability.
- It follows from general determinacy results (e.g., [Martin'98]) that these termination games are determined, meaning they have a value.
- **Central algorithmic problem:** compute the value of the termination game, starting at given vertex, for a given RMDP, RSSG, RCSG.
- Again, we can consider both qualitative and quantitative problems.

Theorem [E.-Yannakakis'05]

For multi-exit RMDPs, even the qualitative termination value problem (is the value $= 1$), is undecidable.
This is so already for linearly-recursive RMDPs with a fixed constant number of exits.
Moreover, even any non-trivial approximation of RMDP optimal termination value is not computable.

- Proof is via a reduction from the emptiness problem for Probabilistic Finite Automata (PFA) [Rabin'63],[Paz'71],[Condon-Lipton'89],[Blondel-Canterini'03].
- We show PFAs can be "embedded" directly in RMDPs.

Theorem ([E.-Yannakakis'05,'07])

- *Quantitative termination value problems for 1-exit RMDPs, (1-exit RSSGs, and 1-exit RCSGs, can be decided in PSPACE using $\exists\mathbb{R}$.*

The proof shows that corresponding to each of these models is a system of monotone nonlinear-minimax equations whose least fixed point gives precisely the game values.

1. $u \in Type_{exit}$: $x_u = 1$.
2. $u \in Type_{random}$: $x_u = \sum_{v \in next(u)} p_{u,v} x_v$
3. $u \in Type_{call}$ (i.e., $u = (b, en')$): $x_u = x_{en'} \cdot x_{(b,ex'),ex}$
4. $u \in Type_{max}$: $x_u = \max_{v \in next(u)} x_v$
5. $u \in Type_{min}$: $x_u = \min_{v \in next(u)} x_v$

Again, yields monotone equations $\mathbf{x} = P(\mathbf{x})$ whose LFP gives the value of the game starting at each vertex.

Qualitative termination problems for 1-exit models

Theorem ([E.-Yannakakis'05,'06,'07])

1. *For 1-exit RMDPs, whether $x_u^* = 1$ is decidable P-time.*

   Proof: *boils down to a spectral optimization problem, solvable via LP.*

2. *For 1-exit RSSGs, we can decide whether $x_u^* = 1$ in NP $\cap$ coNP. (At least as hard as Condon's (1992) problem.)*

   Proof: *A key stackless-memoryless determinacy result for 1-exit RSSGs, (both players have pure, stackless, & memoryless optimal strategies) proved via a strategy improvement argument that relies on subtle analytic properties of certain power series associated with 1-exit RSSGs.*

3. *For 1-exit RCSGs even deciding whether $x_u^* = 1$ is SQRT-SUM-hard. (Player Max need not have any optimal strategies (Min does), but both players do have randomized, stackless, memoryless, (ϵ)-optimal strategies for all $\epsilon > 0$.)*

Model checking

- Given RMDP, RSSG, or RCSG, and given a ω-regular language $L \subseteq \Sigma^\omega$, (e.g., given by a Büchi automaton or LTL formula), we want to know what is the min/max probability value of the event that a run π of the model is in L.
- [E.-Yannakakis'05]]: already for 1-exit RMDPs, even qualitative model checking problems against LTL or ω-regular properties are undecidable.
- [Brazdil,Brozek,Forejt,Kucera'06], show that qualitative almost sure reachability problems for 1-exit RMDPs are decidable in P-time, by reducing this to the qualitative termination problem and the P-time algorithm of [E.-Yannakakis'06].
  They use this to show that qualitative-PCTL branching-time model checking of 1-exit RMDPs is decidable (and EXPTIME-complete).

Other analyses: expected total reward

- It is common to study MDPs with rewards on transitions, where the goal is to maximize/minimize expected total/discounted reward.

Theorem ([E.-Wojtczak-Yannakakis'08])

- *For 1-exit RMDPs with strictly positive rewards we can compute the optimal total expected reward, and an optimal strategy, in P-time.*
- *For 1-exit RSSGs with strictly positive rewards we can decide quantitative value problems for total expected reward in* **NP** $\cap$ **coNP**. *(At least as hard has Condon's problem.)*

- Can be used to analyse maximum/minimum expected running time of abstractions of probabilistic procedural programs with recursion.
- **Note:** We do not even know decidability if strictly positive rewards are replaced by non-negative rewards.

1-counter MDPs and 1-counter Stochastic Games

- [Brazdil-Brozek-E.-Kucera-Wojtczak'10]: For 1-counter MDPs (= 1-box RMDPs = controlled-QBDs), we can decide the qualitative termination problem (in the maximization case) in P-time.
- [Brazdil-Brozek-E.'10]: For 1-counter MDPs (in the minimization case) we can also decide the qualitative termination problem in P-time.

  For 1-counter SSGs we can decide qualitative termination in **NP** $\cap$ **coNP**.
- [Brazdil-Brozek-E.-Kucera'11]: For 1-counter MDPs and 1-counter SSGs can approximation of the termination value is computable (in EXPTIME and NEXPTIME, respectively).
- **Note:** We still don't know, how to decide for a 1-counter MDP whether the optimal termination probability value is $\geq p$.

Brief hints of proof techniques for 1-counter MDPs

- We heavily exploit the connection between 1-counter probabilistic models, and the classic theory of random walks on the integers.
- In a 1-counter MDP, as the counter value gets large, the optimal probability of termination (reaching counter value 0) approaches the optimal probability of forcing the $\liminf$ counter value to $-\infty$.
- We exploit this connection to relate the termination objective to $\liminf$ and mean payoff objectives which we can solve in P-time (via linear programming).
- For approximation, we upper bound how fast the optimal termination probability approaches the optimal $\liminf = -\infty$ probability, using a martingale construction (derived from an LP solution), and Azuma's inequality. (Related to a construction in [Brazdil-Kiefer-Kucera'10].)

Conclusion and overview

- A very rich landscape, with still many many remaining theoretical and practical open questions.
- Many important classes of stochastic processes are captured within the framework of adding recursion to Markov chains, or equivalently by probabilistic extensions to classic infinite-state automata-theoretic models like context-free grammars, pushdown systems, and one-counter automata.
- The algorithmic theory of these recursive stochastic models, and their extensions to MDPs and stochastic games, is an extremely rich subject.