Probabilistic Reachability in Markov Chains

Lecture #19 of Advanced Model Checking

Joost-Pieter Katoen

Lehrstuhl 2: Software Modeling & Verification

E-mail: `katoen@cs.rwth-aachen.de`

July 10, 2009

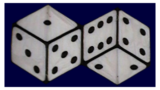 Probabilities help

- When analysing system performance and dependability
  - to quantify arrivals, waiting times, time between failure, QoS, ...
- When modelling uncertainty in the environment
  - to quantify environmental factors in decision support
  - to quantify unpredictable delays, express soft deadlines, ...
- When building protocols for networked embedded systems
  - randomized algorithms
- When analysing large populations
  - number of nodes in the internet, number of end-users, ...

Probabilistic verification so far

- Termination of probabilistic programs (Hart, Sharir & Pnueli, 1983)
  - does a probabilistic program terminate with probability one?
- Markov decision processes (Courcoubetis & Yannakakis, 1988)
  - does a certain (linear) temporal logic formula hold with probability p?
- Discrete-time Markov chains (Hansson & Jonsson, 1990)
  - can we reach a goal state via a given trajectory with probability p?
- Discrete-time Markov decision processes (Bianco & de Alfaro, 1995)
  - what is the maximal (or minimal) probability of doing this?
- Continuous-time Markov chains (Baier, Katoen & Hermanns, 1999)
  - can we do so within a given time interval I?

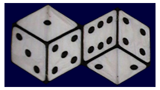

Characteristics

- What is inside?
  - temporal logics and model checking
  - numerical and optimisation techniques from performance and OR
- What can be checked?
  - time-bounded reachability, long-run averages, safety and liveness
- What is its usage?
  - powerful tools: PRISM (10,000 downloads), MRMC, Petri net tools, Probmela
  - applications: distributed systems, security, biology, quantum computing . . .

A synchronous leader election protocol

(Itai & Rodeh, 1990)

- A round-based protocol in a synchronous ring of $N > 2$ nodes
  - the nodes proceed in a lock-step fashion
  - each slot = 1 message is read + 1 state change + 1 message is sent

  $\Rightarrow$ this synchronous computation yields a Markov chain

- Each round starts by each node choosing a uniform id $\in \{ 1, \ldots, K \}$
- Nodes pass their selected id around the ring
- If there is a unique id, the node with the maximum unique id is leader
- If not, start another round and try again $\ldots$

Leader election

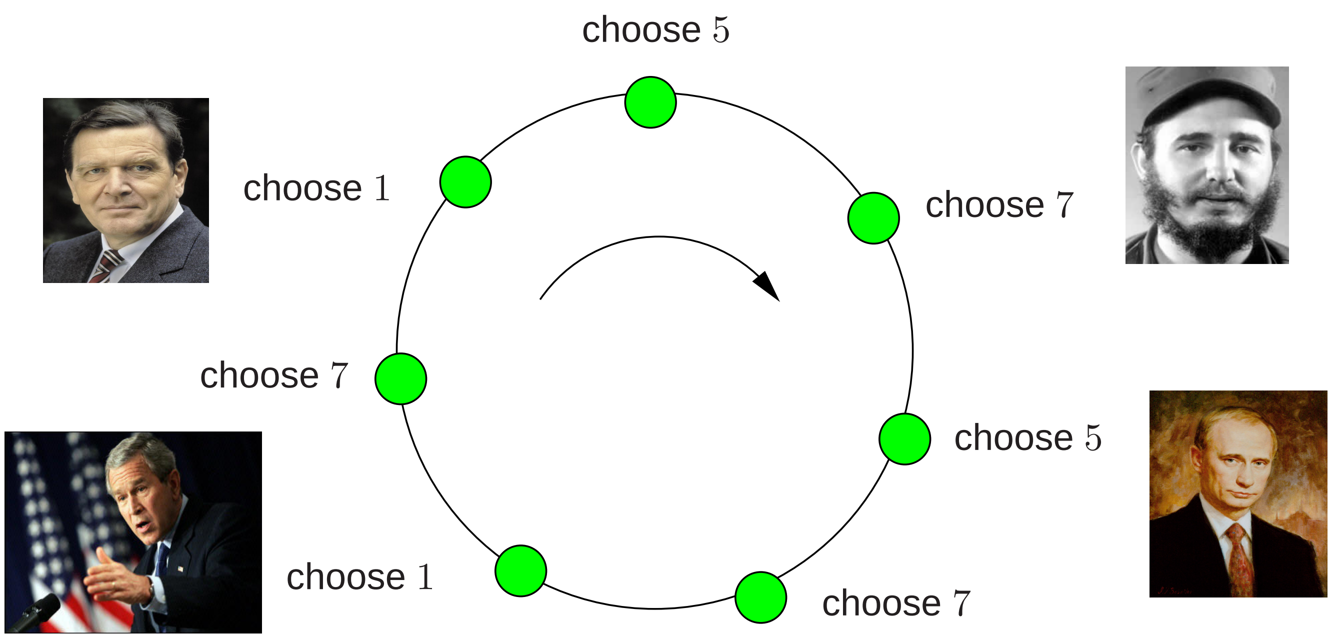

probabilistically choose an id from $[1 \ldots K]$

Leader election

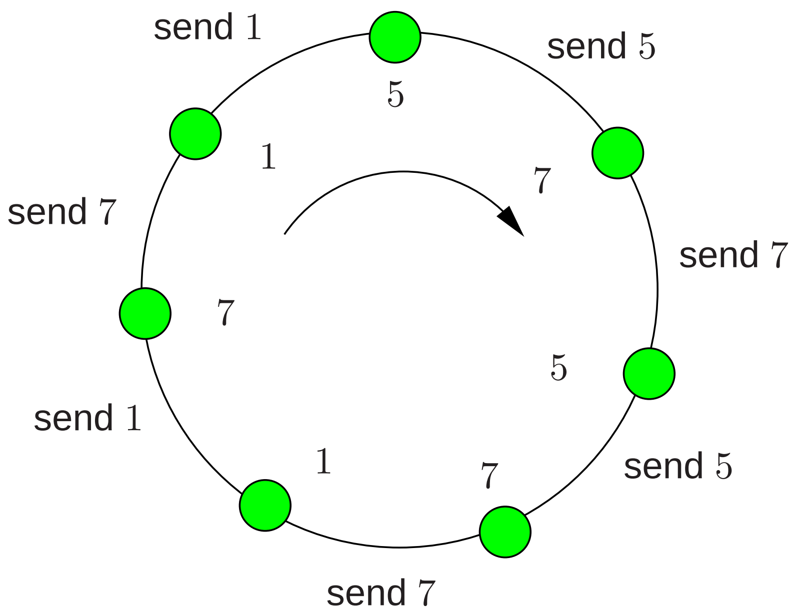

send your selected id to your neighbour

Leader election

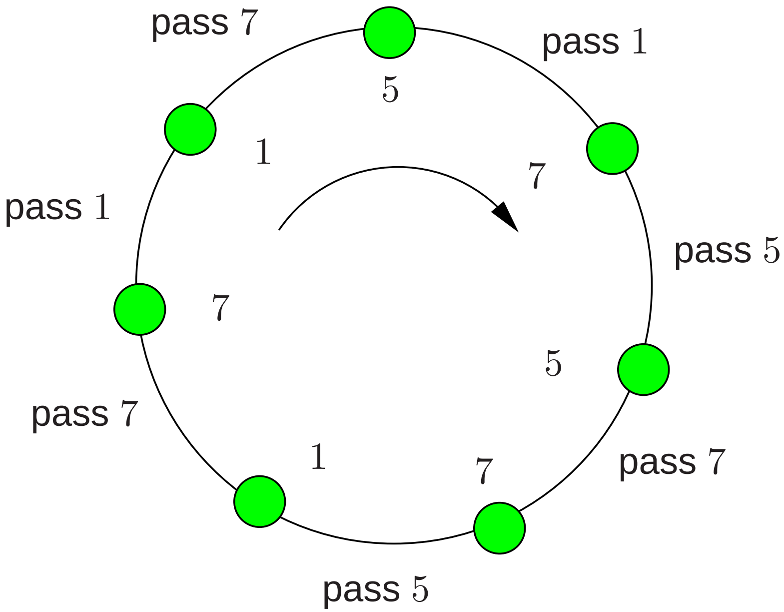

pass the received id, and check uniqueness own id

Leader election

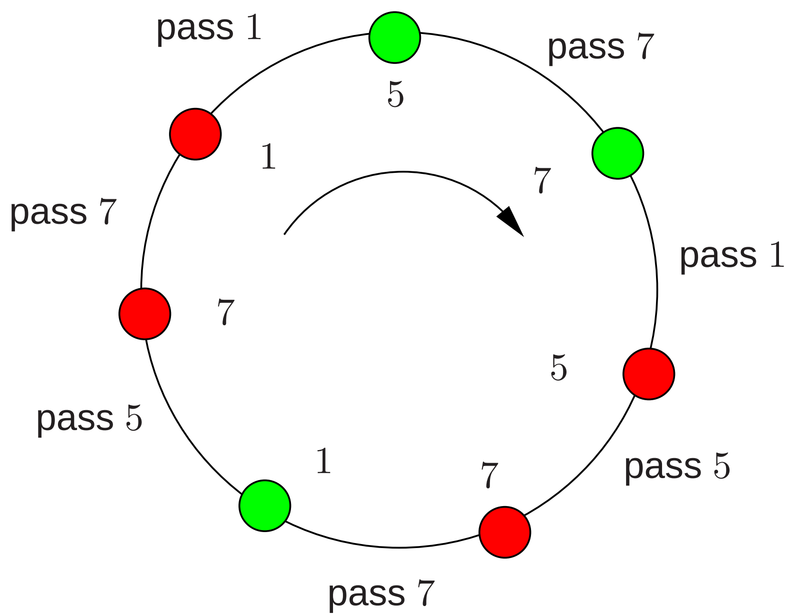

pass the received id, and check uniqueness own id

Leader election

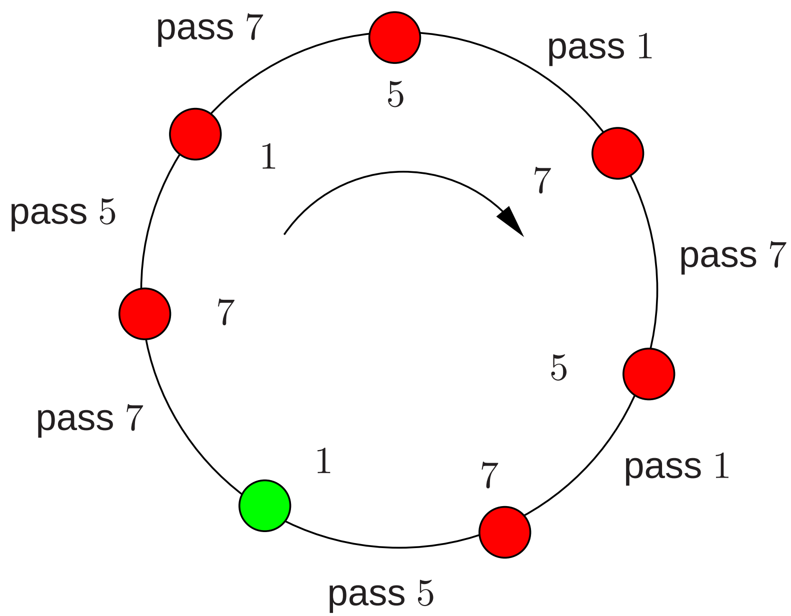

pass the received id, and check uniqueness own id

End of 1st round

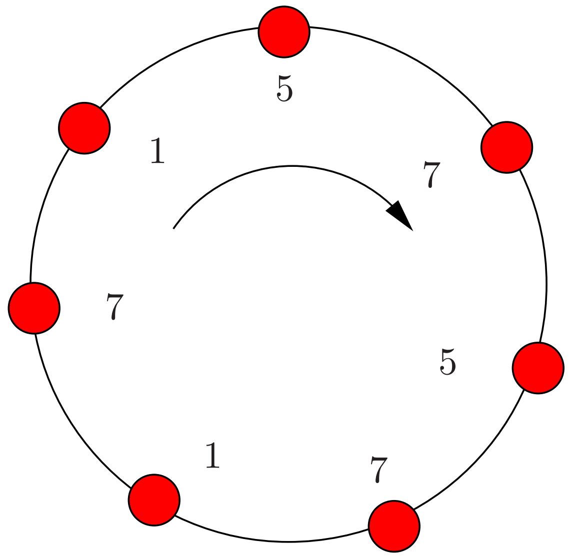

no unique leader has been elected

Start a new round

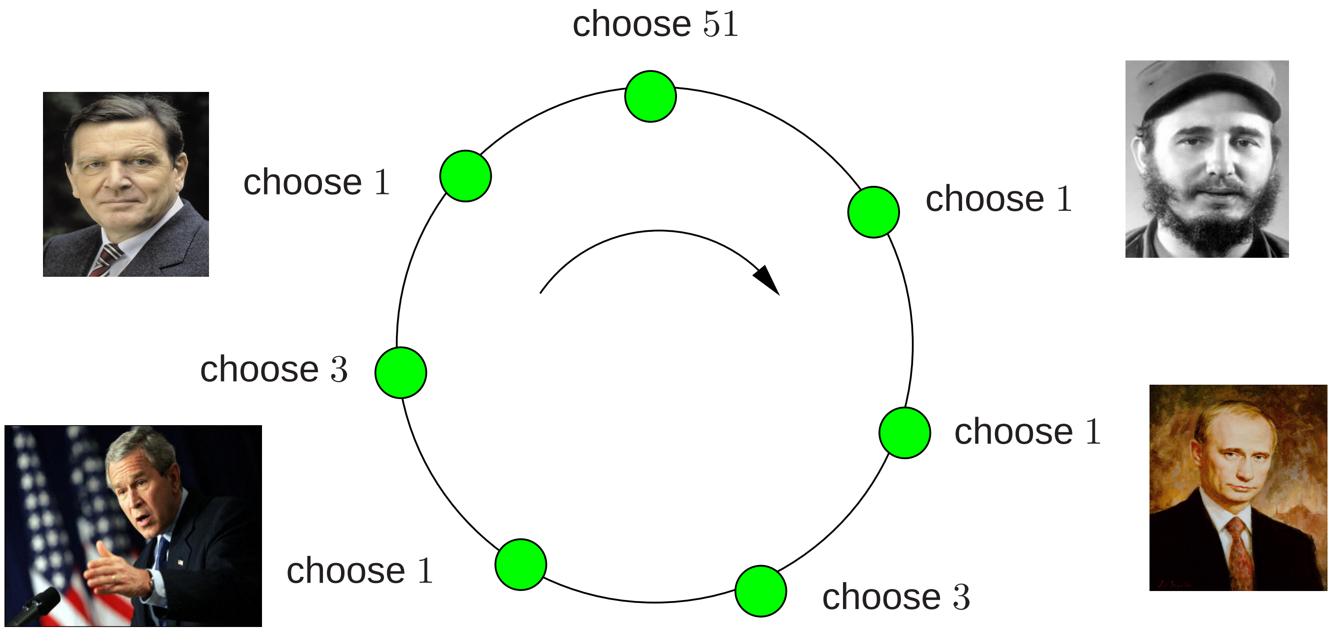

new round and new chances!

Properties of leader election

- Almost surely eventually a leader will be elected:

$$\mathbb{P}_{=1}\left(\Diamond \textit{leader elected}\right)$$

- With probability $\geqslant \frac{4}{5}$, eventually a leader is elected :

$$\mathbb{P}_{\geqslant 0.8}\left(\Diamond \textit{leader elected}\right)$$

- . . . . . . within k steps:

$$\mathbb{P}_{\geqslant 0.8}\left(\Diamond^{\leqslant k} \textit{leader elected}\right)$$

Probability to elect a leader within L rounds

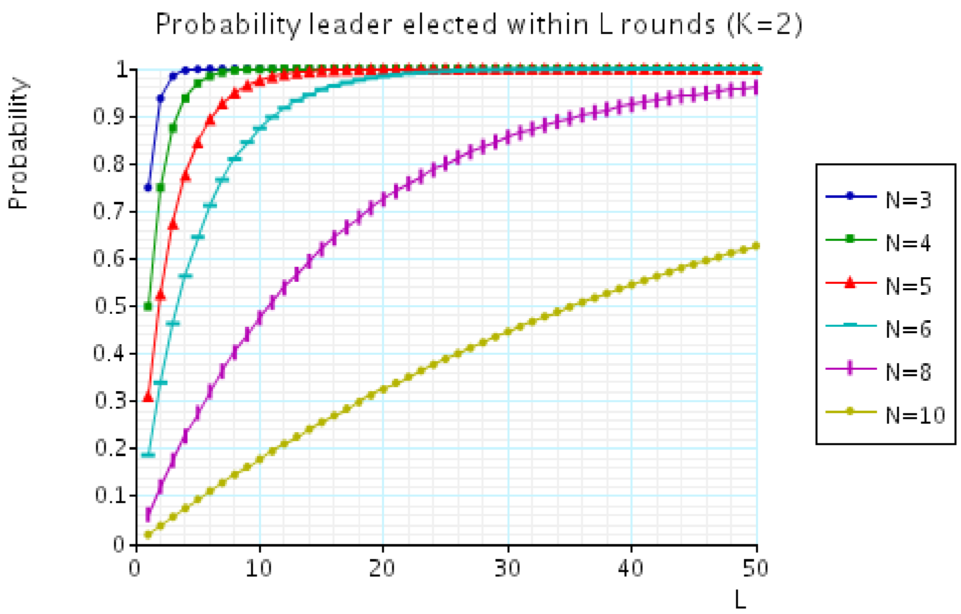

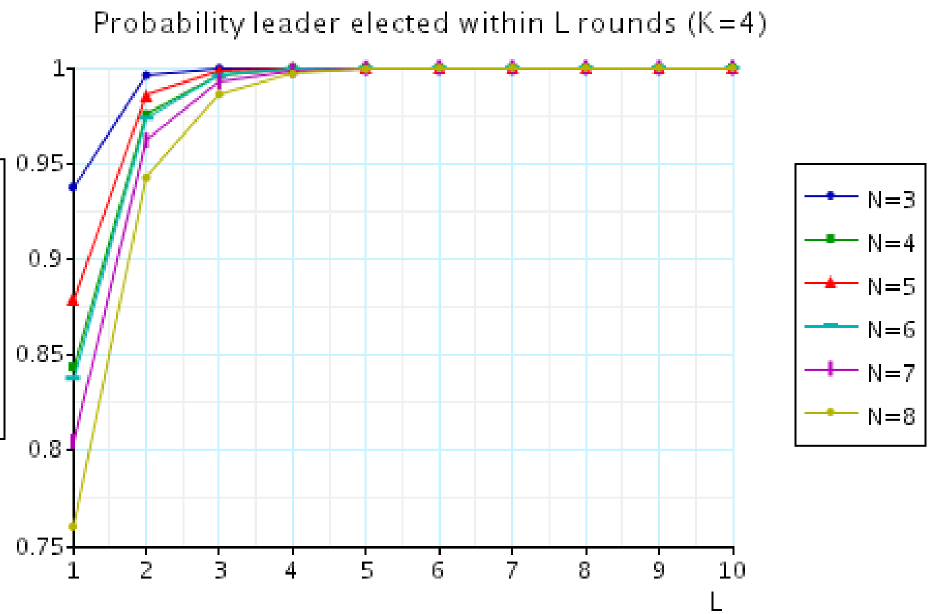

$$\mathbb{P}_{\leqslant q}\left(\Diamond^{\leqslant (N+1)\cdot L}\ \textit{leader elected}\right)$$ (Itai & Rodeh's algorithm)

Probability basics

- First, need an experiment
  - The sample space Ω is the set of possible outcomes
  - An event is a subset of Ω, can form events $A \cap B$, $A \cup B$, $\Omega \setminus A$

- Examples:
  - toss a coin: $\Omega = \{H,T\}$, events: "H", "T"
  - toss two coins: $\Omega = \{(H,H),(H,T),(T,H),(T,T)\}$, event: "at least one H"
  - toss a coin ∞-often: Ω is set of infinite sequences of H/T event: "H in the first 3 throws"
- Probability is:
  - Pr("H") = Pr("T") = $1/2$, Pr("at least one H") = $3/4$
  - Pr("H in the first 3 throws") = $1/2 + 1/4 + 1/8 = 7/8$

Probability basics

- Modelling a 6-sided die using a fair coin
  - algorithm due to Knuth/Yao:
  - start at 0, toss a coin
  - upper branch when H
  - lower branch when T
  - repeat until value chosen

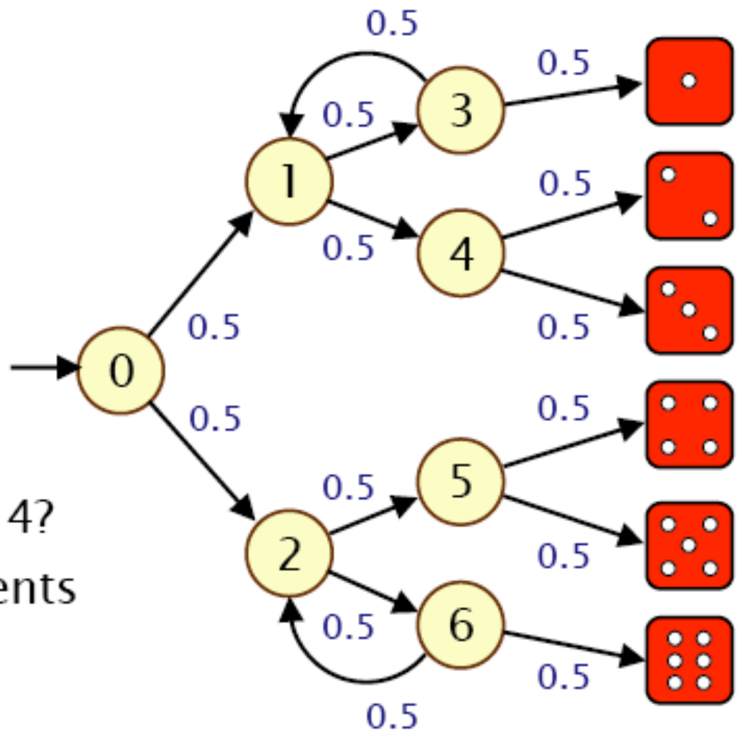

- Is this algorithm correct?
  - e.g. probability of obtaining a 4?
  - Obtain as disjoint union of events
  - THH, TTTHH, TTTTTHH, ...
  - Pr("eventually 4")

    $= (1/2)^3 + (1/2)^5 + (1/2)^7 + \ldots = 1/6$

Discrete-time Markov chains

- State–transition systems augmented with probabilities

- States
  - set of states representing possible configurations of the system being modelled
- Transitions
  - transitions between states model evolution of system's state; occur in discrete time-steps
- Probabilities
  - probabilities of making transitions between states are given by discrete probability distributions

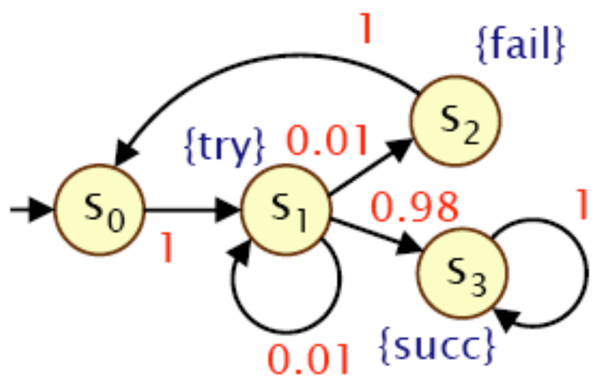

Discrete-time Markov chains

A DTMC $\mathcal{M}$ is a tuple $(S, \mathbf{P}, \iota_{\text{init}}, AP, L)$ with:

- S is a countable nonempty set of states
- $\mathbf{P} : S \times S \to [0, 1]$, transition probability function s.t. $\sum_{s'} \mathbf{P}(s, s') = 1$
  - $\mathbf{P}(s, s')$ is the probability to jump from s to s' in one step
- $\iota_{\text{init}} : S \to [0, 1]$, the initial distribution with $\sum\limits_{s \in S} \iota_{\text{init}}(s) = 1$
  - $\iota_{\text{init}}(s)$ is the probability that system starts in state s
  - state s for which $\iota_{\text{init}}(s) > 0$ is an initial state
- $L : S \to 2^{AP}$, the labelling function

$\Rightarrow$ a DTMC is a transition system with only probabilistic transitions

Example

- Recall Knuth/Yao's die algorithm from earlier:

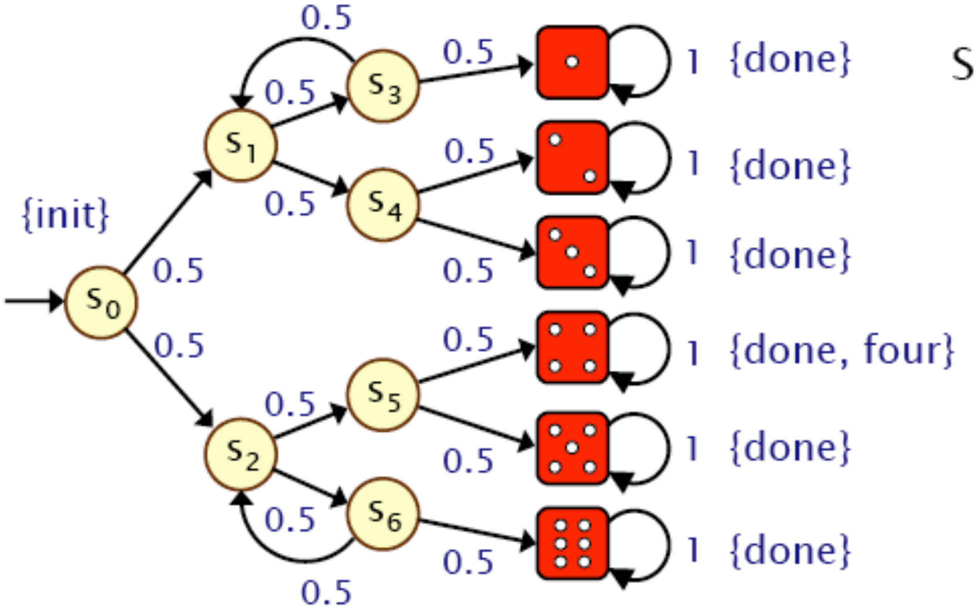

$S = \{ s_0, s_1, \ldots, s_6, 1, 2, \ldots, 6 \}$

$s_{init} = s_0$

$P(s_0,s_1)=0.5$
$P(s_0,s_2)=0.5$
etc.

$L(s_0) = \{init\}$
etc.

Paths in a DTMC

- State graph of DTMC $\mathcal{M}$ is a digraph $G = (V, E)$ with
  - vertices in V are states of $\mathcal{M}$, and $(s, s') \in E$ if and only if $\mathbf{P}(s, s') > 0$
- Paths in $\mathcal{M}$ are maximal (i.e., infinite) paths in its state graph
  - infinite sequence of states $s_0 s_1 s_2 \ldots \ldots$
  - unfolding of the DTMC
- Notations:
  - *Paths*$(\mathcal{M})$ and *Paths*$_{\mathit{fin}}(\mathcal{M})$ denote the set of (finite) paths in $\mathcal{M}$
- Direct successors and predecessors
  - *Post*$(s) = \{ s' \in S \mid \mathbf{P}(s, s') > 0 \}$ and *Pre*$(s) = \{ s' \in S \mid \mathbf{P}(s', s) > 0 \}$
  - *Post*$^*(s)$ and *Pre*$^*(s)$ are reflexive and transitive closure of *Post* and *Pre*

Paths and probabilities

- To reason (quantitatively) about this system
  - need to define a probability space over paths

- Intuitively:
  - sample space: Path(s) = set of all infinite paths from a state s
  - events: sets of infinite paths from s
  - basic events: cylinder sets (or "cones")
  - cylinder set Cyl(ω), for a finite path ω = set of infinite paths with the common finite prefix ω
  - for example: $Cyl(ss_1s_2)$

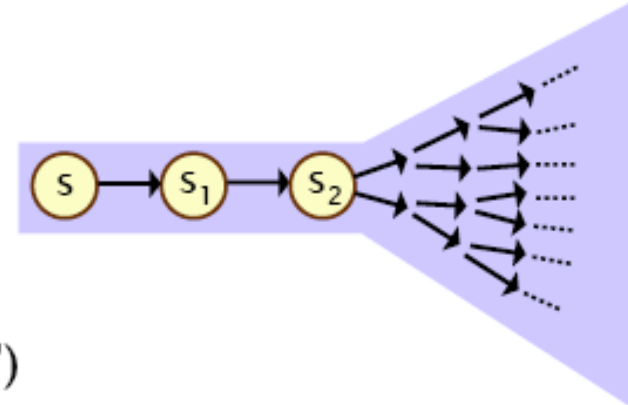

σ-algebra

- Let Ω be an arbitrary non-empty set
- $(\Omega, \mathcal{F})$ with $\mathcal{F} \subseteq 2^{\Omega}$ is a *σ-algebra* on Ω if:
  - $\varnothing \in \mathcal{F}$
  - $E \in \mathcal{F} \quad \Rightarrow \quad \Omega \setminus E \in \mathcal{F}$, and
  - $(\forall i \in \mathbb{N}.\ E_i \in \mathcal{F})$ implies $\bigcup_{i \in \mathbb{N}} E_i \in \mathcal{F}$
- Elements of $\mathcal{F}$ are called *measurable sets* or *events*
- For any family $\mathcal{F}$ of subsets of Ω:
  - there exists a unique smallest σ-algebra on Ω containing $\mathcal{F}$

Probability space

A *probability space* is a structure $(\Omega, \mathcal{F}, Pr)$ with:

- σ-algebra $(\Omega, \mathcal{F})$
- $Pr : \mathcal{F} \rightarrow [0, 1]$ is a *probability measure*, i.e.:
  1. $Pr(\Omega) = 1$, and
  2. $Pr\left(\bigcup_{i=1}^{\infty} E_i\right) = \sum_{i=1}^{\infty} Pr(E_i)$ for $E_i \in \mathcal{F}$ and $E_i \cap E_j = \varnothing$ for $i \neq j$

$Pr(E)$ is the probability of E, i.e., E is measurable

Probability space example

- Sample space Ω
  - $\Omega = \{1,2,3\}$
- Event set Σ
  - e.g. powerset of Ω
  - $\Sigma = \{ \emptyset, \{1\}, \{2\}, \{3\}, \{1,2\}, \{1,3\}, \{2,3\}, \{1,2,3\} \}$
  - (closed under complement/countable union, contains $\emptyset$)
- Probability measure Pr
  - e.g. $Pr(1) = Pr(2) = Pr(3) = 1/3$
  - $Pr(\{1,2\}) = 1/3+1/3 = 2/3$, etc.

Properties of probability measures

- An event E with $Pr(E) = 1$ is called *almost sure*
  - $Pr(D) \; = \; Pr(E \cap D) + \underbrace{Pr(D \setminus E)}_{=0} \; = \; Pr(E \cap D)$
- $E_1, \ldots, E_n$ are almost sure implies $\bigcap_{1 \leqslant i \leqslant n} E_i$ is almost sure
- For any Ω and $\mathcal{F} \subseteq 2^{\Omega}$ there exists a *smallest* σ-algebra containing $\mathcal{F}$
  - it is obtained by taking the intersection over all σ-algebras on Ω that contain $\mathcal{F}$

Probability space on DTMC paths

- Events are *infinite paths* in the DTMC $\mathcal{M}$, i.e., $\Omega = \mathit{Paths}(\mathcal{M})$
- σ-algebra on $\mathcal{M}$ is generated by *cylinder sets* of finite paths $\hat{\pi}$:

$$\mathit{Cyl}(\hat{\pi}) \;=\; \{\, \pi \in \mathit{Paths}(\mathcal{M}) \mid \hat{\pi} \text{ is a prefix of } \pi \,\}$$

  - cylinder sets serve as events of the smallest σ-algebra on $\mathit{Paths}(\mathcal{M})$

- *Pr* is the *probability measure* on the σ-algebra on $\mathit{Paths}(\mathcal{M})$:

$$\mathit{Pr}\big(\mathit{Cyl}(s_0 \ldots s_n)\big) \;=\; \iota_{\text{init}}(s_0) \cdot \mathbf{P}(s_0 \ldots s_n)$$

  - where $\mathbf{P}(s_0\, s_1 \ldots s_n) \;=\; \prod_{0 \leqslant i < n} \mathbf{P}(s_i, s_{i+1})$ if $n > 0$
  - and $\mathbf{P}(s_0) = 1$ for paths containing a single state

Probability space example

- Paths where sending fails the first time
  - $\omega = s_0 s_1 s_2$
  - $Cyl(\omega)$ = all paths starting $s_0 s_1 s_2 \ldots$
  - $P_{s0}(\omega) = P(s_0,s_1) \cdot P(s_1,s_2)$
    $= 1 \cdot 0.01 = 0.01$
  - $Pr_{s0}(Cyl(\omega)) = P_{s0}(\omega) = 0.01$

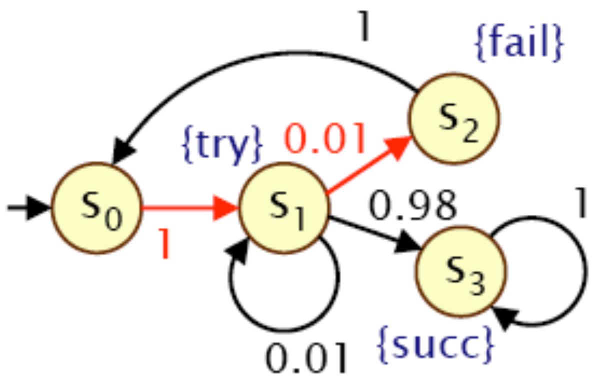

- Paths which are eventually successful and with no failures
  - $Cyl(s_0 s_1 s_3) \cup Cyl(s_0 s_1 s_1 s_3) \cup Cyl(s_0 s_1 s_1 s_1 s_3) \cup \ldots$
  - $Pr_{s0}(\, Cyl(s_0 s_1 s_3) \cup Cyl(s_0 s_1 s_1 s_3) \cup Cyl(s_0 s_1 s_1 s_1 s_3) \cup \ldots \,)$
    $= P_{s0}(s_0 s_1 s_3) + P_{s0}(s_0 s_1 s_1 s_3) + P_{s0}(s_0 s_1 s_1 s_1 s_3) + \ldots$
    $= 1 \cdot 0.98 + 1 \cdot 0.01 \cdot 0.98 + 1 \cdot 0.01 \cdot 0.01 \cdot 0.98 + \ldots$
    $= 0.9898989898\ldots$
    $= 98/99$

Reachability probabilities

- What is the probability to reach a set of states $B \subseteq S$ in DTMC $\mathcal{M}$?
  - B could be certain *bad* states which should be visited only seldomly

- Which event does $\Diamond B$ mean formally?
  - the union of all cylinders $\mathit{Cyl}(s_0 \ldots s_n)$ where
  - $s_0 \ldots s_n$ is an initial path fragment in $\mathcal{M}$ with $s_0, \ldots, s_{n-1} \notin B$ and $s_n \in B$

$$\begin{aligned} \mathit{Pr}(\Diamond B) &= \sum_{s_0 \ldots s_n \in \mathit{Paths}_{\mathit{fin}}(\mathcal{M}) \cap (S \setminus B)^* B} \mathit{Pr}\big(\mathit{Cyl}(s_0 \ldots s_n)\big) \\ &= \sum_{s_0 \ldots s_n \in \mathit{Paths}_{\mathit{fin}}(\mathcal{M}) \cap (S \setminus B)^* B} \iota_{\text{init}}(s_0) \cdot \mathbf{P}(s_0 \ldots s_n) \end{aligned}$$

Computing reachability probabilities

- Compute as (infinite) sum...
- $\Sigma_{s0,\ldots,sn \in \text{Reachfin}(s,\,T)} \Pr_{s0}(Cyl(s_0,\ldots,s_n))$

  $= \Sigma_{s0,\ldots,sn \in \text{Reachfin}(s,\,T)} P(s_0,\ldots,s_n)$
- Example:
  – ProbReach(s_0, {4})

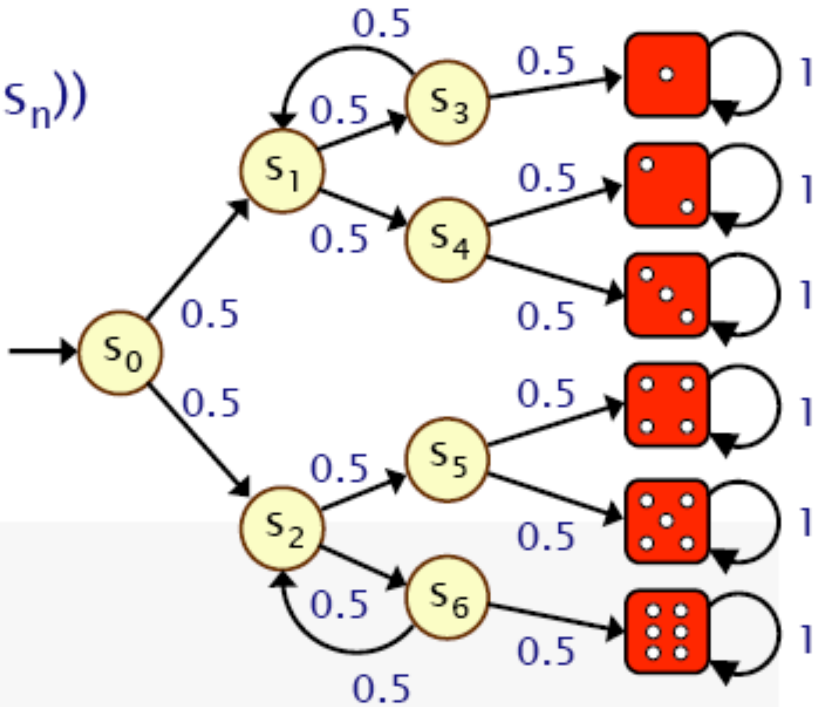

Computing reachability probabilities

- ProbReach(s_0, {s_6}) : compute as infinite sum?
  - doesn't scale...

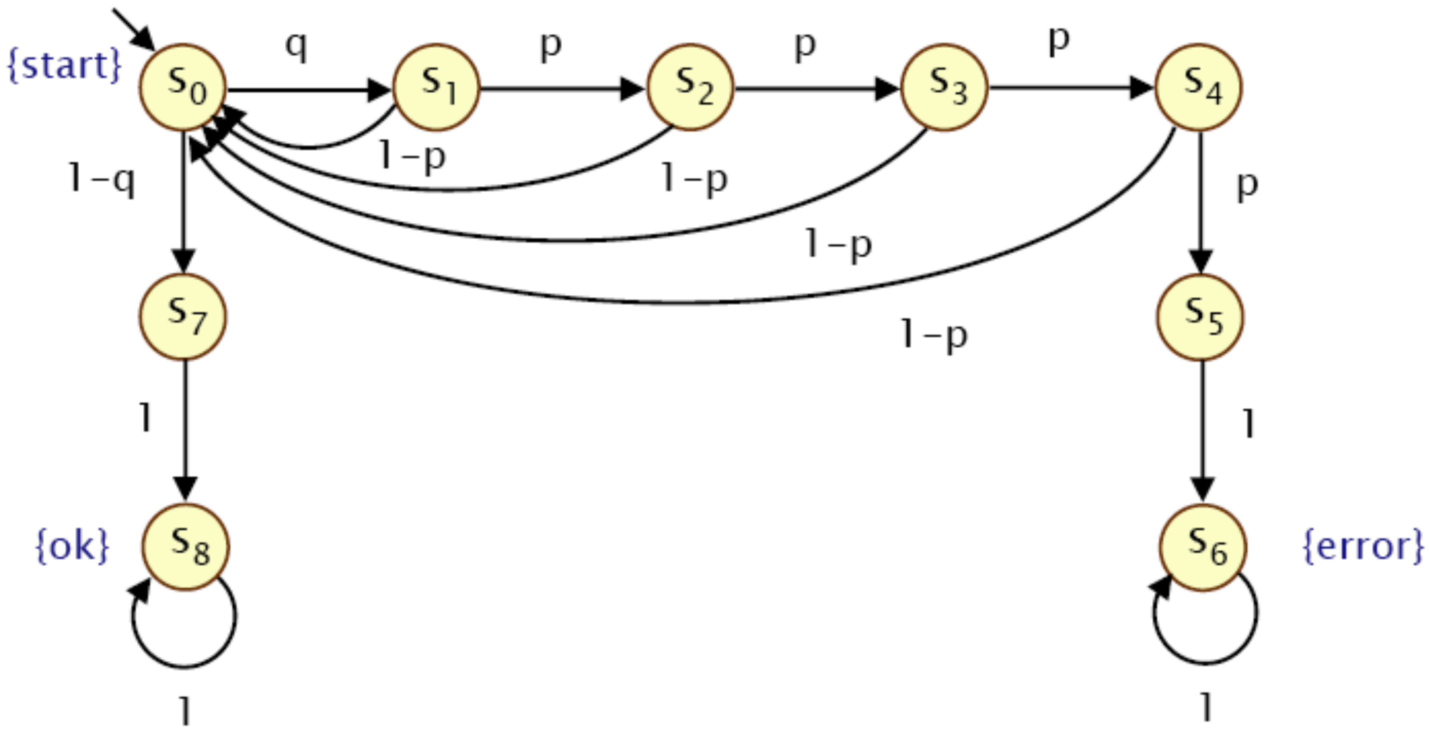

Reachability probabilities in finite DTMCs

- Let $Pr(s \models \Diamond B) = Pr_s(\Diamond B) = Pr_s\{\pi \in Paths(s) \mid \pi \models \Diamond B\}$
  - where Pr_s is the probability measure in $\mathcal{M}$ with only initial state s
- Let variable $x_s = Pr(s \models \Diamond B)$ for any state s
  - if B is not reachable from s then $x_s = 0$
  - if $s \in B$ then $x_s = 1$
- For any state $s \in Pre^*(B) \setminus B$:

$$x_s = \underbrace{\sum_{t \in S \setminus B} \mathbf{P}(s,t) \cdot x_t}_{\text{reach } B \text{ via } t} + \underbrace{\sum_{u \in B} \mathbf{P}(s,u)}_{\text{reach } B \text{ in one step}}$$

Example

- Compute ProbReach(s_0, {4})

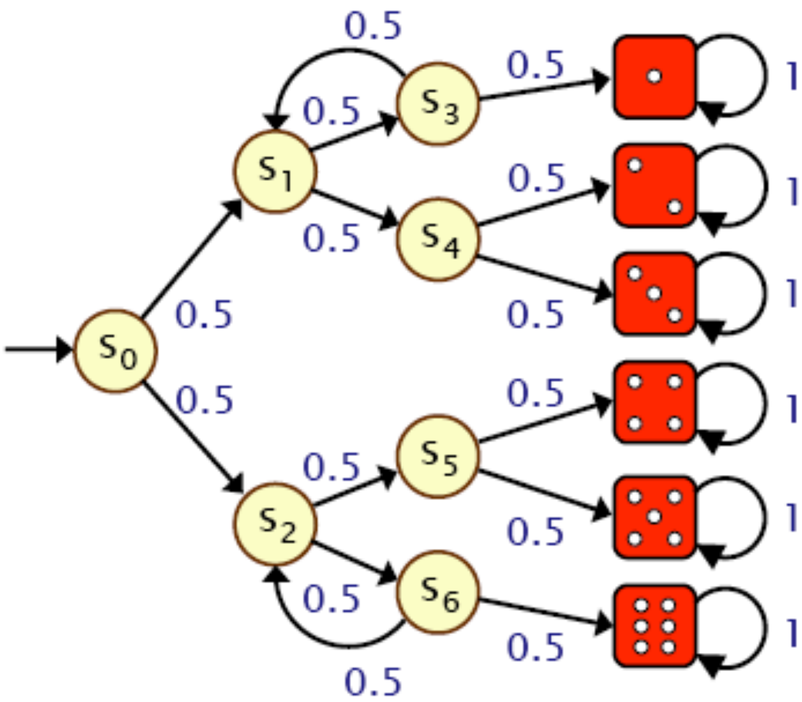

Unique solution

- Why the need to identify states that can reach T?

- Consider this simple DTMC:
  - compute probability of reaching $\{s_0\}$ from s_1

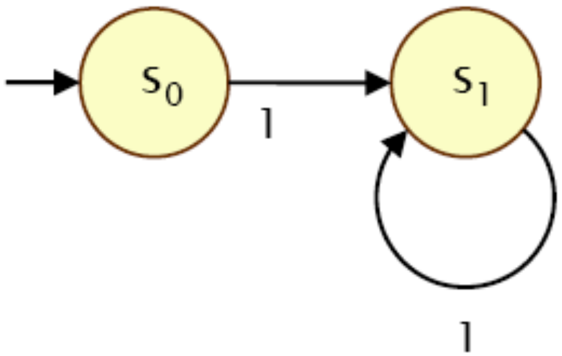

  - linear equation system: $x_{s0} = 1$, $x_{s1} = x_{s1}$
  - multiple solutions: $(x_{s0}, x_{s1}) = (1,p)$ for any p

Linear equation system

- These equations can be rewritten into the following form:

$$\mathbf{x} = \mathbf{A}\mathbf{x} + \mathbf{b}$$

  - where vector $\mathbf{x} = (x_s)_{s \in \tilde{S}}$ with $\tilde{S} = \mathit{Pre}^*(B) \setminus B$
  - $\mathbf{A} = \left(\mathbf{P}(s,t) \right)_{s,t \in \tilde{S}}$, the transition probabilities in $\tilde{S}$
  - $\mathbf{b} = \left(b_s \right)_{s \in \tilde{S}}$ contains the probabilities to reach B within one step

- *Linear equation system*: $(\mathbf{I} - \mathbf{A})\mathbf{x} = \mathbf{b}$
  - note: more than one solution may exist if $\mathbf{I} - \mathbf{A}$ has no inverse (i.e., is singular)

$\Rightarrow$ characterize the desired probability as least fixed point

Example

Let $B = \{\textit{delivered}\}$

$\tilde{S} = \{\textit{init}, \textit{try}, \textit{lost}\}$ and the equations:

$$\begin{aligned} x_{\textit{init}} &= x_{\textit{try}} \\ x_{\textit{try}} &= \frac{1}{10} \cdot x_{\textit{lost}} + \frac{9}{10} \\ x_{\textit{lost}} &= x_{\textit{try}} \end{aligned}$$

which can be rewritten as:

$$\begin{pmatrix} 1 & -1 & 0 \\ 0 & 1 & -\frac{1}{10} \\ 0 & -1 & 1 \end{pmatrix} \cdot \mathbf{x} = \begin{pmatrix} 0 \\ \frac{9}{10} \\ 0 \end{pmatrix}$$

and yields the (unique) solution: $x_{\textit{try}} = x_{\textit{init}} = x_{\textit{lost}} = 1$.

Constrained reachability

- Let $\mathcal{M} = (S, \mathbf{P}, \iota_{\text{init}}, AP, L)$ be a (possibly infinite) DTMC and $B, C \subseteq S$

- $C \, \mathsf{U}^{\leqslant n} B$ is the union of the basic cylinders of path fragments:
  - $s_0 \, s_1 \ldots s_k$ with $k \leqslant n$ and $s_i \in C$ for all $0 \leqslant i < k$ and $s_k \in B$

- Let $S_{=0}$, $S_{=1}$, $S_?$ be a partition of S such that:
  - $B \subseteq S_{=1} \subseteq \{ s \in S \mid Pr(s \models C \, \mathsf{U} \, B) = 1 \}$
  - $S \setminus (C \cup B) \subseteq S_{=0} \subseteq \{ s \in S \mid Pr(s \models C \, \mathsf{U} \, B) = 0 \}$
  - so: all states in $S_?$ belong to $C \setminus B$

- Let $\mathbf{A} = \big(\mathbf{P}(s,t) \big)_{s,t \in S_?}$ and $\big(b_s \big)_{s \in S_?}$ where $b_s = \mathbf{P}(s, S_{=1})$

Least fixed point characterization

The vector $\mathbf{x} = \big(Pr(s \models C \,\mathsf{U}\, B) \big)_{s \in S_?}$ is the *least fixed point* of the operator

$$\Upsilon : [0,1]^{S_?} \rightarrow [0,1]^{S_?} \quad \text{given by} \quad \Upsilon(\mathbf{y}) = \mathbf{A} \cdot \mathbf{y} + \mathbf{b}$$

Furthermore, for $\mathbf{x}^{(0)} = \mathbf{0}$ and $\mathbf{x}^{(n+1)} = \Upsilon(\mathbf{x}^{(n)})$ for $n \geqslant 0$:

- $\mathbf{x}^{(n)} = (x_s^{(n)})_{s \in S_?}$ where for any s: $x_s^{(n)} = Pr(s \models C \,\mathsf{U}^{\leqslant n} S_{=1})$
- $\mathbf{x}^{(0)} \leqslant \mathbf{x}^{(1)} \leqslant \mathbf{x}^{(2)} \leqslant \ldots \leqslant \mathbf{x}$, and
- $\mathbf{x} = \lim\limits_{n \rightarrow \infty} \mathbf{x}^{(n)}$

partial ordering is: $\mathbf{y} \leqslant \mathbf{y}'$ iff $y_s \leqslant y'_s$ for all $s \in S_?$

Proof

Constrained reachability probabilities

- So: $\mathbf{x}$ is the *least* solution of $\mathbf{Ax} + \mathbf{b} = \mathbf{x}$ in $[0,1]^{S_?}$
- And: can be approximated by:

$$\mathbf{x}^{(0)} = \mathbf{0} \quad \text{and} \quad \mathbf{x}^{(n+1)} = \mathbf{A}\mathbf{x}^{(n)} + \mathbf{b} \;\text{ for }\; n \geqslant 0$$

- *Power method*: compute vectors $\mathbf{x}^{(0)}, \mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \ldots$ and abort if:

$$\max_{s \in S_?} \left| x_s^{(n+1)} - x_s^{(n)} \right| < \varepsilon \qquad \text{for some small tolerance } \varepsilon$$

  - convergence guaranteed
  - alternative techniques: e.g., Jacobi or Gauss-Seidel, successive overrelaxation

Unique solution

Let $\mathcal{M}$ be a finite DTMC with state space S partitioned into:

- $S_{=0} \;=\; \textit{Sat}(\neg\exists(C \,\mathsf{U}\, B))$
- $S_{=1}$ a subset of $\{s \in S \mid \textit{Pr}(s \models C \,\mathsf{U}\, B) = 1\}$ that contains B
- $S_{?} \;=\; S \setminus (S_{=0} \cup S_{=1})$

For B, $C \subseteq S$, the vector

$$\big(\,\textit{Pr}(s \models C \,\mathsf{U}\, B)\,\big)_{s \in S_?}$$

is the *unique* solution of the linear equation system:

$$\mathbf{x} \;=\; \mathbf{A}\mathbf{x} + \mathbf{b} \quad \text{where} \quad \mathbf{A} \;=\; \big(\,\mathbf{P}(s,t)\,\big)_{s,t \in S_?} \;\text{ and }\; \mathbf{b} \;=\; \big(\,\mathbf{P}(s, S_{=1})\,\big)_{s \in S_?}$$

Computing constrained reachability probabilities

- The probabilities of the events $C \, \mathsf{U}^{\leqslant n} B$ can be obtained iteratively:

$$\mathbf{x}^{(0)} = \mathbf{0} \quad \text{and} \quad \mathbf{x}^{(i+1)} = \mathbf{A}\mathbf{x}^{(i)} + \mathbf{b} \text{ for } 0 \leqslant i < n$$

- where $\mathbf{A} = \big(\mathbf{P}(s,t)\big)_{s,t \in C \setminus B}$ and $\mathbf{b} = \big(\mathbf{P}(s,B)\big)_{s \in C \setminus B}$

- Then: $\mathbf{x}^{(n)}(s) = Pr(s \models C \, \mathsf{U}^{\leqslant n} B)$ for $s \in C \setminus B$

Bounded reachability probabilities

- $\text{ProbReach}(s_0, \{1,2,3,4,5,6\}) = 1$
- $\text{ProbReach}^{\leq k}(s_0, \{1,2,3,4,5,6\}) = \ldots$

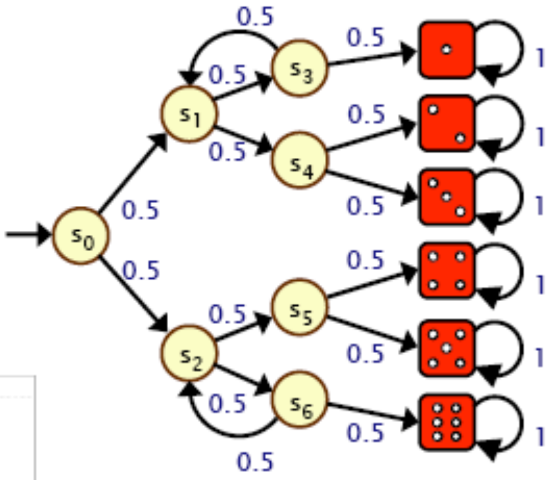

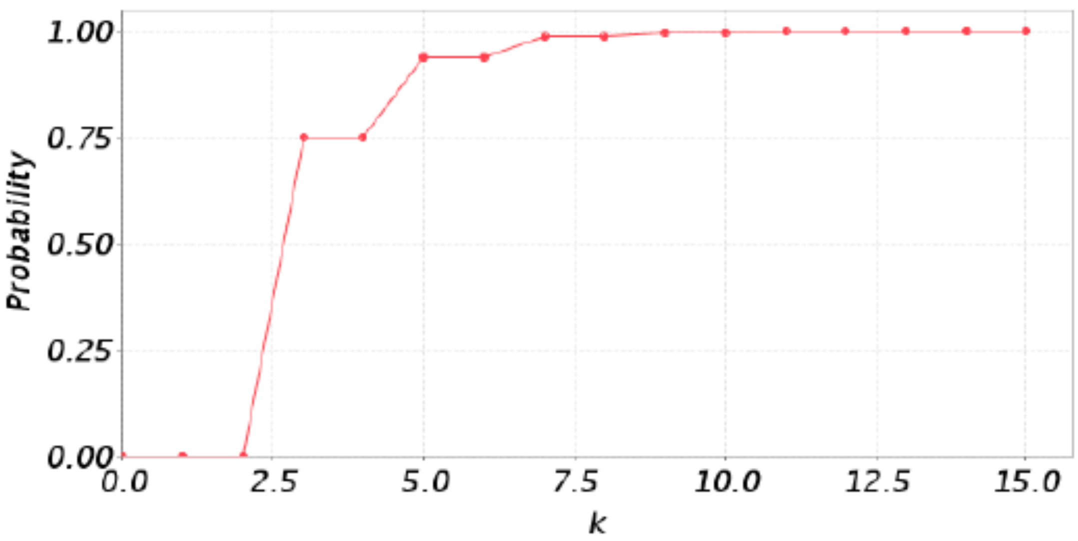

Transient probabilities

- Given that $\mathbf{A}^n(s,t) \; = \; Pr(s \models S_? \, \mathsf{U}^{=n} \, t)$
  - if $B = \varnothing$, $C = S$, we have $S_{=1} = S_{=0} = \varnothing$ and $S_? = S$ and $\mathbf{A} \; = \; \mathbf{P}$
  - $\mathbf{P}^n(s,t)$ is the probability to be in state t after n steps once started in s

- Transient probability: $\Theta_n^{\mathcal{M}}(t) \; = \; \sum_{s \in S} \iota_{\text{init}}(s) \cdot \mathbf{P}^n(s,t)$

- $\Theta_n^{\mathcal{M}} \; = \; \underbrace{\mathbf{P} \cdot \mathbf{P} \cdot \ldots \cdot \mathbf{P}}_{n \text{ times}} \cdot \iota_{\text{init}} \; = \; \mathbf{P}^n \cdot \iota_{\text{init}}$
  - where the initial distribution ι_{init} is viewed as column-vector

- Compute $\Theta_n^{\mathcal{M}}$ by successive vector-matrix multiplication:

$$\Theta_0^{\mathcal{M}} \; = \; \iota_{\text{init}}, \qquad \Theta_n^{\mathcal{M}} \; = \; \mathbf{P} \cdot \Theta_{n-1}^{\mathcal{M}} \quad \text{for } n \geqslant 1$$

Reachability = transient probabilities

- Suppose we want to compute probabilities for $\Diamond^{\leqslant n} B$ in $\mathcal{M}$
  - observe: once B is reached, remaining behaviour is not important
- Adapt $\mathcal{M}$ by making all states in B absorbing
  - $\mathbf{P}_B(s,t) = \mathbf{P}(s,t)$ if $s \notin B$ and $\mathbf{P}_B(s,s) = 1$ for $s \in B$
  - all outgoing transitions of $s \in B$ are replaced by a single self-loop at s
- Then:

$$\underbrace{Pr^{\mathcal{M}}(\Diamond^{\leqslant n} B)}_{\text{reachability in } \mathcal{M}} = \underbrace{\sum_{s' \in B} \Theta_n^{\mathcal{M}_B}(s')}_{\text{transient probability in } \mathcal{M}_B}$$

Constrained reachability = transient probabilities

- Suppose we want to compute probabilities for $C \, \mathsf{U}^{\leqslant n} B$ in $\mathcal{M}$
  - observe: once B is reached, remaining behaviour is not important
  - observe: once $s \in S \setminus (C \cup B)$ is reached, remaining behaviour not important
- Adapt $\mathcal{M}$ by making all states in B and $S \setminus (C \cup B)$ absorbing
  - $\mathbf{P}_B(s,t) = \mathbf{P}(s,t)$ if $s \notin B$ and $\mathbf{P}_B(s,s) = 1$ for $s \in B$ or $s \in C \cup B$
- Then:

$$\underbrace{Pr^{\mathcal{M}}(C \, \mathsf{U}^{\leqslant n} B)}_{\text{reachability in } \mathcal{M}} = \underbrace{\sum_{s' \in B} \Theta_n^{\mathcal{M}_{C,B}}(s')}_{\text{transient probability in } \mathcal{M}_{C,B}}$$