Probabilistic Systems

Part 1: discrete time Markov chains

Dr. Gethin Norman

School of Computing Science
University of Glasgow

gethin.norman@glasgow.ac.uk

Discrete-time Markov chains

Discrete-time Markov chains (DTMCs)

Paths and probabilities for DTMCs

Probabilistic reachability for DTMCs

Rewards and expected reachability for DTMCs

Discrete-time Markov chains

Kripke structures augmented with probabilities

States:

- represent possible configurations of the system being modelled
- labelled by atomic propositions (properties that hold in the states)

Transitions:

- model evolution of a system's state
- occur in discrete time-steps

Probabilities:

- likelihood of making transitions between states are given by discrete probability distributions

Kripke structures

A Kripke structure is a tuple (S, s_0, T, L) where

- S is a finite set of states
- s_0 is the initial state
- $T \subseteq S \times S$ is the transition relation

  if $(s, s') \in T$, then there is a transition from s to s'
- $L : S \rightarrow 2^{AP}$ is the labelling function where AP is a set of atomic propositions

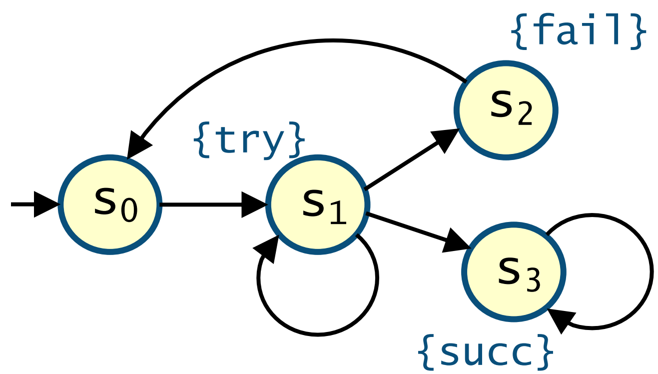

Discrete-time Markov chains

A discrete-time Markov chain is a tuple $(S, s_0, \mathbf{P}, L)$ where:

- S is a finite set of states
- s_0 is the initial state
- $\mathbf{P} : S \times S \rightarrow [0,1]$ is the transition probability matrix

  where $\mathbf{P}(s,s')$ is the probability of making a transition from s to s'

  - we require that $\Sigma_{s' \in S} \mathbf{P}(s,s')=1$ for all states $s \in S$
- $L : S \rightarrow 2^{AP}$ is the labelling function where AP is a set of atomic propositions

i.e. require that the total probability of making a transition from any state is 1

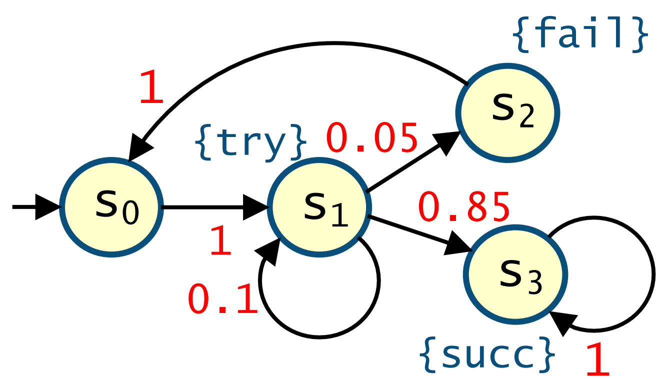

Simple DTMC example

Modelling a very simple communication protocol

- after one step, process starts trying to send a message
- with probability 0.1, channel unready so wait a step
- with probability 0.85, send message successfully and stop
- with probability 0.05, sending fails, then in next step it restarts

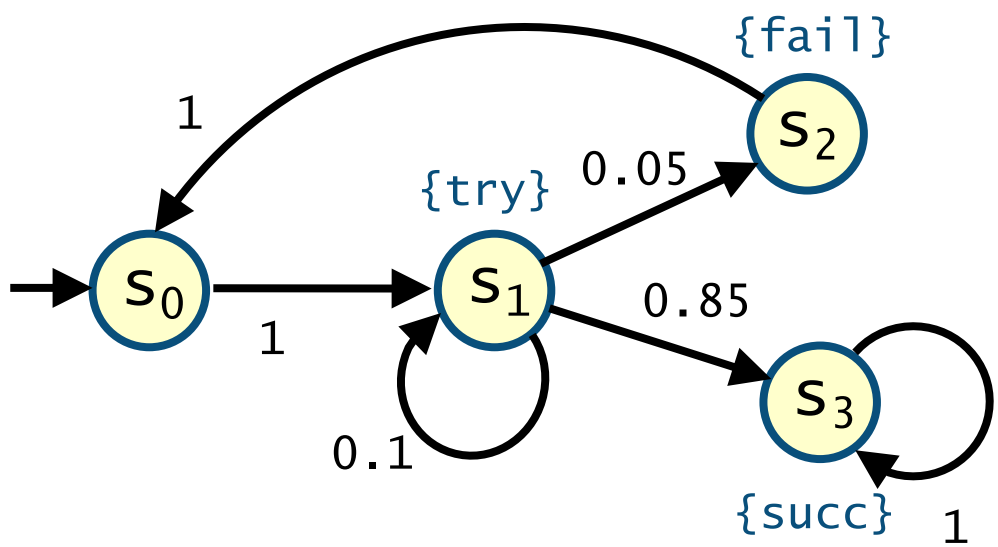

Simple DTMC example

States: $S = \{s_0, s_1, s_2, s_3\}$

Initial state: s_0

Probability transition matrix:

$$P = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0.1 & 0.05 & 0.85 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

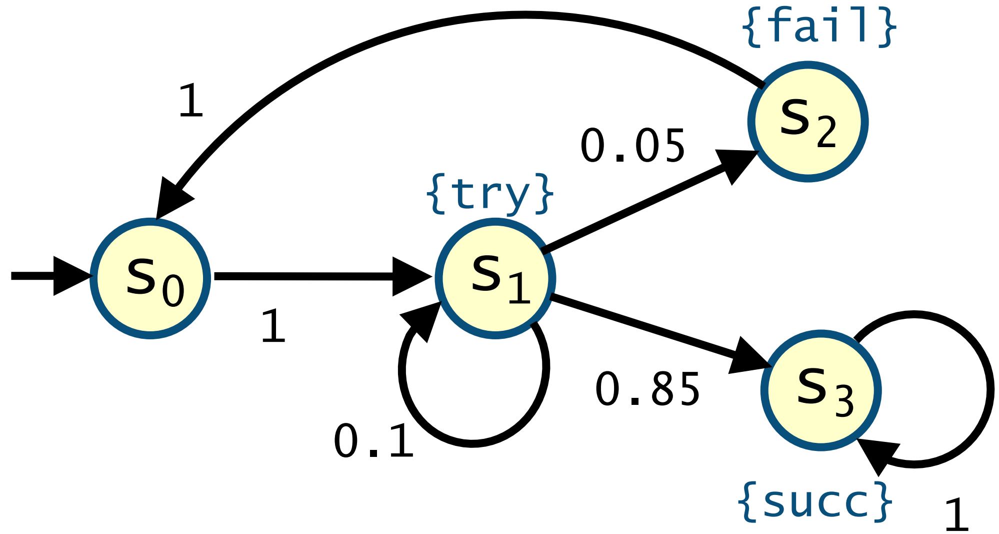

Atomic propositions: $AP = \{try, fail, succ\}$

– labelling: $L(s_0)=\emptyset$, $L(s_1)=\{try\}$, $L(s_2)=\{fail\}$ and $L(s_3)=\{succ\}$

DTMC example 2 – Coins and dice

Modelling a 6-sided die using a fair coin

- algorithm due to Knuth/Yao
- start at s_0, flip a coin
- upper branch when flip H
- lower branch when flip T
- repeat until value chosen

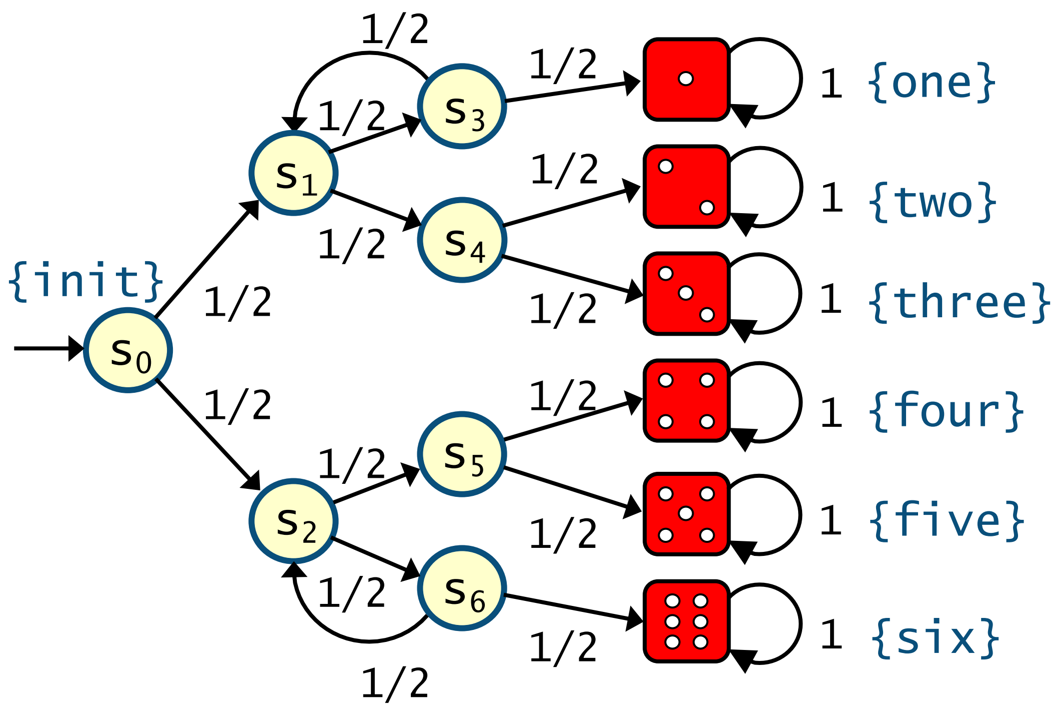

DTMC example 2 - Coins and dice

Is this model correct?

- e.g. probability of obtaining a 4 equals 1/6
- is it guaranteed to terminate?

How efficient is it?

- what is the probability of needing more than four coin flips?
- on average, how many coin flips are needed?

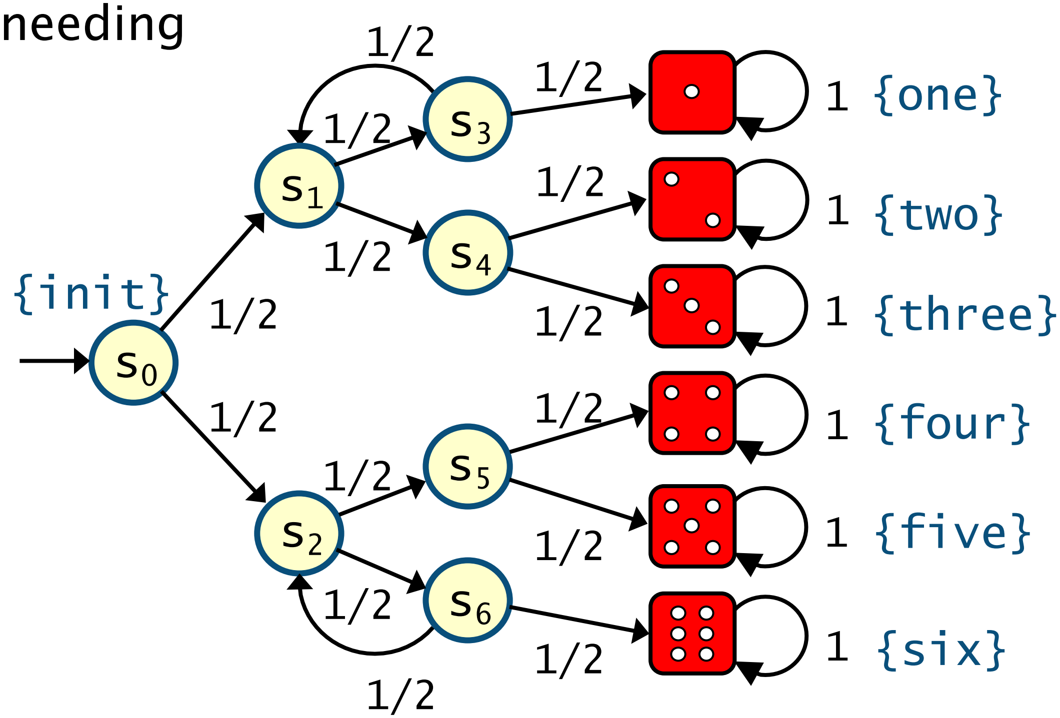

DTMC example 2 – Coins and dice

Is this model correct?

- e.g. probability of obtaining a 4 equals 1/6
- is it guaranteed to terminate?

How efficient is it?

- what is the probability of needing more than four coin flips?
- on average, how many coin flips are needed?

Probabilistic model checking provides a framework for answering these kinds of questions

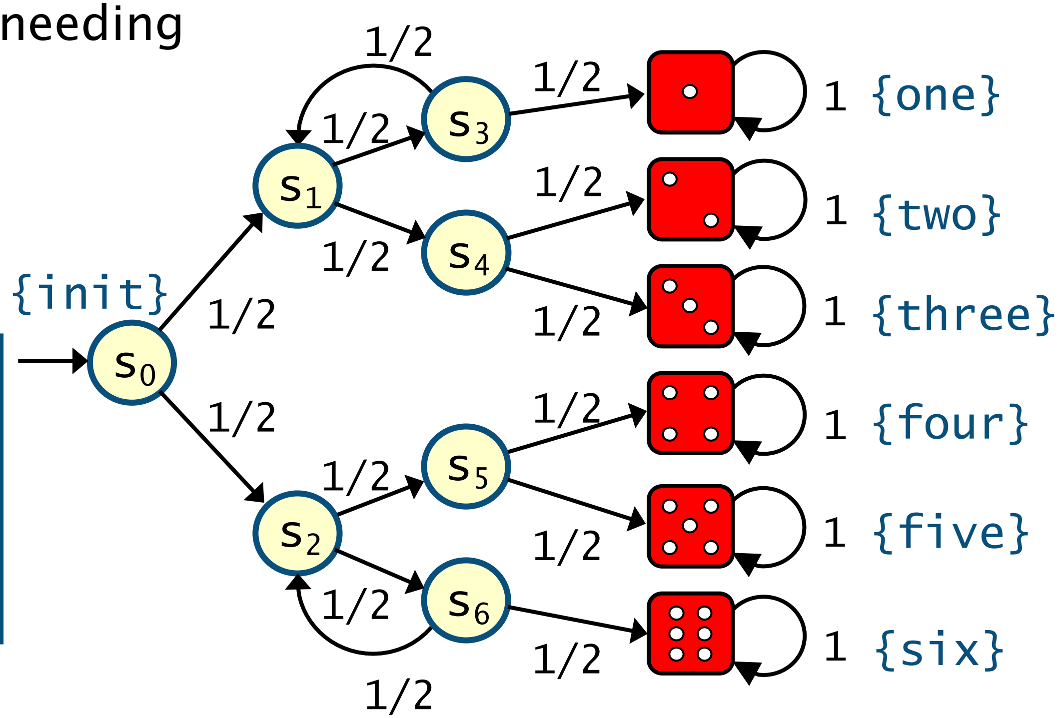

Discrete time Markov chains

Discrete-time Markov chains (DTMCs)

Paths and probabilities for DTMCs

Probabilistic reachability for DTMCs

Rewards and expected reachability for DTMCs

First some probability basics

Need an experiment…

- the sample space is the set of possible outcomes of the experiment
- an event is a subset of the sample space
- the probability of an event is the degree of certainty an event will occur

Example: toss two coins

- sample space: {(H,H), (H,T), (T,H), (T,T)}
- event: "at least one H"
- probability: $1/2 + (1/2)\cdot(1/2) = 3/4$

Example: toss a coin infinitely often

- sample space: set of infinite sequences of H/T
- event: "H in the first 3 throws"
- probability: $1/2 + (1/2)\cdot(1/2) + (1/2)\cdot(1/2)\cdot(1/2) = 7/8$

Probability space $(\Omega, \Sigma, \mathbf{Prob})$

Sample space Ω is an arbitrary non-empty set

Event set Σ is family of subsets of Ω which is

- closed under complementation
  - if A is in Σ, then the complement $\Omega \setminus A$ is in Σ
- closed under countable union
  - if A_i is in Σ for $i \in \mathbb{N}$, then the union $\cup_i A_i$ is in Σ
- contains the empty set ($\emptyset$ is in Σ)

Elements of Σ are called **measurable sets** and Σ a **σ-algebra** on Ω

Probability measure **Prob** is a function $\mathbf{Prob}: \Sigma \to [0,1]$ such that

- $\mathbf{Prob}(\Omega) = 1$
- $\mathbf{Prob}(\cup A_i) = \Sigma_i \ \mathbf{Prob}(A_i)$

  for any disjoint family of measurable sets $A_1, A_2, \ldots$

Probability space – Simple example

Sample space $\Omega = \mathbb{N} = \{ 0, 1, 2, 3, 4, \ldots \}$

- the natural numbers

Event set $\Sigma = \{ \emptyset, \text{“odd”}, \text{“even”}, \mathbb{N} \}$

- (closed under complement/countable union, contains $\emptyset$)
- e.g. “odd”$\cup$“even” $= \mathbb{N}$ and $\mathbb{N} \setminus$“odd” $=$ “even”

Probability measure Prob

- e.g. corresponding to picking a number uniformly at random
- **Prob**(“odd”)=1/2, **Prob**(“even”)=1/2, …

Back to DTMCs – Paths

A **path** in a DTMC represents an **execution** of the system

- i.e. one possible behaviour

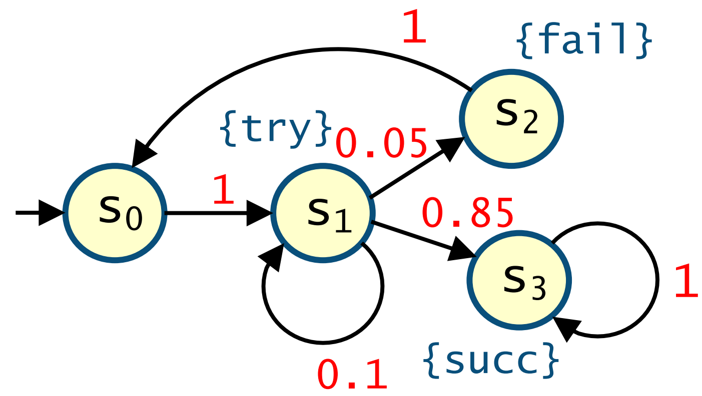

Formally:

- infinite sequence of states $s_0 s_1 s_2 s_3 \ldots$ such that $\mathbf{P}(s_i, s_{i+1}) > 0$ for all $i \geq 0$

Back to DTMCs – Paths

A path in a DTMC represents an execution of the system

- i.e. one possible behaviour

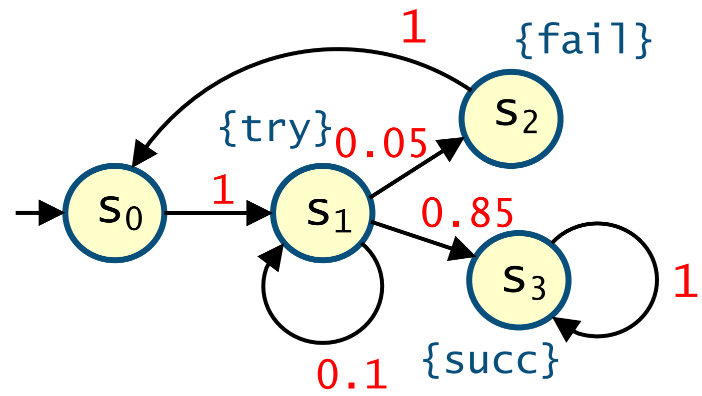

Formally:

- infinite sequence of states $s_0s_1s_2s_3\ldots$ such that $\mathbf{P}(s_i, s_{i+1})>0$ for all $i \geq 0$

Example execution:

- start, wait, fail, retry, start, succeed

Back to DTMCs – Paths

A path in a DTMC represents an execution of the system

- i.e. one possible behaviour

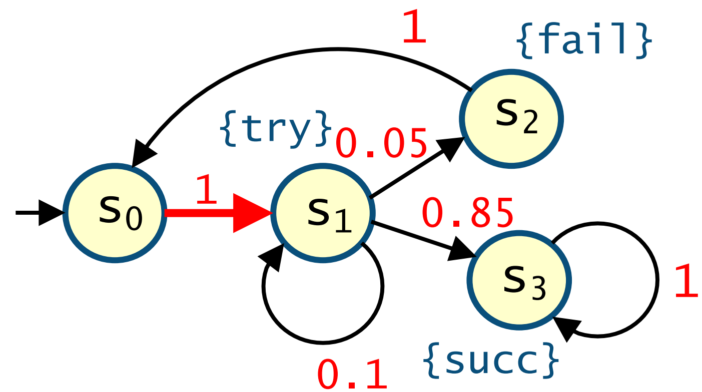

Formally:

- infinite sequence of states $s_0 s_1 s_2 s_3 \ldots$
  such that $P(s_i, s_{i+1}) > 0$ for all $i \geq 0$

Example execution:

- start, wait, fail, retry, start, succeed: $s_0 s_1$

Back to DTMCs – Paths

A path in a DTMC represents an execution of the system

- i.e. one possible behaviour

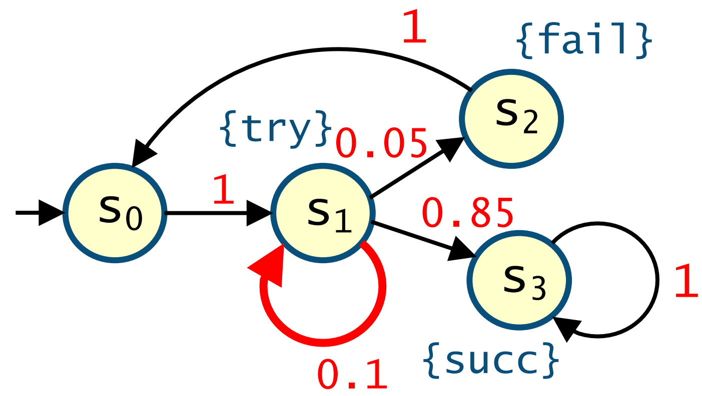

Formally:

- infinite sequence of states $s_0s_1s_2s_3\ldots$ such that $\mathbf{P}(s_i, s_{i+1})>0$ for all $i \geq 0$

Example execution:

- start, wait, fail, retry, start, succeed: $s_0s_1s_1$

Back to DTMCs – Paths

A path in a DTMC represents an execution of the system

- i.e. one possible behaviour

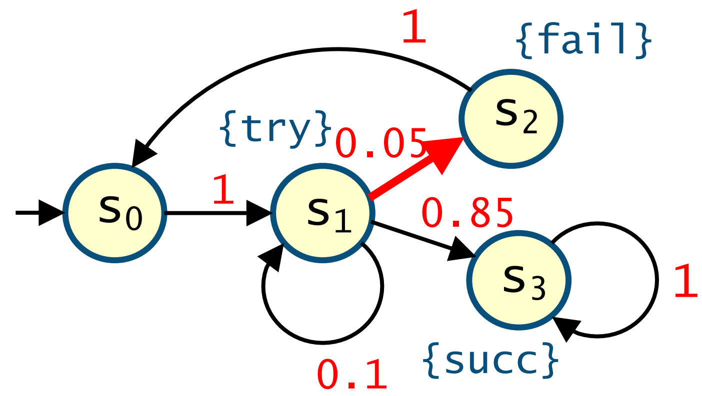

Formally:

- infinite sequence of states $s_0 s_1 s_2 s_3 \ldots$
  such that $P(s_i, s_{i+1}) > 0$ for all $i \geq 0$

Example execution:

- start, wait, fail, retry, start, succeed: $s_0 s_1 s_1 s_2$

Back to DTMCs – Paths

A **path** in a DTMC represents an **execution** of the system

- i.e. one possible behaviour

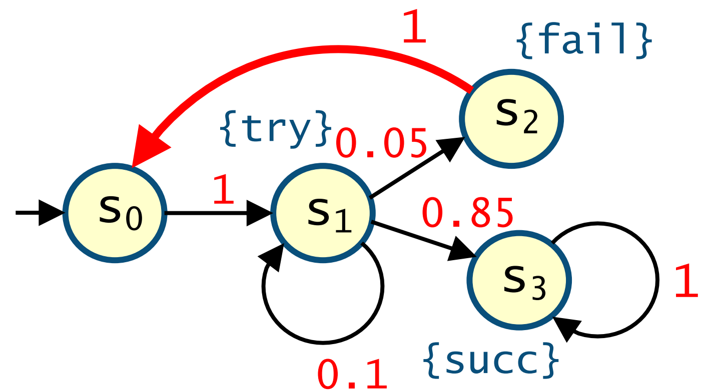

Formally:

- infinite sequence of states $s_0s_1s_2s_3\ldots$
  such that $\mathbf{P}(s_i,s_{i+1})>0$ for all $i\geq0$

Example execution:

- start, wait, fail, retry, start, succeed: $s_0s_1s_1s_2s_0$

Back to DTMCs – Paths

A **path** in a DTMC represents an **execution** of the system

- i.e. one possible behaviour

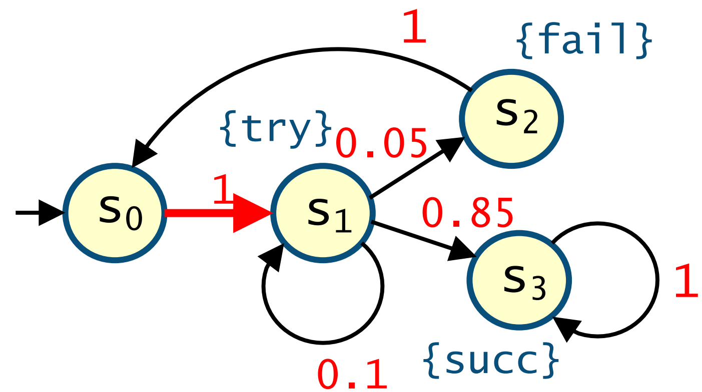

Formally:

- infinite sequence of states $s_0s_1s_2s_3\ldots$ such that $P(s_i,s_{i+1})>0$ for all $i\geq0$

Example execution:

- start, wait, fail, retry, start, succeed: $s_0s_1s_1s_2s_0s_1$

Back to DTMCs – Paths

A path in a DTMC represents an execution of the system

- i.e. one possible behaviour

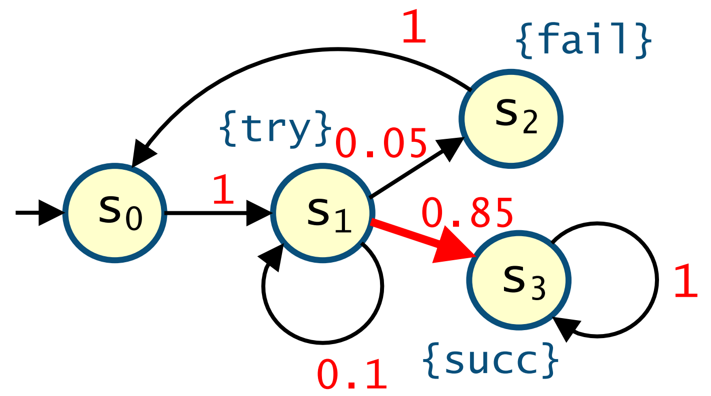

Formally:

- infinite sequence of states $s_0s_1s_2s_3\ldots$ such that $\mathbf{P}(s_i,s_{i+1})>0$ for all $i \geq 0$

Example execution:

- start, wait, fail, retry, start, succeed: $s_0s_1s_1s_2s_0s_1s_3\ldots$

Back to DTMCs – Paths

A path in a DTMC represents an execution of the system

- i.e. one possible behaviour

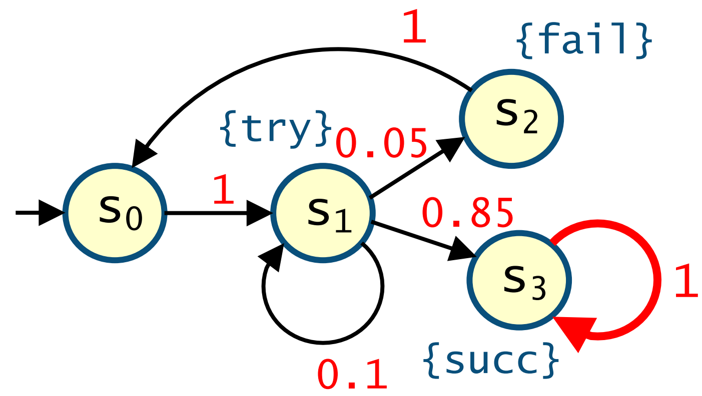

Formally:

- infinite sequence of states $s_0 s_1 s_2 s_3 \ldots$ such that $\mathbf{P}(s_i, s_{i+1}) > 0$ for all $i \geq 0$

Example execution:

- start, wait, fail, retry, start, succeed: $s_0 s_1 s_1 s_2 s_0 s_1 s_3 s_3 \ldots$

Probability space over paths

To reason about a DTMC when starting from some state s

- need to define a probability space over paths starting from the state s

Intuitively:

- sample space: infinite paths starting from the state s
- events: sets of infinite paths
- basic events: cylinder sets
- cylinder Cyl(ω) for a finite path ω equals the set of infinite paths that have ω as a prefix
- e.g. Cyl(ss_1s_2)

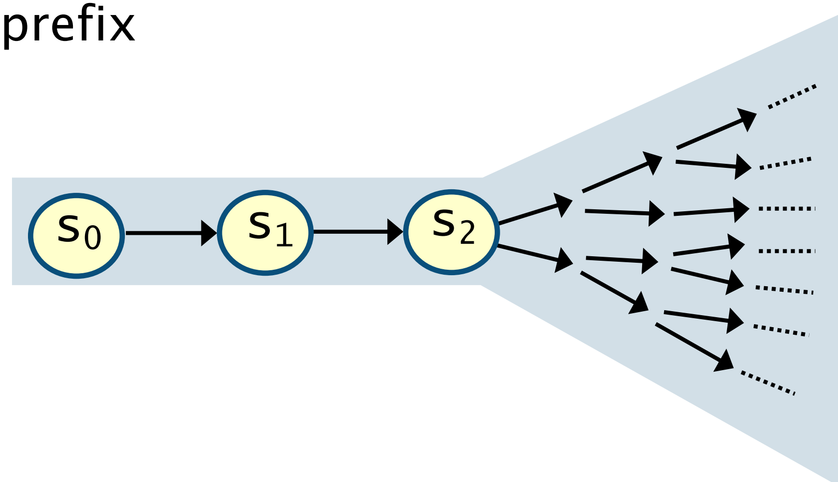

Probability space over paths

Probability space $(\mathtt{Path_s}, \Sigma_s, \mathtt{Prob_s})$

Sample space: all infinite paths starting from the state s

Probability space over paths

Probability space $(\mathtt{Path}_s, \Sigma_s, \mathtt{Prob}_s)$

Sample space: all infinite paths starting from the state **s**

Event set: least σ–algebra including the cylinder $\mathtt{Cyl}(\omega)$ of every finite path $\omega = s s_1 s_2 \ldots s_n$

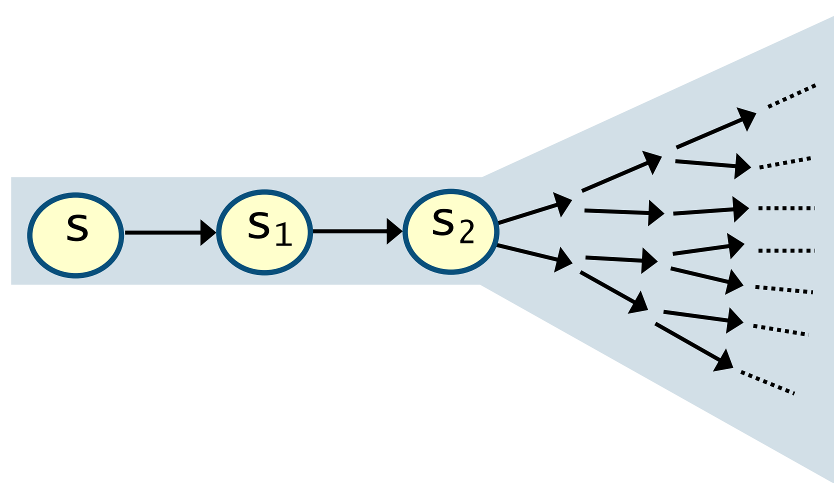

Probability space over paths

Probability space $(Path_s, \Sigma_s, Prob_s)$

Sample space: all infinite paths starting from the state s

Event set: least σ–algebra including the cylinder $Cyl(\omega)$ of every finite path $\omega = s s_1 s_2 \ldots s_n$

Probability measure: unique extension of function $Prob_s$ over cylinders where $Prob_s(Cyl(\omega)) = P(s,s_1) \cdot P(s_1,s_2) \cdots P(s_{n-1},s_n)$

probability of a cylinder given by multiplying the probability of each transition of the finite path

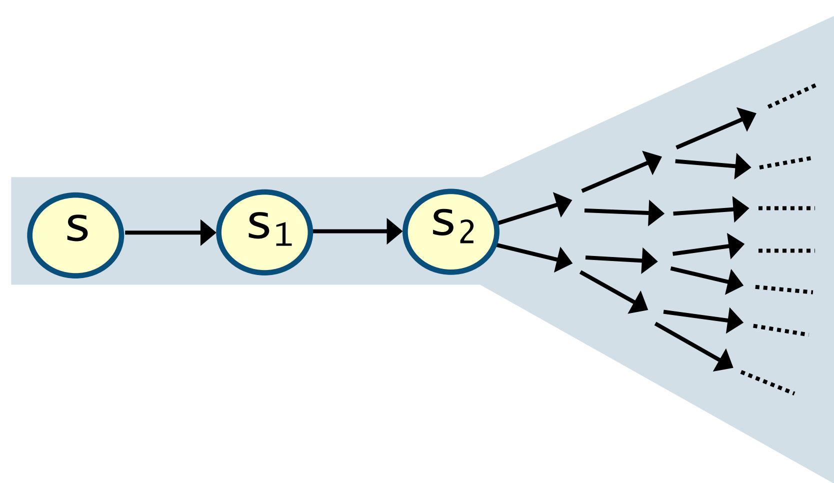

Paths and probabilities – Example

Paths where sending fails the first time

– all paths starting $s_0s_1s_2$, i.e. the cylinder $\mathtt{Cyl}(s_0s_1s_2)$

Probability:

$$\mathbf{Prob}_{s0}(\mathtt{Cyl}(s_0s_1s_2)) = \mathbf{P}(s_0,s_1)\cdot\mathbf{P}(s_1,s_2) = 1 \cdot 0.05 = 0.05$$

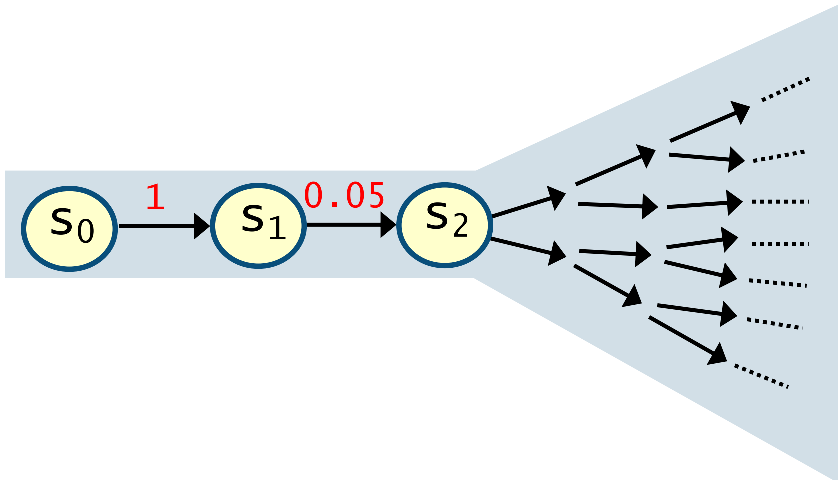

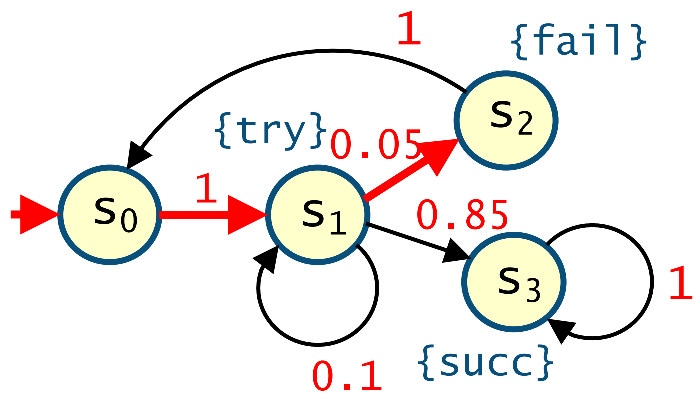

Paths and probabilities – Example

Paths which are eventually successful with no failures

- infinite paths of the form $s_0\ (s_1)^*\ s_3{}^{\omega}$
- i.e. the (disjoint) union of the cylinders:

$$\mathrm{Cyl}(s_0s_1s_3) \cup \mathrm{Cyl}(s_0s_1s_1s_3) \cup \mathrm{Cyl}(s_0s_1s_1s_1s_3) \cup \ldots$$

Probability:

$$\mathbf{Prob}_{s0}(\mathrm{Cyl}(s_0s_1s_3) \cup \mathrm{Cyl}(s_0s_1s_1s_3) \cup \mathrm{Cyl}(s_0s_1s_1s_1s_3) \cup \ldots)$$

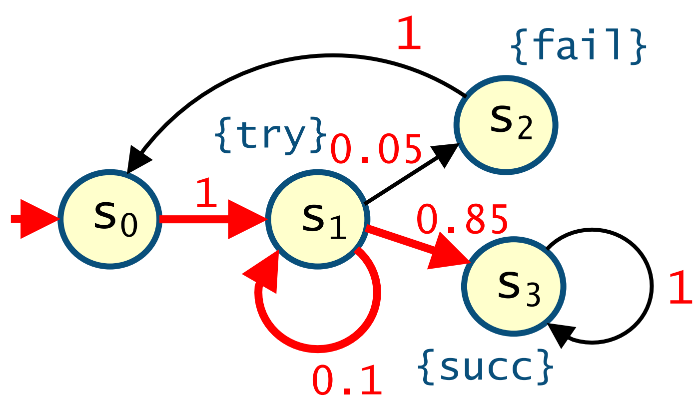

Paths and probabilities – Example

Paths which are eventually successful with no failures

- infinite paths of the form $s_0 \ (s_1)^* \ s_3^\omega$
- i.e. the (disjoint) union of the cylinders:

$$\mathrm{Cyl}(s_0s_1s_3) \cup \mathrm{Cyl}(s_0s_1s_1s_3) \cup \mathrm{Cyl}(s_0s_1s_1s_1s_3) \cup \ldots$$

Probability:

$$\mathbf{Prob}_{s0}(\mathrm{Cyl}(s_0s_1s_3) \cup \mathrm{Cyl}(s_0s_1s_1s_3) \cup \mathrm{Cyl}(s_0s_1s_1s_1s_3) \cup \ldots) =$$

$$\mathbf{Prob}_{s0}(\mathrm{Cyl}(s_0s_1s_3))+\mathbf{Prob}_{s0}(\mathrm{Cyl}(s_0s_1s_1s_3))+\mathbf{Prob}_{s0}(\mathrm{Cyl}(s_0s_1s_1s_1s_3))+\cdots$$

since the sets are disjoint

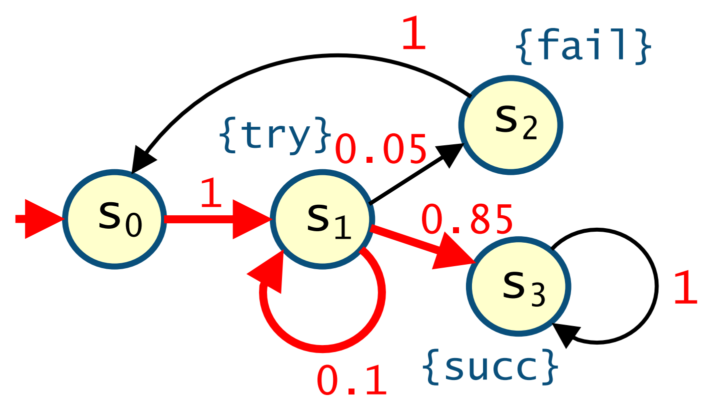

Paths and probabilities – Example

Paths which are eventually successful with no failures

- infinite paths of the form $s_0\ (s_1)^*\ s_3^\omega$
- i.e. the (disjoint) union of the cylinders:

$$\mathrm{Cyl}(s_0s_1s_3) \cup \mathrm{Cyl}(s_0s_1s_1s_3) \cup \mathrm{Cyl}(s_0s_1s_1s_1s_3) \cup \ldots$$

Probability:

$$\mathbf{Prob}_{s0}(\mathrm{Cyl}(s_0s_1s_3) \cup \mathrm{Cyl}(s_0s_1s_1s_3) \cup \mathrm{Cyl}(s_0s_1s_1s_1s_3) \cup \ldots\ ) =$$

$$\mathbf{Prob}_{s0}(\mathrm{Cyl}(s_0s_1s_3))+\mathbf{Prob}_{s0}(\mathrm{Cyl}(s_0s_1s_1s_3))+\mathbf{Prob}_{s0}(\mathrm{Cyl}(s_0s_1s_1s_1s_3))+\cdots$$

$$= 1\cdot 0.85 + 1\cdot(0.1)\cdot 0.85 + 1\cdot(0.1\cdot 0.1)\cdot 0.85 + \cdots + 1\cdot(0.1)^n\cdot 0.85 + \cdots$$

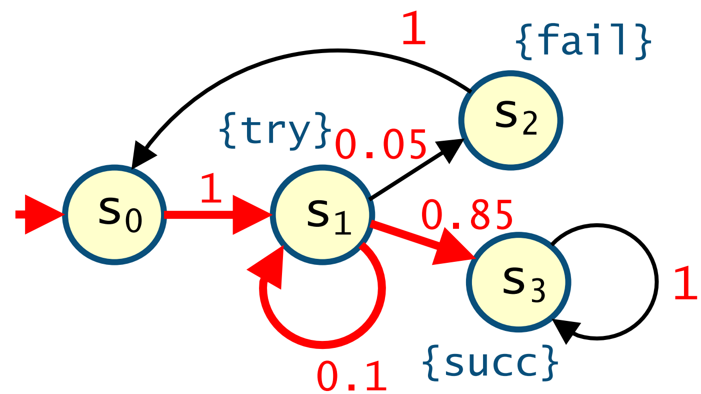

Paths and probabilities – Example

Paths which are eventually successful with no failures

- infinite paths of the form $s_0\ (s_1)^*\ s_3^{\omega}$
- i.e. the (disjoint) union of the cylinders:

  $Cyl(s_0s_1s_3) \cup Cyl(s_0s_1s_1s_3) \cup Cyl(s_0s_1s_1s_1s_3) \cup \ldots$

Probability:

$\mathbf{Prob}_{s0}(Cyl(s_0s_1s_3) \cup Cyl(s_0s_1s_1s_3) \cup Cyl(s_0s_1s_1s_1s_3) \cup \ldots) =$

$\mathbf{Prob}_{s0}(Cyl(s_0s_1s_3))+\mathbf{Prob}_{s0}(Cyl(s_0s_1s_1s_3))+\mathbf{Prob}_{s0}(Cyl(s_0s_1s_1s_1s_3))+\cdots$

$= 1\cdot 0.85 + 1\cdot(0.1)\cdot 0.85 + 1\cdot(0.1\cdot 0.1)\cdot 0.85 + \cdots + 1\cdot(0.1)^n\cdot 0.85 + \cdots$

$= 0.85\cdot(1 + 0.1 + \cdots + 0.1^n + \cdots)$

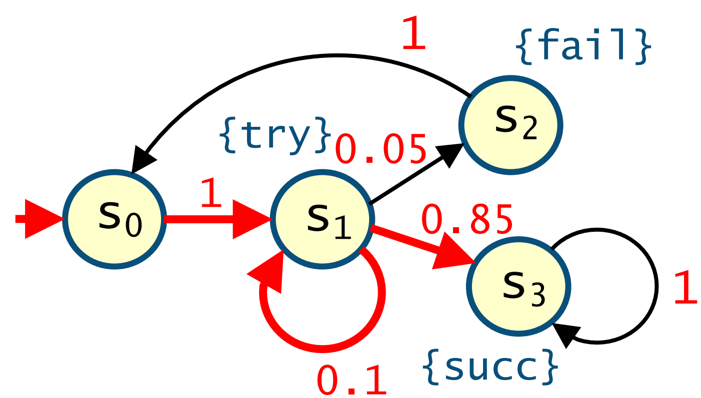

Paths and probabilities – Example

Paths which are eventually successful with no failures

- infinite paths of the form $s_0\ (s_1)^*\ s_3^{\omega}$
- i.e. the (disjoint) union of the cylinders:

$$\mathtt{Cyl}(s_0s_1s_3) \cup \mathtt{Cyl}(s_0s_1s_1s_3) \cup \mathtt{Cyl}(s_0s_1s_1s_1s_3) \cup \ldots$$

Probability:

$$\mathbf{Prob}_{s0}(\mathtt{Cyl}(s_0s_1s_3) \cup \mathtt{Cyl}(s_0s_1s_1s_3) \cup \mathtt{Cyl}(s_0s_1s_1s_1s_3) \cup \ldots) =$$

$$\mathbf{Prob}_{s0}(\mathtt{Cyl}(s_0s_1s_3))+\mathbf{Prob}_{s0}(\mathtt{Cyl}(s_0s_1s_1s_3))+\mathbf{Prob}_{s0}(\mathtt{Cyl}(s_0s_1s_1s_1s_3))+\cdots$$

$$= 1\cdot 0.85 + 1\cdot(0.1)\cdot 0.85 + 1\cdot(0.1\cdot 0.1)\cdot 0.85 + \cdots + 1\cdot(0.1)^n\cdot 0.85 + \cdots$$

$$= 0.85\cdot(1 + 0.1 + \cdots + 0.1^n + \cdots)$$

$$= 0.85 \cdot 10/9 = 17/18$$

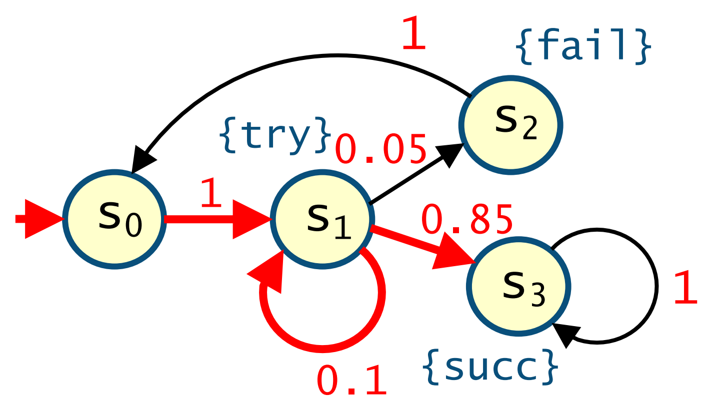

Discrete time Markov chains

Discrete-time Markov chains (DTMCs)

Paths and probabilities for DTMCs

Probabilistic reachability for DTMCs

Rewards and expected reachability for DTMCs

Probabilistic reachability

Fundamental property of DTMCs: probabilistic reachability

- probability of a path reaching some target set of states T
  - **P**(s,T) probability of reaching T from state s
  - vector: **P**(T) values for all states of a DTMC
- e.g. "probability of the algorithm terminating successfully?"
- e.g. "probability that an error occurs during execution?"

Dual of reachability: invariance

- probability of remaining within some class of states
- **Prob**("remain in set I") = 1 - **Prob**("reach set S\I")
- e.g. "probability that an error never occurs"

Also other variants of reachability

- step-bounded, constrained ("until"), …

Probabilistic reachability – Example

Modelling a 6-sided die using a fair coin

- algorithm due to Knuth/Yao
- start at s_0, toss a coin
- upper branch when H
- lower branch when T
- repeat until value chosen

Is this algorithm correct?

- e.g. probability of reaching "4" equals 1/6
- event: all possible ways of reaching "4" from s_0

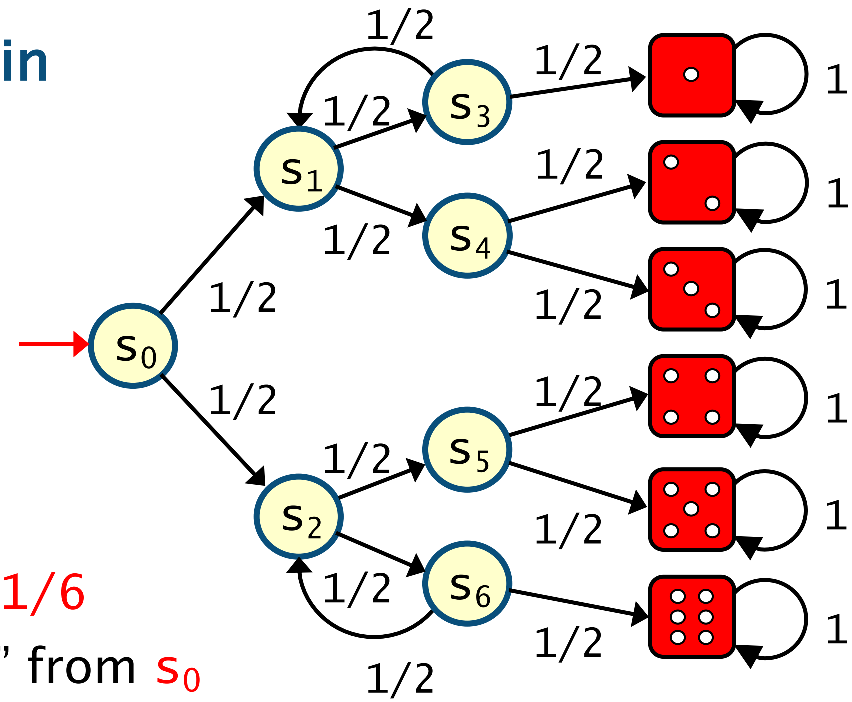

Probabilistic reachability – Example

Modelling a 6-sided die using a fair coin

- algorithm due to Knuth/Yao
- start at s_0, toss a coin
- upper branch when H
- lower branch when T
- repeat until value chosen

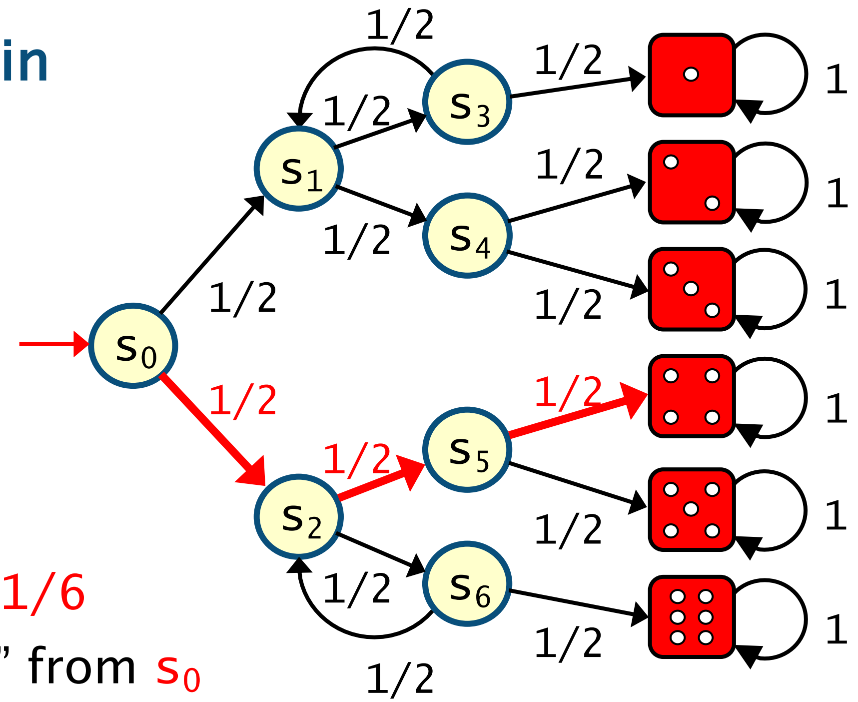

Is this algorithm correct?

- e.g. probability of reaching "4" equals 1/6
- event: all possible ways of reaching "4" from s_0
- ways of reaching "4" :

  THH,
- probability of reaching "4" :

  $(1/2)^3$ +

Probabilistic reachability – Example

Modelling a 6-sided die using a fair coin

- algorithm due to Knuth/Yao
- start at s_0, toss a coin
- upper branch when H
- lower branch when T
- repeat until value chosen

Is this algorithm correct?

- e.g. probability of reaching "4" equals 1/6
- event: all possible ways of reaching "4" from s_0
- ways of reaching "4" :

  THH, TTTHH,

- probability of reaching "4" :

  $(1/2)^3 + (1/2)^5 +$

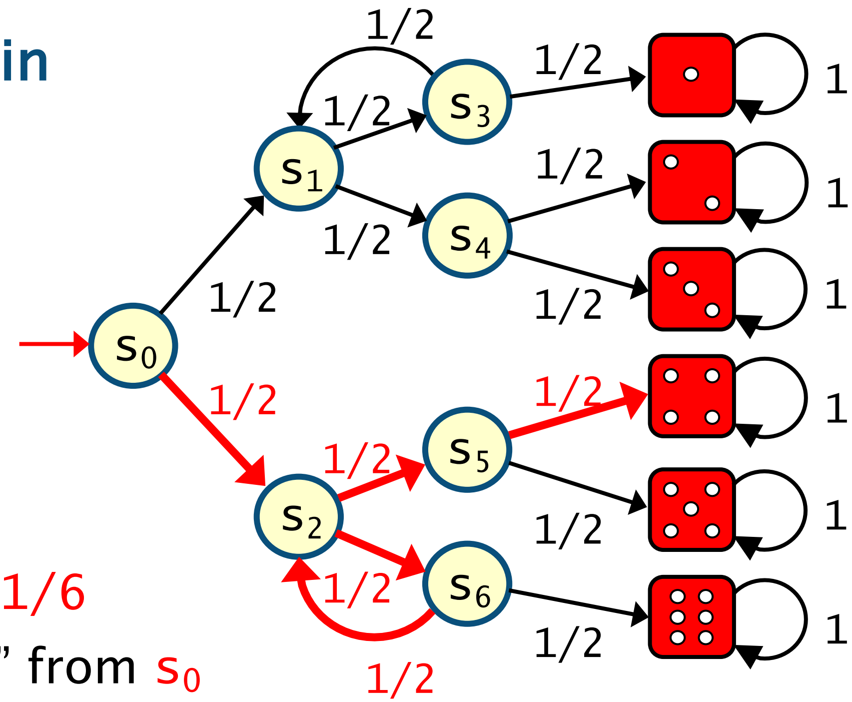

Probabilistic reachability – Example

Modelling a 6-sided die using a fair coin

- algorithm due to Knuth/Yao
- start at s_0, toss a coin
- upper branch when H
- lower branch when T
- repeat until value chosen

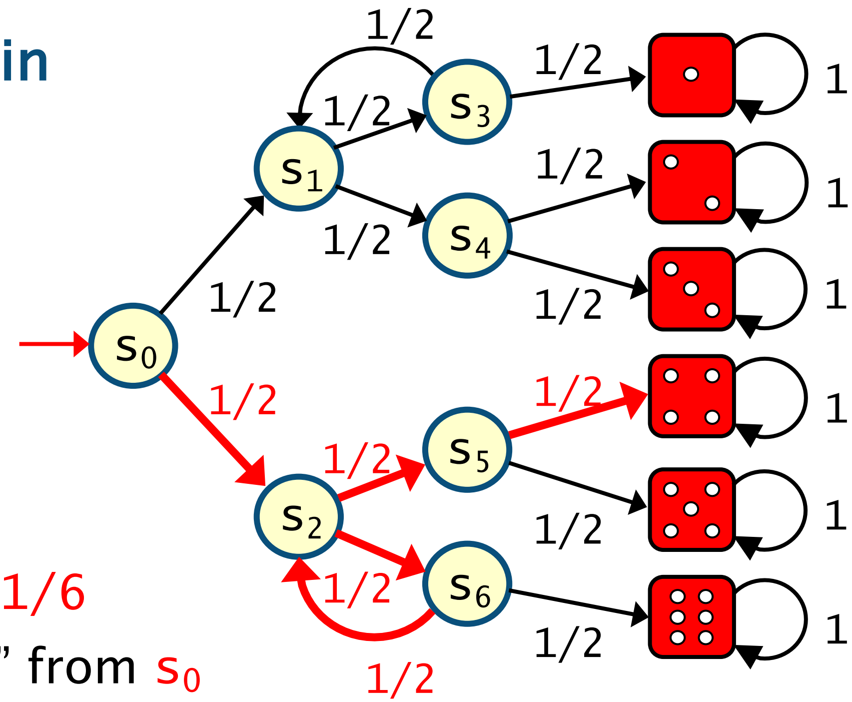

Is this algorithm correct?

- e.g. probability of reaching "4" equals 1/6
- event: all possible ways of reaching "4" from s_0
- ways of reaching "4" :

  THH, TTTHH, TTTTTHH, …
- probability of reaching "4" :

  $(1/2)^3 + (1/2)^5 + (1/2)^7 +$

Probabilistic reachability – Example

Modelling a 6-sided die using a fair coin

- algorithm due to Knuth/Yao
- start at s_0, toss a coin
- upper branch when H
- lower branch when T
- repeat until value chosen

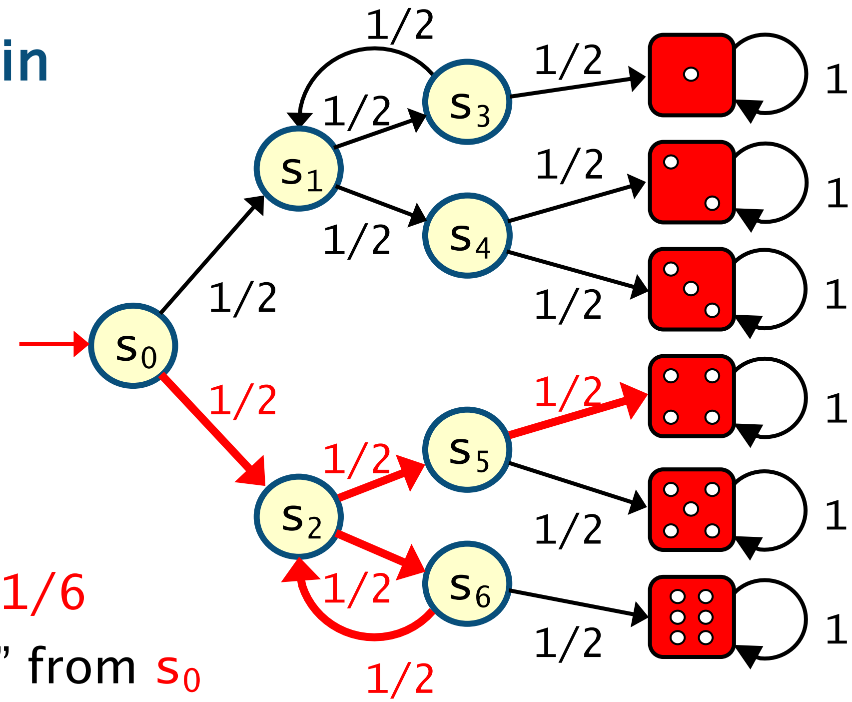

Is this algorithm correct?

- e.g. probability of reaching "4" equals 1/6
- event: all possible ways of reaching "4" from s_0
- ways of reaching "4" :

  THH, TTTHH, TTTTTHH, … , $T(TT)^nHH$, …

- probability of reaching "4" :

$$(1/2)^3 + (1/2)^5 + (1/2)^7 + \cdots + (1/2)^{2n+3} + \cdots = 1/6$$

Probabilistic reachability – Computation

Computing an infinite sum not not feasible in practice

Alternative to calculate P(s,T): derive a linear equation system

- calculate probabilities for all states $s \in S$ simultaneously

Let x_s denote the probability of reaching T from state s

Probabilistic reachability – Computation

Computing an infinite sum not not feasible in practice

Alternative to calculate **P(s,T)**: derive a linear equation system

- calculate probabilities for all states $s \in S$ simultaneously

Let x_s denote the probability of reaching **T** from state **s**

- if $s \in T$, then $x_s = 1$

if the state is in the target set then the probability of reaching the target is 1

Probabilistic reachability – Computation

Computing an infinite sum not not feasible in practice

Alternative to calculate $P(s,T)$: derive a linear equation system

– calculate probabilities for all states $s \in S$ simultaneously

Let x_s denote the probability of reaching T from state s

– if $s \in T$, then $x_s = 1$

– if T is not reachable from s, then $x_s = 0$

i.e. no (finite) path from s to a state in T

if one cannot reach the target, then the probability of reaching the target is 0

Probabilistic reachability – Computation

Computing an infinite sum not not feasible in practice

Alternative to calculate $P(s,T)$: derive a linear equation system

– calculate probabilities for all states $s \in S$ simultaneously

Let x_s denote the probability of reaching T from state s

– if $s \in T$, then $x_s = 1$

– if T is not reachable from s, then $x_s = 0$

i.e. no (finite) path from s to a state in T

– otherwise $x_s = \Sigma_{s' \in S} P(s,s') \cdot x_{s'}$

probability defined recursively using the transition probabilities: summation over all states s' of the probability of making a transition to s' multiplied by the probability of reaching the target from s'

Probabilistic reachability – Computation

Can view as a least fixed point computation over vectors $y \in [0,1]^S$

- consider the function $F : [0,1]^S \rightarrow [0,1]^S$ where

$$F(y)(s) = \begin{cases} 1 & \text{if } s \in T \\ \sum_{s' \in S} P(s,s') \cdot y(s') & \text{otherwise} \end{cases}$$

a vector of probabilities in each state

If we let $x^{(0)}=0$ and $x^{(n+1)}=F(x^{(n)})$ then we have that

- $x^{(0)} \leq x^{(1)} \leq x^{(2)} \leq x^{(3)} \leq \ldots$
- $P(T) = \lim_{n\rightarrow\infty} x^{(n)}$
  - recall $P(T)$ is the vector of probabilities $(P(s,T))_{s\in S}$
- $P(T)$ is the least fixed point of F

Probabilistic reachability – Example

Probability of reaching **{4}**

- i.e. tossing a four

x_s denotes the probability of reaching T from s

- if $s \in T$, then $x_s = 1$
- if T is not reachable from s, then $x_s = 0$
- otherwise $x_s = \Sigma_{s' \in S}\ P(s,s') \cdot x_{s'}$

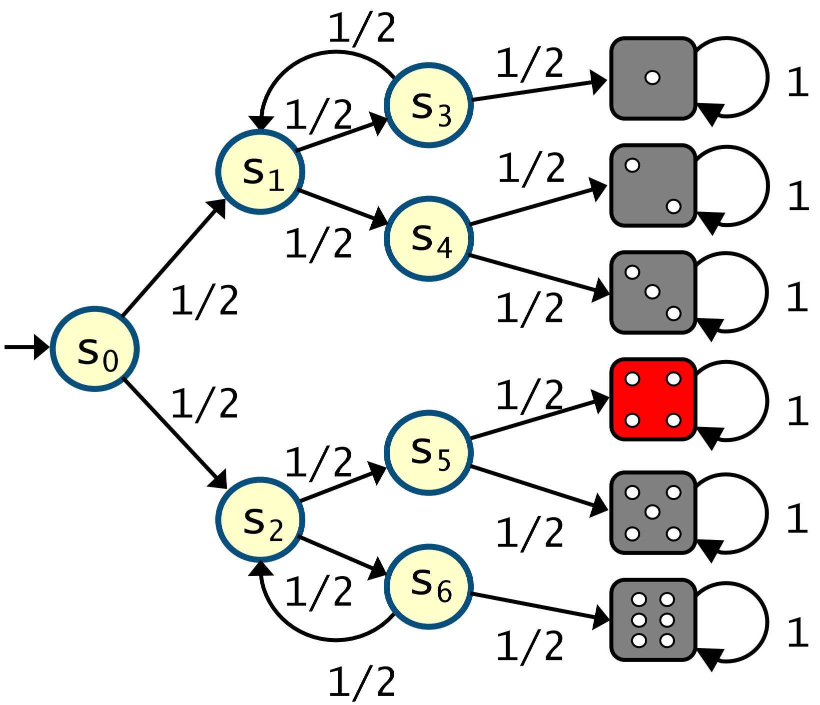

Probabilistic reachability – Example

Probability of reaching {4}

– i.e. tossing a four

x_s denotes the probability of reaching T from s

- if $s \in T$, then $x_s = 1$
- if T is not reachable from s, then $x_s = 0$
- otherwise $x_s = \Sigma_{s' \in S} \mathbf{P}(s,s') \cdot x_{s'}$

- $x_4 = 1$

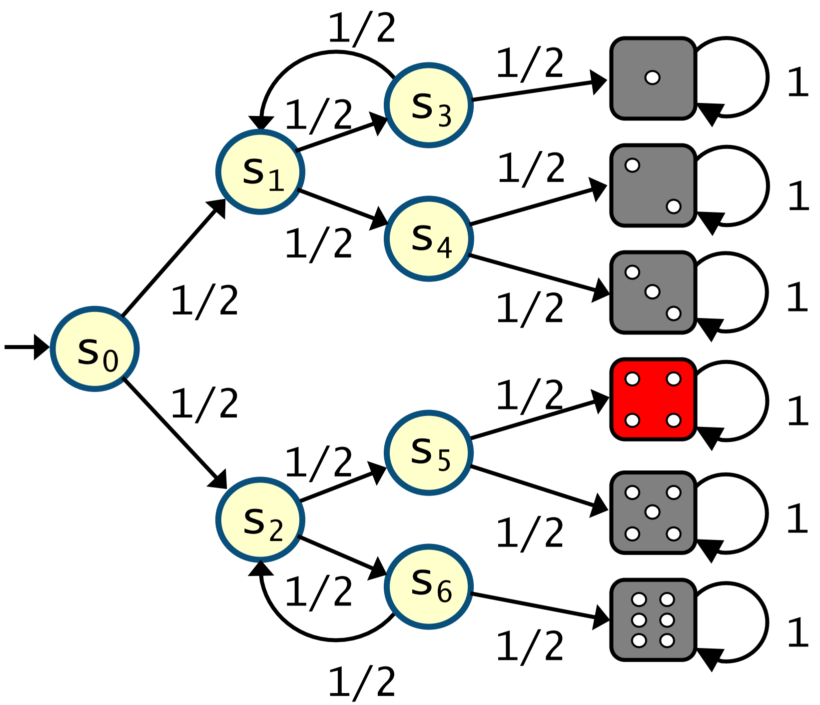

Probabilistic reachability – Example

Probability of reaching {4}

– i.e. tossing a four

x_s denotes the probability of reaching T from s

- if $s \in T$, then $x_s = 1$
- if T is not reachable from s, then $x_s = 0$
- otherwise $x_s = \sum_{s' \in S} \mathbf{P}(s,s') \cdot x_{s'}$

- $x_4 = 1$
- only s_0, s_2, s_5 and s_6 reach {4} therefore $x_s = 0$ for all other states

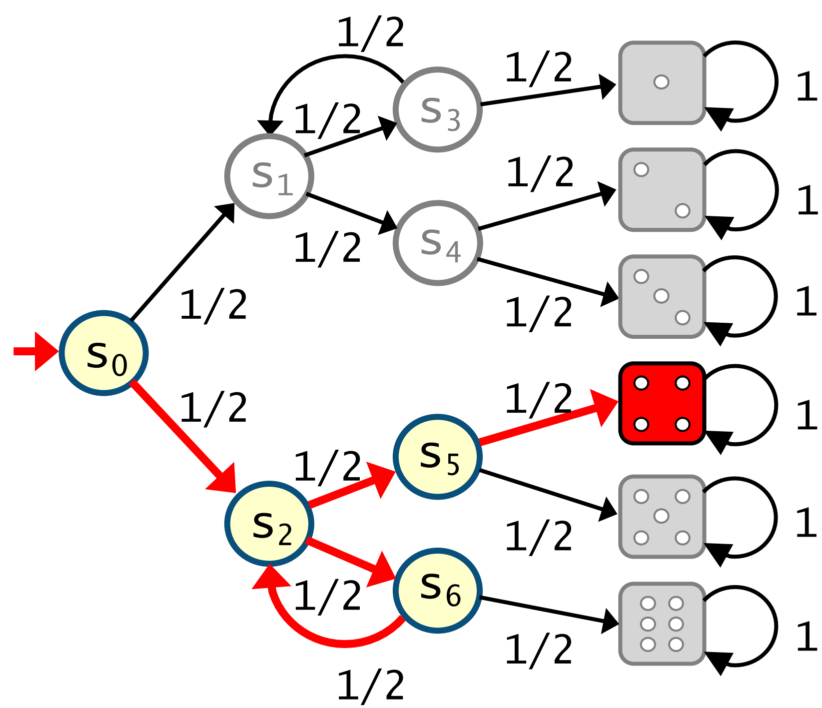

Probabilistic reachability – Example

Probability of reaching {4}

- i.e. tossing a four

x_s denotes the probability of reaching T from s

- if $s \in T$, then $x_s = 1$
- if T is not reachable from s, then $x_s = 0$
- otherwise $x_s = \Sigma_{s' \in S} \, P(s,s') \cdot x_{s'}$

- $x_4 = 1$
- only s_0, s_2, s_5 and s_6 reach {4}
  therefore $x_s = 0$ for all other states
- $x_{s0} = 1/2 \cdot 0 + 1/2 \cdot x_{s2}$
- $x_{s2} = 1/2 \cdot x_{s5} + 1/2 \cdot x_{s6}$
- $x_{s5} = 1/2 \cdot 1 + 1/2 \cdot 0$
- $x_{s6} = 1/2 \cdot x_{s2} + 1/2 \cdot 0$

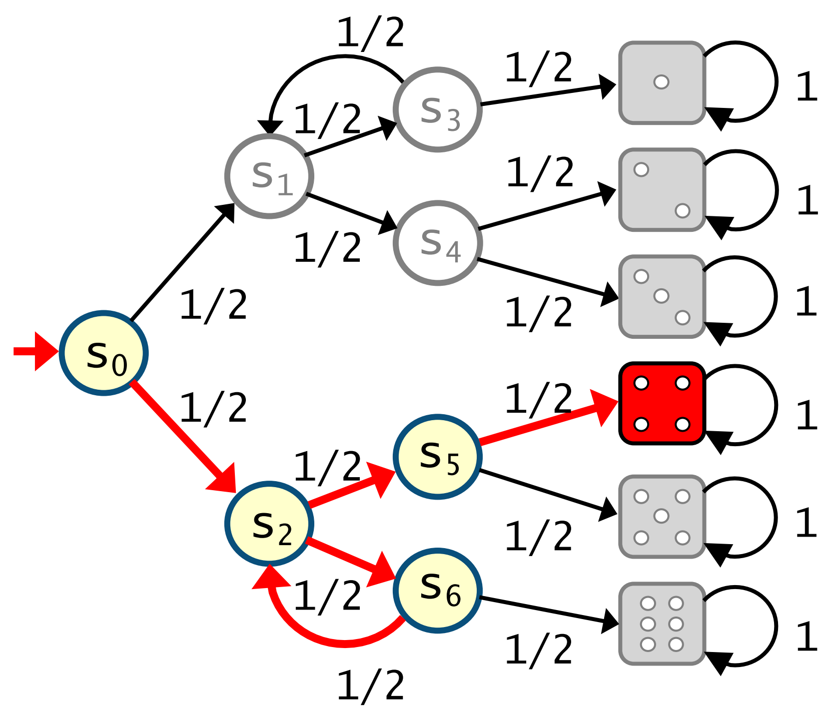

Probabilistic reachability – Example

Probability of reaching {4}

– i.e. tossing a four

x_s denotes the probability of reaching T from s

- if $s \in T$, then $x_s = 1$
- if T is not reachable from s, then $x_s = 0$
- otherwise $x_s = \Sigma_{s' \in S}\ \mathbf{P}(s,s') \cdot x_{s'}$

- $x_4 = 1$
- only s_0, s_2, s_5 and s_6 reach {4}
  therefore $x_s = 0$ for all other states
- $x_{s0} = 1/2 \cdot x_{s2}$
- $x_{s2} = 1/2 \cdot x_{s5} + 1/2 \cdot x_{s6}$
- $x_{s5} = 1/2$
- $x_{s6} = 1/2 \cdot x_{s2}$

simplifying

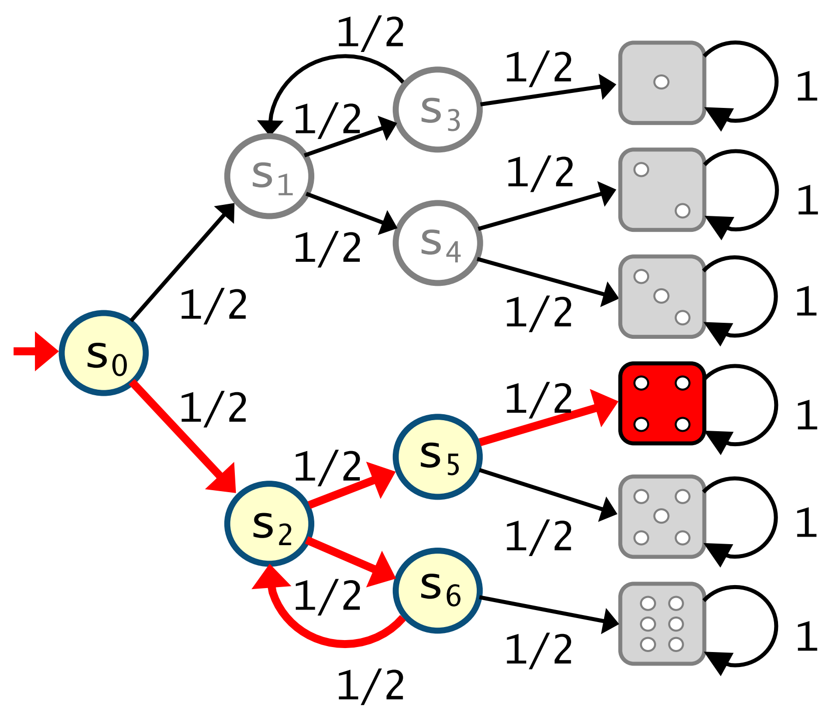

Probabilistic reachability – Example

Probability of reaching {4}

– i.e. tossing a four

x_s denotes the probability of reaching T from s

- if $s \in T$, then $x_s = 1$
- if T is not reachable from s, then $x_s = 0$
- otherwise $x_s = \Sigma_{s' \in S}\ \mathbf{P}(s,s') \cdot x_{s'}$

- $x_4 = 1$
- only s_0, s_2, s_5 and s_6 reach {4}
  therefore $x_s = 0$ for all other states
- $x_{s0} = 1/2 \cdot x_{s2}$
- $x_{s2} = (1/2)\cdot(1/2) + (1/2)\cdot(1/2 \cdot x_{s2})$
- $x_{s5} = 1/2$
- $x_{s6} = 1/2 \cdot x_{s2}$

substituting the values of x_{s5} and x_{s6} into x_{s2}

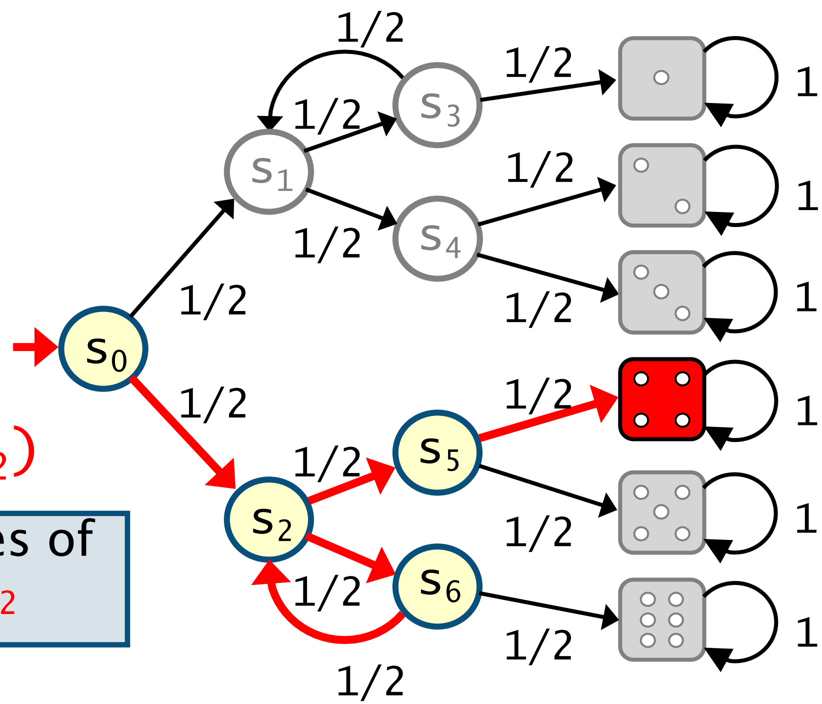

Probabilistic reachability – Example

Probability of reaching {4}

– i.e. tossing a four

x_s denotes the probability of reaching T from s

- if $s \in T$, then $x_s = 1$
- if T is not reachable from s, then $x_s = 0$
- otherwise $x_s = \Sigma_{s' \in S} \mathbf{P}(s,s') \cdot x_{s'}$

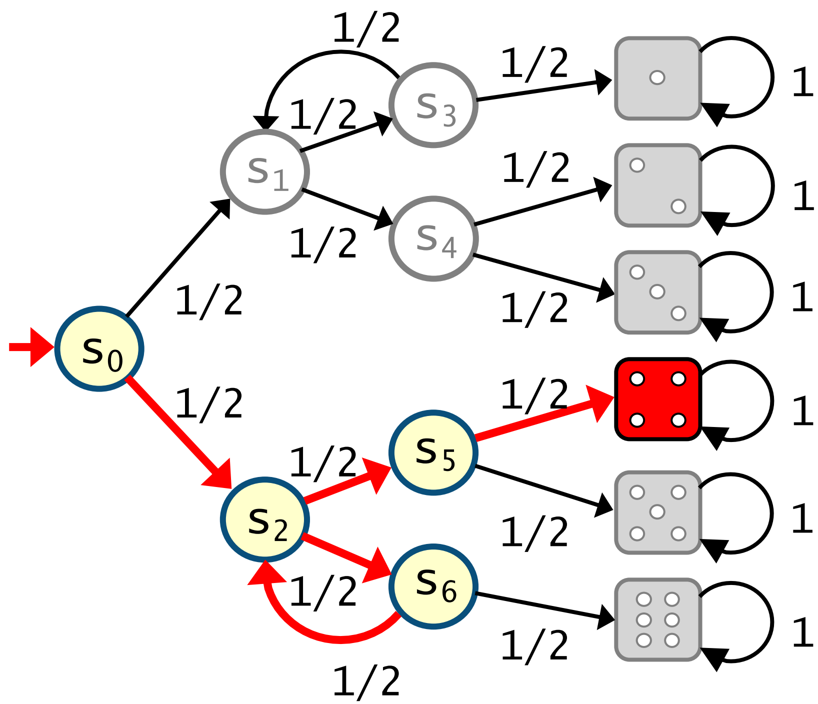

- $x_4 = 1$
- only s_0, s_2, s_5 and s_6 reach {4}
  therefore $x_s = 0$ for all other states
- $x_{s0} = 1/2 \cdot x_{s2}$
- $(3/4) \cdot x_{s2} = 1/4$
- $x_{s5} = 1/2$
- $x_{s6} = 1/2 \cdot x_{s2}$

simplifying

Probabilistic reachability – Example

Probability of reaching {4}

- i.e. tossing a four

x_s denotes the probability of reaching T from s

- if $s \in T$, then $x_s = 1$
- if T is not reachable from s, then $x_s = 0$
- otherwise $x_s = \Sigma_{s' \in S} \mathbf{P}(s,s') \cdot x_{s'}$

- $x_4 = 1$
- only s_0, s_2, s_5 and s_6 reach {4} therefore $x_s = 0$ for all other states
- $x_{s0} = 1/2 \cdot x_{s2}$
- $x_{s2} = (1/4)/(3/4) = 1/3$
- $x_{s5} = 1/2$
- $x_{s6} = 1/2 \cdot x_{s2}$

simplifying again

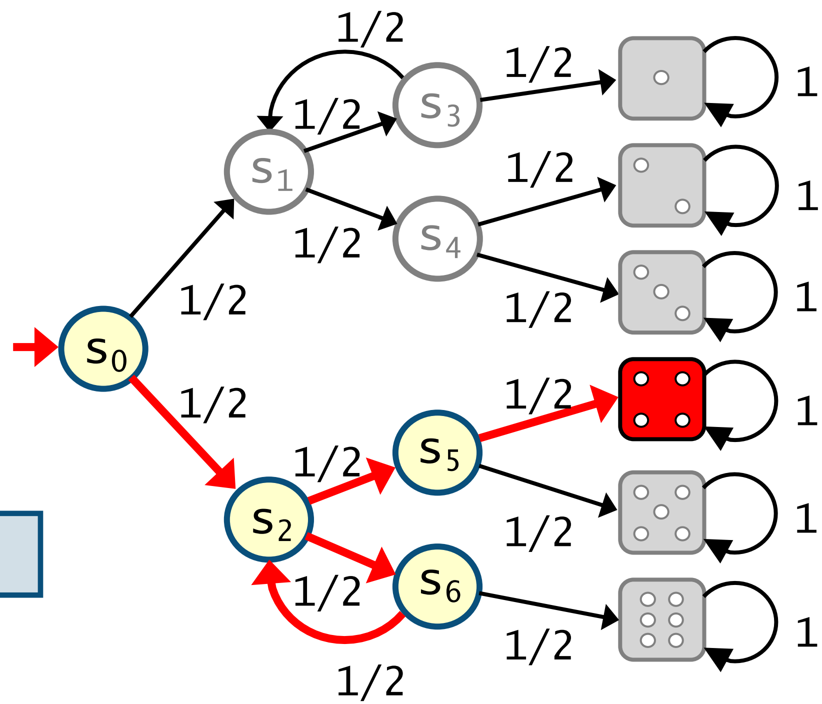

Probabilistic reachability – Example

Probability of reaching {4}

– i.e. tossing a four

x_s denotes the probability of reaching T from s

- if $s \in T$, then $x_s = 1$
- if T is not reachable from s, then $x_s = 0$
- otherwise $x_s = \Sigma_{s' \in S}\ \mathbf{P}(s,s') \cdot x_{s'}$

- $x_4 = 1$
- only s_0, s_2, s_5 and s_6 reach {4}
  therefore $x_s = 0$ for all other states
- $x_{s0} = 1/2 \cdot x_{s2} = (1/2) \cdot (1/3) = 1/6$
- $x_{s2} = 1/3$
- $x_{s5} = 1/2$
- $x_{s6} = 1/2 \cdot x_{s2} = (1/2) \cdot (1/3) = 1/6$

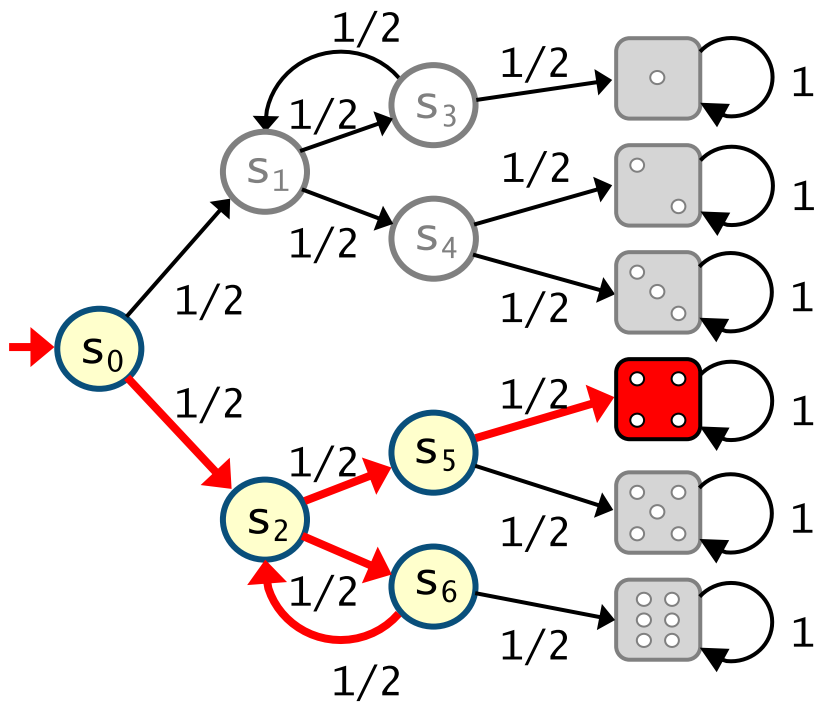

substituting the value of x_{s2} into the other equations

Probabilistic reachability – Example

Probability of reaching {4}

– i.e. tossing a four

x_s denotes the probability of reaching T from s

- if $s \in T$, then $x_s = 1$
- if T is not reachable from s, then $x_s = 0$
- otherwise $x_s = \Sigma_{s' \in S}$ $\mathbf{P}(s,s') \cdot x_{s'}$

- $x_4 = 1$
- only s_0, s_2, s_5 and s_6 reach {4} therefore $x_s = 0$ for all other states
- $x_{s0} = 1/2 \cdot x_{s2} = 1/2 \cdot 1/3 = 1/6$
- $x_{s2} = 1/3$
- $x_{s5} = 1/2$
- $x_{s6} = 1/2 \cdot x_{s2} = 1/2 \cdot 1/3 = 1/6$

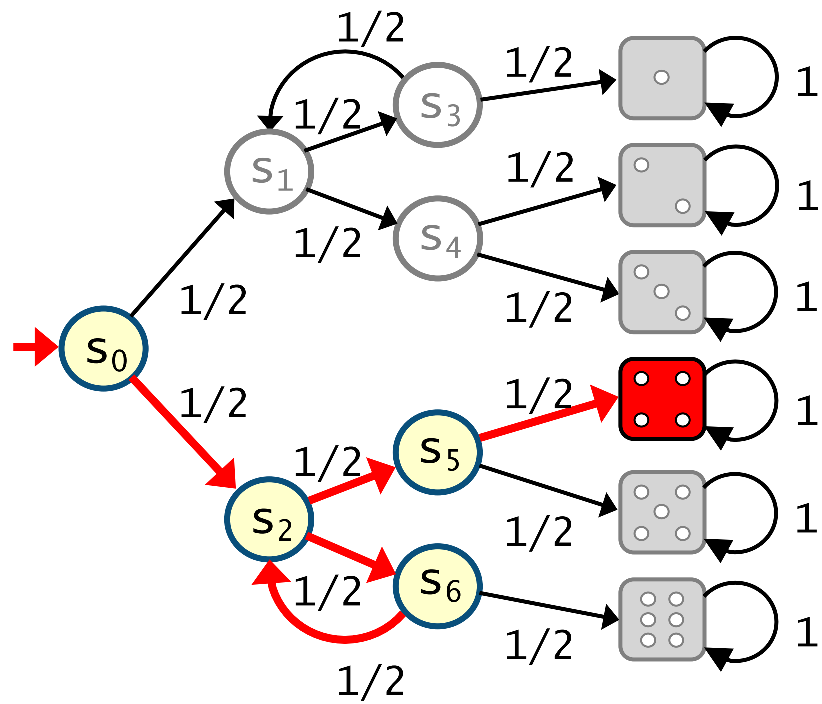

probability of a tossing a four is 1/6

Probabilistic reachability – Example

Probability of reaching {4}, as a least fixpoint

– $x^{(0)}=0$ and $x^{(n+1)}=F(x^{(n)})$ where

$$F(y)(s) = \begin{cases} 1 & \text{if } s \in T \\ \sum_{s' \in S} P(s,a)(s') \cdot y(s') & \text{otherwise} \end{cases}$$

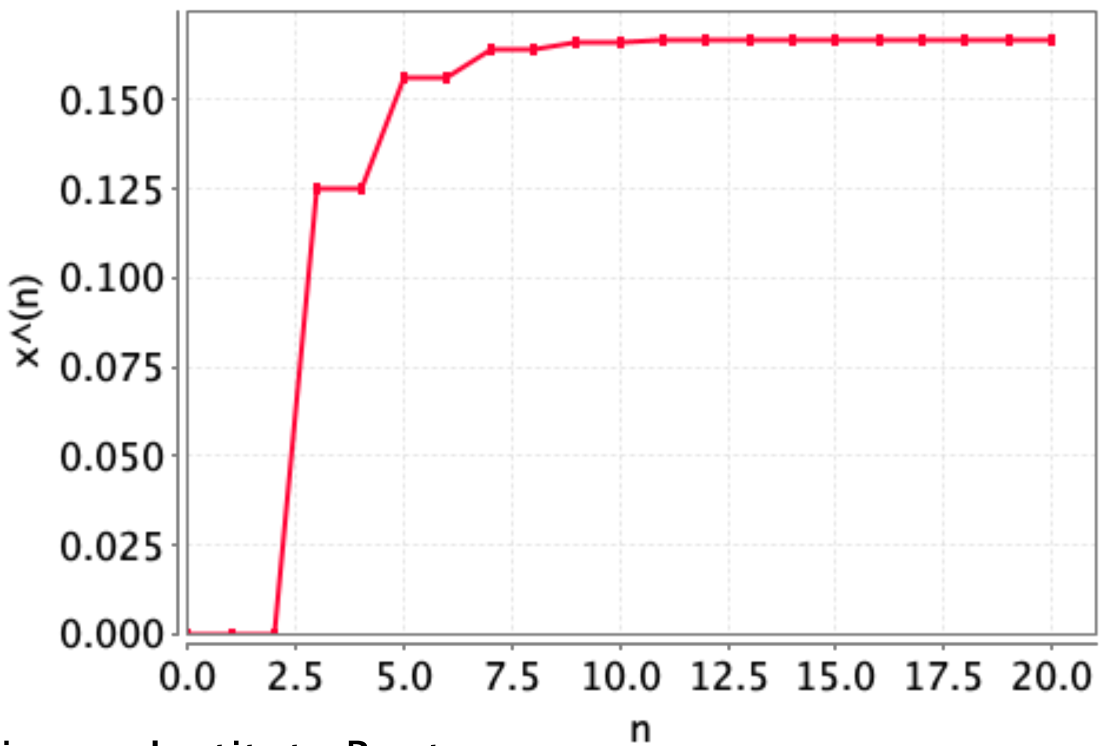

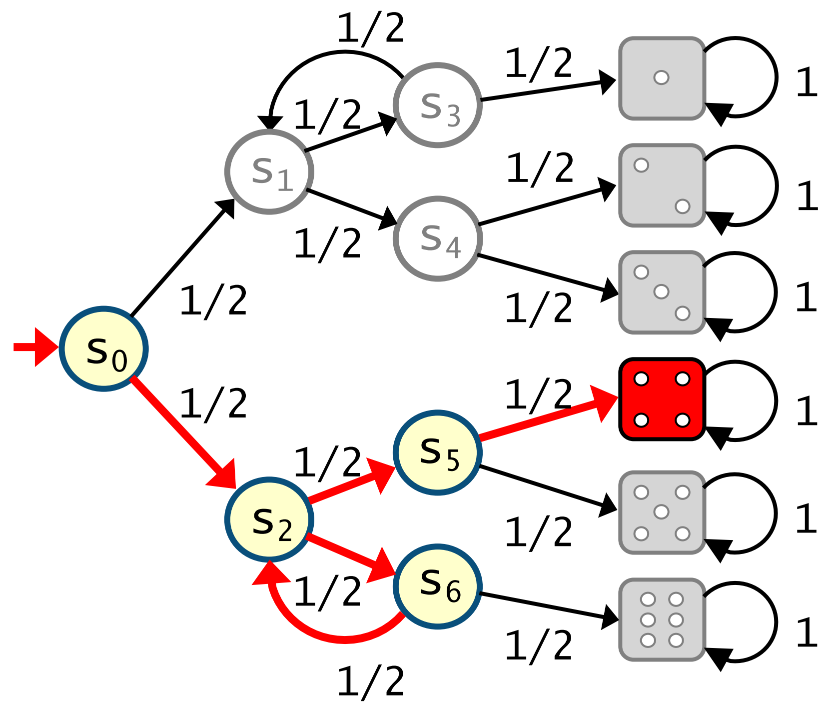

Probabilistic reachability – Complexity

Computing reachability probabilities for DTMCs reduces to

- graph-based analysis (finding the states that can reach the target)
- solving a linear equation system

Graph based analysis

- linear in the size of the DTMC (simple backwards traversal to find the states that can reach the target)

Solving a system of linear equations

- polynomial (cubic) in the size of the DTMC (Gaussian elimination)

In practice iterative methods are used for solving large linear equation systems

- power method (i.e. as a least fixed point), Gauss-Siedel

More general probabilistic properties

For example can compute the **probability** an LTL formula **ψ** is true

– $P_{=?}$[ψ] "what is the probability that ψ holds?"

– need to compute **Prob** { ω ∈ Path | ω satisfies ψ }

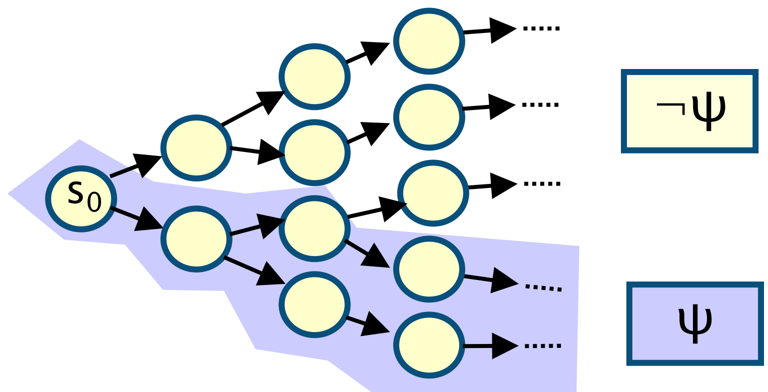

Such sets of path are measurable (elements of the event set)

– therefore probability is well defined

LTL model checking for DTMCs

Model check LTL specification ψ against a DTMC

1. Generate a deterministic Rabin automaton (DRA) for ψ
   - build nondeterministic Büchi automaton (NBA) for ψ [VW94]
   - convert the NBA to a DRA [Saf88]
2. Construct product DTMC **D⊗A**
3. Identify accepting BSCCs of **D⊗A**
   - BSCC: bottom strongly connected components
     - these are sets of states such that any state can be reached from any other state and once entered one cannot leave the set
4. Compute probability of reaching accepting BSCCs
   - from all states of the D⊗A

Discrete time Markov chains

Discrete-time Markov chains (DTMCs)

Paths and probabilities for DTMCs

Probabilistic reachability for DTMCs

Rewards and expected reachability for DTMCs

Reward structures

We augment DTMCs with rewards (or, conversely, costs)

- real-valued quantities assigned to states and/or transitions
- these can have a wide range of possible interpretations

Some examples:

- elapsed time, power consumption, size of message queue, number of messages successfully delivered, net profit, …

Costs? or rewards?

- mathematically, no distinction between rewards and costs
- when interpreted normally desirable to minimise costs and maximise rewards

DTMC reward structures

For a DTMC a reward structure is a pair $(\underline{r}, \mathbf{R})$

- $\underline{r} : S \rightarrow \mathbb{R}_{\geq 0}$ is the state reward function (vector over states)
- $\mathbf{R} : S \times S \rightarrow \mathbb{R}_{\geq 0}$ is the transition reward function (matrix over states)

> $\underline{r}(s)$ – the reward associated with state s
>
> $\mathbf{R}(s,s')$ – the reward associated with the transition from s to s'

DTMC reward structures

For a DTMC a reward structure is a pair **(r,R)**

- $r : S \rightarrow \mathbb{R}_{\geq 0}$ is the state reward function (vector over states)
- $\mathbf{R} : S \times S \rightarrow \mathbb{R}_{\geq 0}$ is the transition reward function (matrix over states)

Examples

- "time-steps": r is not used and R returns 1 for all transitions
- "number of messages lost": r is not used and R maps transitions corresponding to a message loss to 1
- "power consumption": r is defined as the per-time-step energy consumption in each state and R as the energy cost of each transition

Probability basics – Expectations

Recall a probability space is a tuple $(\Omega, \Sigma, \text{Prob})$

- Ω is the sample space
- Σ is the event set
- $\text{Prob}: \Sigma \to [0,1]$ is the probability measure

Real valued random variable X over the probability space is a (measurable) function $X: \Omega \to \mathbb{R}$

maps elements of the sample space to real values

Probability basics – Expectations

Recall a probability space is a tuple **$(\Omega, \Sigma, \mathrm{Prob})$**

- Ω is the sample space
- Σ is the event set
- $\mathrm{Prob}: \Sigma \to [0,1]$ is the probability measure

Real valued **random variable X** over the probability space is a (measurable) function **$X: \Omega \to \mathbb{R}$**

Expected ("average") value of the random variable:

$$\mathrm{Exp}(X) = \int_{\omega \in \Omega} X(\omega)\ d\mathrm{Prob}$$

measurability needed for integral to be well-defined

Probability basics – Expectations

Recall a probability space is a tuple $(\Omega, \Sigma, Prob)$

- Ω is the sample space
- Σ is the event set
- $Prob: \Sigma \to [0,1]$ is the probability measure

Real valued **random variable X** over the probability space is a (measurable) function $X: \Omega \to \mathbb{R}$

Expected ("average") value of the random variable:

$$Exp(X) = \int_{\omega \in \Omega} X(\omega) \, dProb$$
$$= \Sigma_{\omega \in \Omega} X(\omega) \cdot Prob(\omega)$$

if the probability space is discrete, e.g. finite

Probability basics – Expectations

Example: expected (average) value of a die when tossed

Probability space

- sample space: possible values {“1”,“2”,“3”,“4”,“5”,“6”}
- events: subsets of the sample space
- probability measure: **Prob**(“1”) = **Prob**(“2”) = ⋯ = **Pr**(“6”) = 1/6

Random variable **X:Ω→ℝ**

- the value of the die: X(“1”) = 1, X(“2”) = 2,… , X(“6”) = 6

Expected value of the random variable

- i.e. the expected (average) value of the die when thrown
- E(X) = **Prob**(“1”)·X(“1”)+**Prob**(“2”)·X(“2”)+ ⋯ +**Prob**(“6”)·X(“2”)

  = 1/6·1 + 1/6·2 + ⋯ + 1/6·6 = 21/6 = 3⅓

Expected reward properties

Probability space for DTMCs

- sample space is the set infinite paths Path
- therefore random variables of the form $X\colon \mathit{Path} \to \mathbb{R}$

function from infinite paths to real values

Consider any infinite path $\omega = s_0 s_1 s_2 \ldots$

Cumulative (reachability)

- "reward cumulated before reaching a target set T"
- random variable X where

$X(\omega)$ equals $r(s_0) + \cdots + r(s_{k-1}) + \mathbf{R}(s_0, s_1) + \cdots + \mathbf{R}(s_{k-1}, s_k)$

if $k = \min\{\, j \mid s_j \in T \,\}$ exists

find the first time that a state in T is reached along the path

Expected reward properties

Probability space for DTMCs

- sample space is the set infinite paths Path
- therefore random variables of the form $X\colon Path \to \mathbb{R}$

function from infinite paths to real values

Consider any infinite path $\omega = s_0 s_1 s_2 \ldots$

Cumulative (reachability)

- "reward cumulated before reaching a target set T"
- random variable X where

$X(\omega)$ equals $r(s_0) + \cdots + r(s_{k-1}) + \mathbf{R}(s_0,s_1) + \cdots + \mathbf{R}(s_{k-1},s_k)$

if $k = \min\{ j \mid s_j \in T \}$ exists

summation of rewards up until T is reached for the first time

Expected reward properties

Probability space for DTMCs

- sample space is the set infinite paths Path
- therefore random variables of the form $X:\text{Path}\rightarrow\mathbb{R}$

function from infinite paths to real values

Consider any infinite path $\omega = s_0 s_1 s_2 \ldots$

Cumulative (reachability)

- "reward cumulated before reaching a target set T"
- random variable X where

$X(\omega)$ equals $\underline{r}(s_0) + \cdots + \underline{r}(s_{k-1}) + \mathbf{R}(s_0,s_1) + \cdots + \mathbf{R}(s_{k-1},s_k)$

if $k = \min\{ j \mid s_j \in T \}$ exists and infinity otherwise

summation of rewards up until T is reached for the first time
however, if T is never reached the cumulated reward is infinity

Computing the expected rewards

Computing expected cumulated reward before reaching a state in T

- graph-based analysis and solving a system of linear equations
- compute the expectations for all states

Let y_s denote the value of E(X) when starting from state s

Computing the expected rewards

Computing expected cumulated reward before reaching a state in T

- graph-based analysis and solving a system of linear equations
- compute the expectations for all states

Let y_s denote the value of E(X) when starting from state s

- if s is in T, then $y_s = 0$

we have reached a state in T so no rewards to cumulate

Computing the expected rewards

Computing expected cumulated reward before reaching a state in T

- graph-based analysis and solving a system of linear equations
- compute the expectations for all states

Let y_s denote the value of E(X) when starting from state s

- if s is in T, then $y_s = 0$
- if s does not reach T with probability 1, then $y_s = \infty$

follows from the fact that if no state in T is reached
we set the cumulated reward to infinity for the path

Computing the expected rewards

Computing expected cumulated reward before reaching a state in T

- graph-based analysis and solving a system of linear equations
- compute the expectations for all states

Let y_s denote the value of E(X) when starting from state s

- if s is in T, then $y_s = 0$
- if s does not reach a state in T with probability 1, then $y_s = \infty$
- otherwise $y_s = r(s) + \sum_{s' \in S} P(s,s') \cdot R(s,s') + \sum_{s' \in S} P(s,s') \cdot y_{s'}$

state reward associated with s

average (expected) transition reward when leaving state s

average of the expectations in successor states

Expected rewards – Example

Expected number of coin flips

if s is in T, then $y_s = 0$
if s does not reach T with probability 1, then $y_s = \infty$
otherwise $y_s = r(s) + \sum_{s' \in S} \mathbf{P}(s,s') \cdot \mathbf{R}(s,s') + \sum_{s' \in S} \mathbf{P}(s,s') \cdot y_{s'}$

– reward structure $\mathbf{R}(s,s')=1$ for all states labelled done

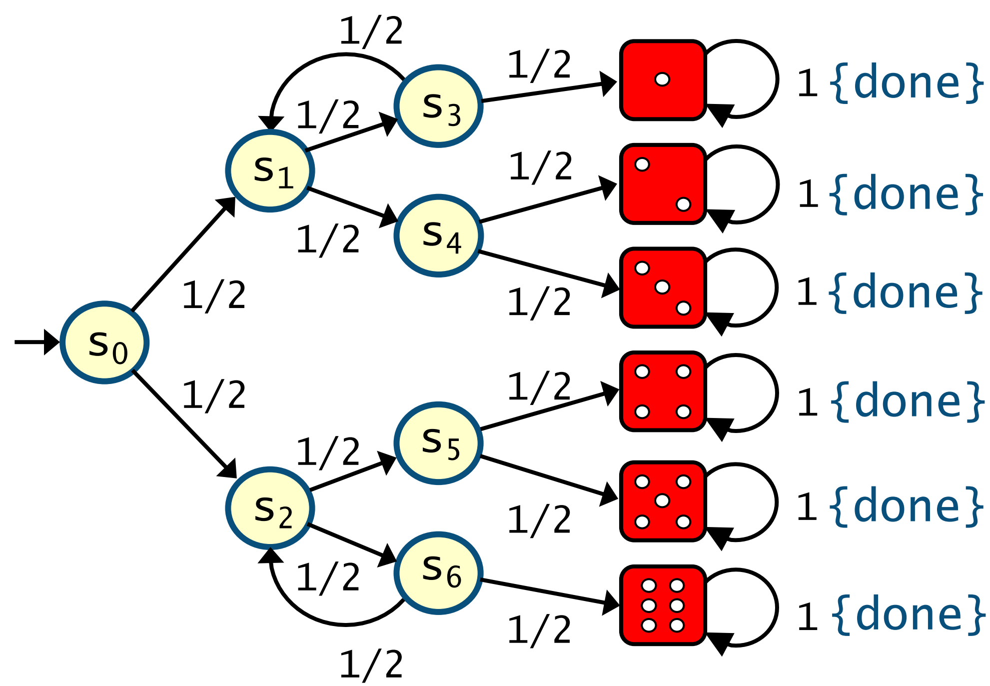

Expected rewards – Example

Expected number of coin flips

if s is in T, then $y_s = 0$
if s does not reach T with probability 1, then $y_s = \infty$
otherwise $y_s = r(s) + \Sigma_{s' \in S} \mathbf{P}(s,s') \cdot \mathbf{R}(s,s') + \Sigma_{s' \in S} \mathbf{P}(s,s') \cdot y_{s'}$

– reward structure $\mathbf{R}(s,s')=1$ for all states labelled done

- $y_s = 0$ for all states labelled done

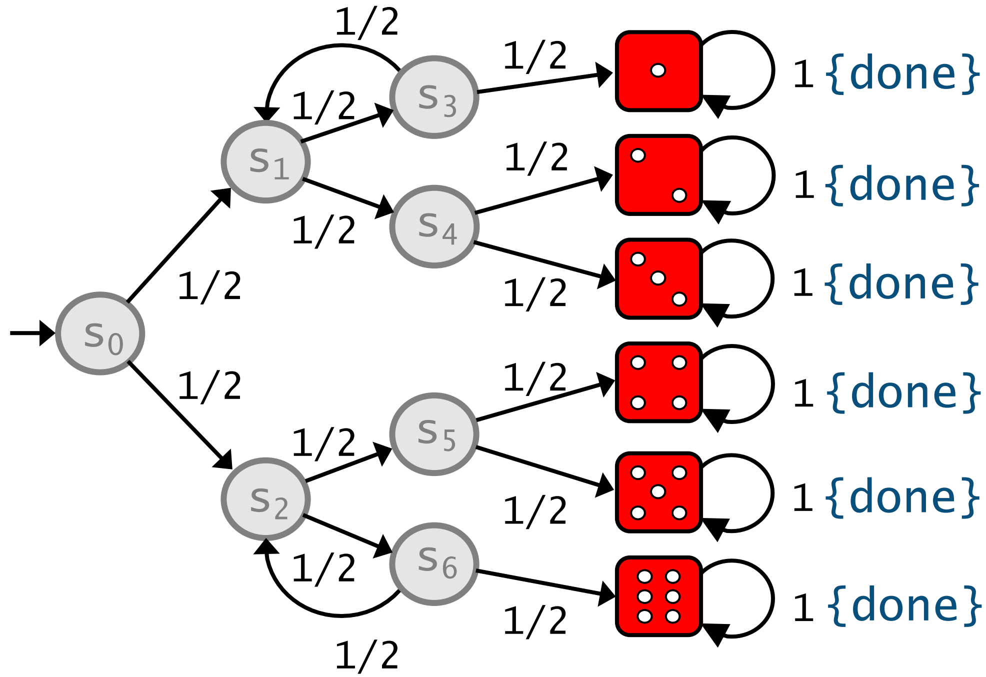

Expected rewards – Example

Expected number of coin flips

> if s is in T, then $y_s = 0$
> if s does not reach T with probability 1, then $y_s = \infty$
> otherwise $y_s = r(s) + \Sigma_{s' \in S} \mathbf{P}(s,s') \cdot \mathbf{R}(s,s') + \Sigma_{s' \in S} \mathbf{P}(s,s') \cdot y_{s'}$

– reward structure $\mathbf{R}(s,s')=1$ for all states labelled done

- $y_s = 0$ for all states labelled done
- all states reach done with probability 1 therefore no state has value infinity

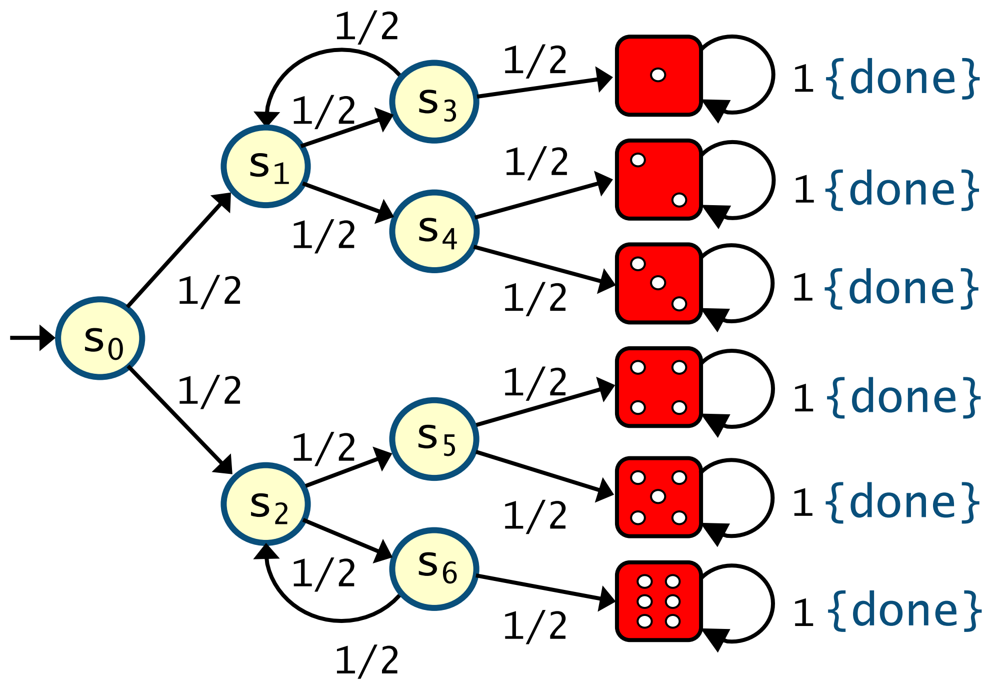

Expected rewards – Example

Expected number of coin flips

if s is in T, then $y_s = 0$

if s does not reach T with probability 1, then $y_s = \infty$

otherwise $y_s = r(s) + \Sigma_{s' \in S} \mathbf{P}(s,s') \cdot \mathbf{R}(s,s') + \Sigma_{s' \in S} \mathbf{P}(s,s') \cdot y_{s'}$

– reward structure $\mathbf{R}(s,s')=1$ for all states labelled done

- $y_s = 0$ for all states labelled done
- $y_{s0} = 1/2 \cdot 1 + 1/2 \cdot 1 + 1/2 \cdot y_{s1} + 1/2 \cdot y_{s2}$

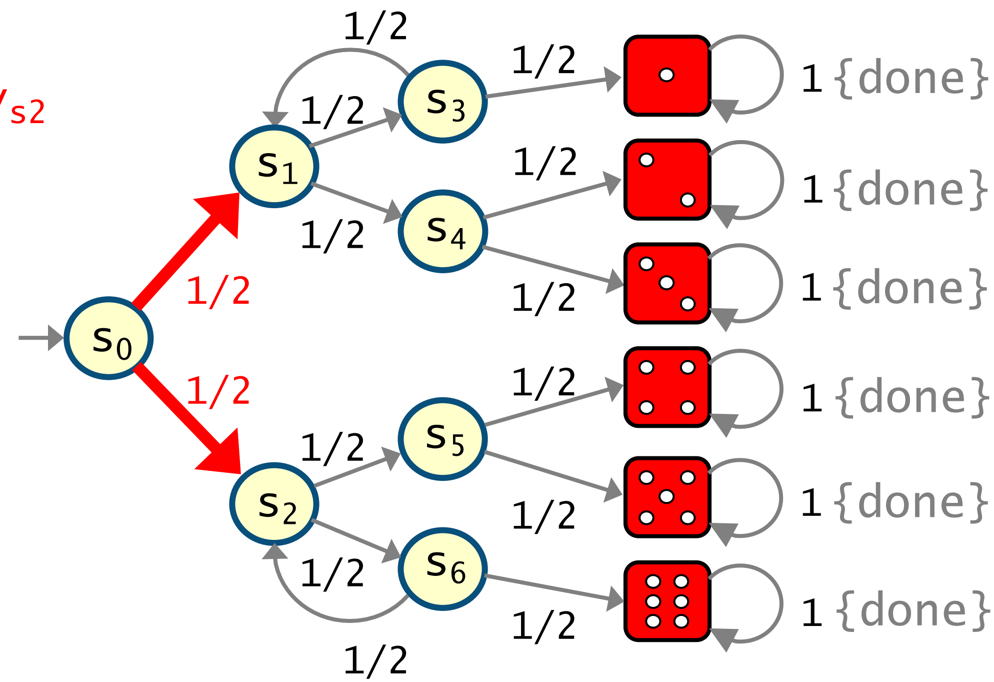

Expected rewards – Example

Expected number of coin flips

> if s is in T, then $y_s = 0$
> if s does not reach T with probability 1, then $y_s = \infty$
> otherwise $y_s = r(s) + \Sigma_{s' \in S} P(s,s') \cdot R(s,s') + \Sigma_{s' \in S} P(s,s') \cdot y_{s'}$

– reward structure $R(s,s')=1$ for all states labelled done

- $y_s = 0$ for all states labelled done
- $y_{s0} = 1 + 1/2 \cdot y_{s1} + 1/2 \cdot y_{s2}$

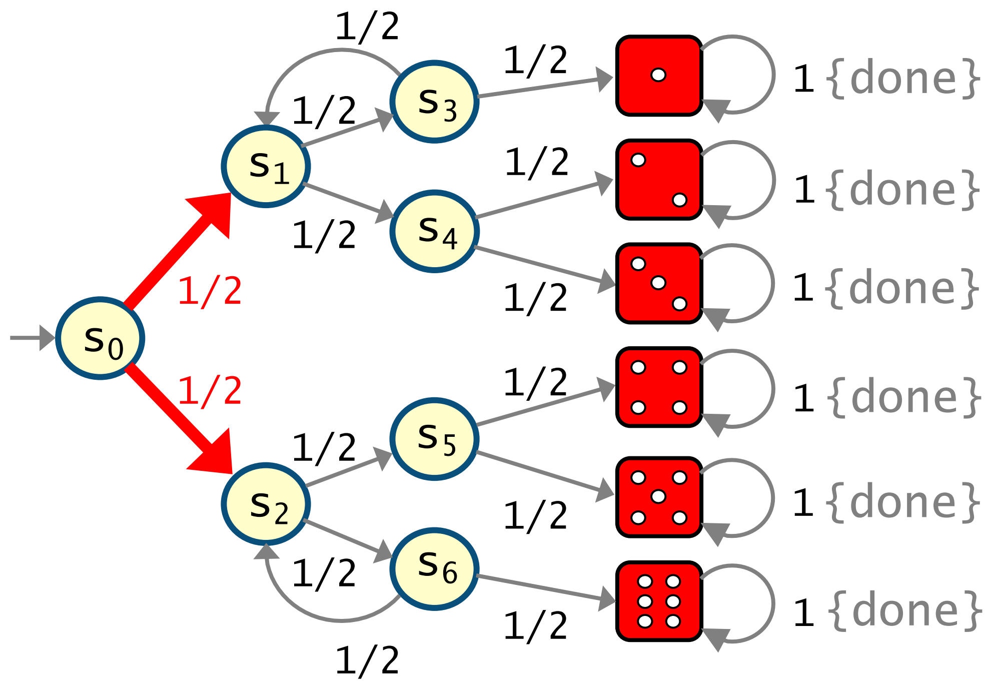

Expected rewards – Example

Expected number of coin flips

> if s is in T, then $y_s = 0$
> if s does not reach T with probability 1, then $y_s = \infty$
> otherwise $y_s = r(s) + \Sigma_{s' \in S} \mathbf{P}(s,s') \cdot \mathbf{R}(s,s') + \Sigma_{s' \in S} \mathbf{P}(s,s') \cdot y_{s'}$

– reward structure $\mathbf{R}(s,s')=1$ for all states labelled done

- $y_s = 0$ for all states labelled done
- $y_{s0} = 1 + 1/2 \cdot y_{s1} + 1/2 \cdot y_{s2}$
- $y_{s1} = 1 + 1/2 \cdot y_{s3} + 1/2 \cdot y_{s4}$
- $y_{s2} = 1 + 1/2 \cdot y_{s5} + 1/2 \cdot y_{s6}$
- $y_{s3} = 1 + 1/2 \cdot y_{s1} + 1/2 \cdot 0$
- $y_{s4} = 1 + 1/2 \cdot 0 + 1/2 \cdot 0$
- $y_{s5} = 1 + 1/2 \cdot 0 + 1/2 \cdot 0$
- $y_{s6} = 1/2 + 1/2 + 1/2 \cdot 0 + 1/2 \cdot y_{s2}$

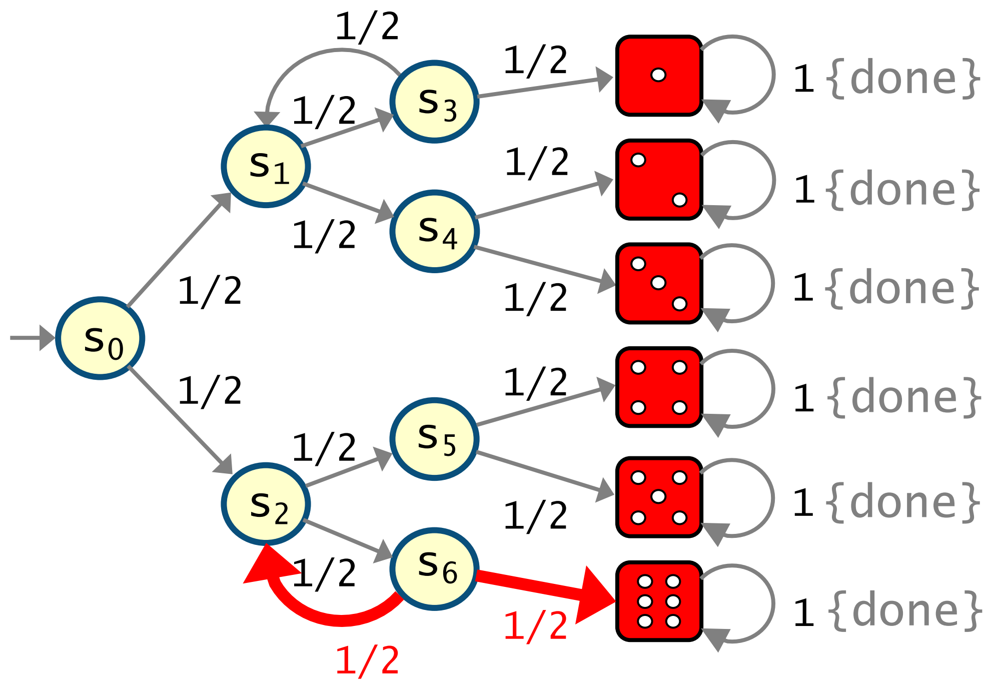

Expected rewards – Example

Expected number of coin flips

> if s is in T, then $y_s = 0$
> if s does not reach T with probability 1, then $y_s = \infty$
> otherwise $y_s = r(s) + \Sigma_{s' \in S} \mathbf{P}(s,s') \cdot \mathbf{R}(s,s') + \Sigma_{s' \in S} \mathbf{P}(s,s') \cdot y_{s'}$

– reward structure $\mathbf{R}(s,s')=1$ for all states labelled done

- $y_s = 0$ for all states labelled done
- $y_{s0} = 1 + 1/2 \cdot y_{s1} + 1/2 \cdot y_{s2}$
- $y_{s1} = 1 + 1/2 \cdot y_{s3} + 1/2 \cdot y_{s4}$
- $y_{s2} = 1 + 1/2 \cdot y_{s5} + 1/2 \cdot y_{s6}$
- $y_{s3} = 1 + 1/2 \cdot y_{s1}$
- $y_{s4} = 1$
- $y_{s5} = 1$
- $y_{s6} = 1 + 1/2 \cdot y_{s2}$

simplifying

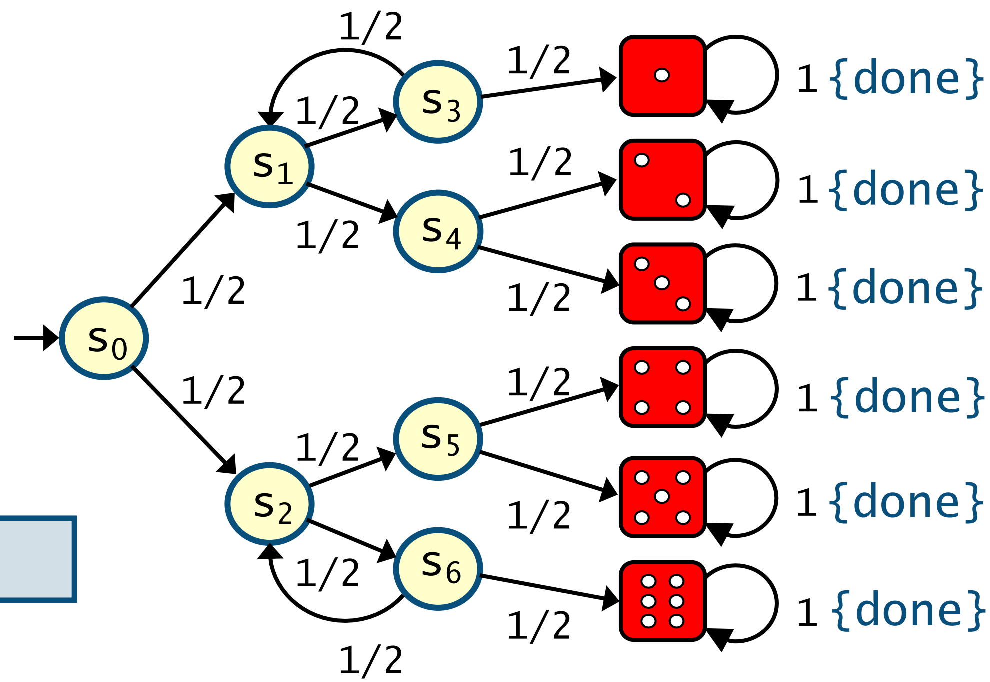

Expected rewards – Example

Expected number of coin flips

if s is in T, then $y_s = 0$
if s does not reach T with probability 1, then $y_s = \infty$
otherwise $y_s = r(s) + \Sigma_{s'\in S}\mathbf{P}(s,s')\cdot\mathbf{R}(s,s') + \Sigma_{s'\in S}\mathbf{P}(s,s')\cdot y_{s'}$

– reward structure $\mathbf{R}(s,s')=1$ for all states labelled done

- $y_s = 0$ for all states labelled done
- $y_{s0} = 1 + 1/2\cdot y_{s1} + 1/2\cdot y_{s2}$
- $y_{s1} = 1 + 1/2\cdot(1+1/2\cdot y_{s1}) + 1/2$
- $y_{s2} = 1 + 1/2 + 1/2\cdot(1+1/2\cdot y_{s2})$
- $y_{s3} = 1 + 1/2\cdot y_{s1}$
- $y_{s4} = 1$
- $y_{s5} = 1$
- $y_{s6} = 1 + 1/2\cdot y_{s2}$

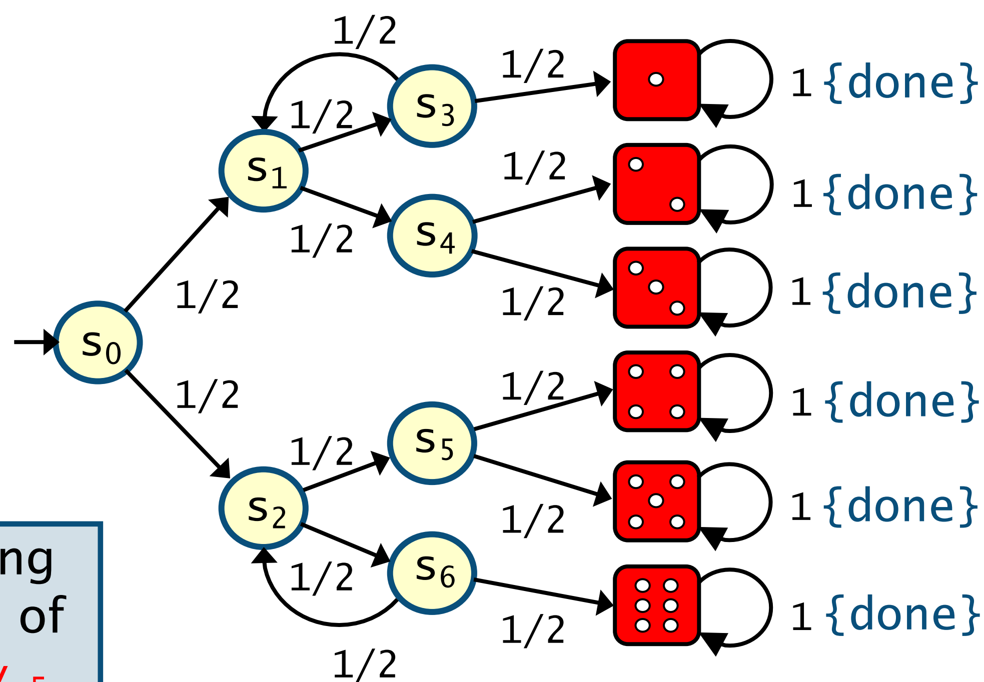

substituting the values of y_{s3}, y_{s4}, y_{s5} and y_{s6}

Expected rewards – Example

Expected number of coin flips

> if s is in T, then $y_s = 0$
> if s does not reach T with probability 1, then $y_s = \infty$
> otherwise $y_s = r(s) + \Sigma_{s' \in S} \mathbf{P}(s,s') \cdot \mathbf{R}(s,s') + \Sigma_{s' \in S} \mathbf{P}(s,s') \cdot y_{s'}$

– reward structure $\mathbf{R}(s,s')=1$ for all states labelled done

- $y_s = 0$ for all states labelled done
- $y_{s0} = 1 + 1/2 \cdot y_{s1} + 1/2 \cdot y_{s2}$
- $y_{s1} = 8/3$
- $y_{s2} = 8/3$
- $y_{s3} = 1 + 1/2 \cdot y_{s1}$
- $y_{s4} = 1$
- $y_{s5} = 1$
- $y_{s6} = 1 + 1/2 \cdot y_{s2}$

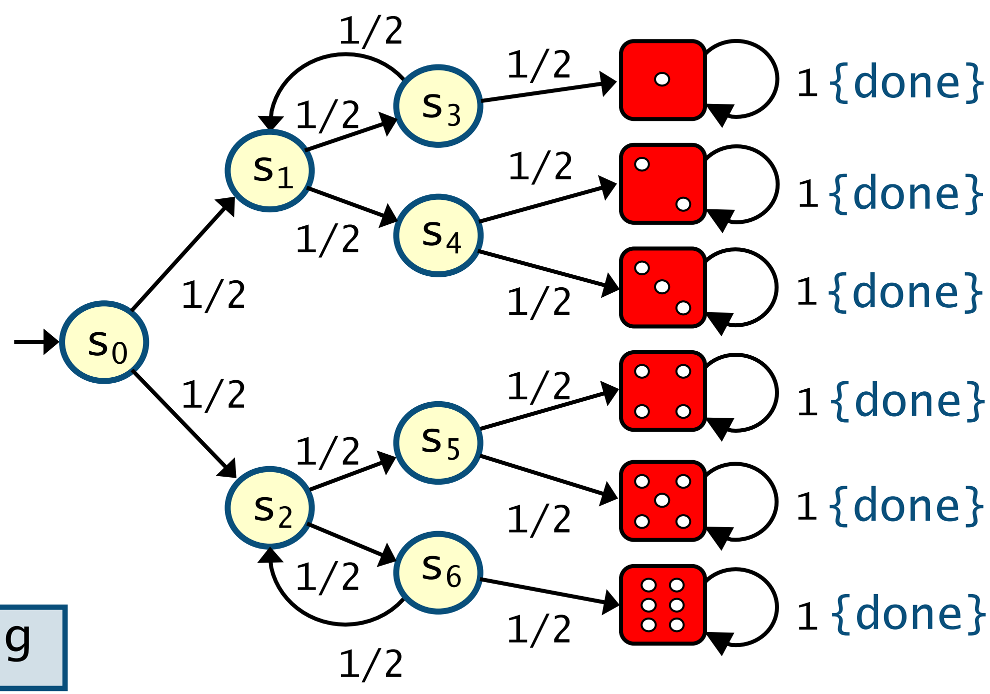

simplifying

Expected rewards – Example

Expected number of coin flips

> if s is in T, then $y_s = 0$
> if s does not reach T with probability 1, then $y_s = \infty$
> otherwise $y_s = r(s) + \Sigma_{s' \in S} \mathbf{P}(s,s') \cdot \mathbf{R}(s,s') + \Sigma_{s' \in S} \mathbf{P}(s,s') \cdot y_{s'}$

– reward structure $\mathbf{R}(s,s')=1$ for all states labelled done

- $y_s = 0$ for all states labelled done
- $y_{s0} = 1 + 1/2 \cdot (8/3) + 1/2 \cdot (8/3)$
- $y_{s1} = 8/3$
- $y_{s2} = 8/3$
- $y_{s3} = 1 + 1/2 \cdot (8/3)$
- $y_{s4} = 1$
- $y_{s5} = 1$
- $y_{s6} = 1 + 1/2 \cdot (8/3)$

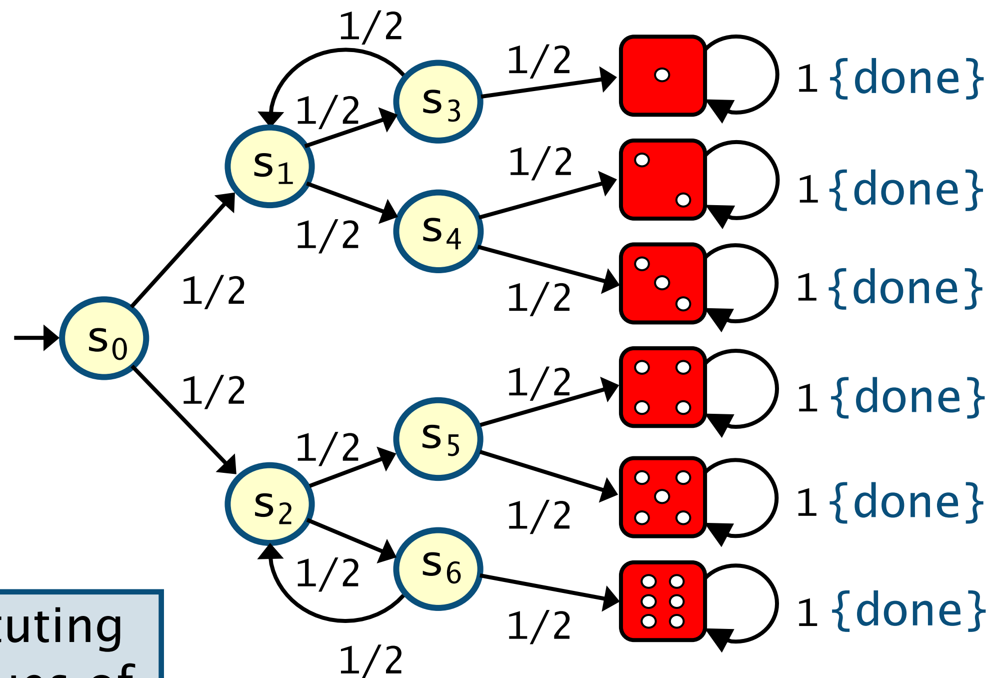

substituting the values of y_{s1} and y_{s2}

Expected rewards – Example

Expected number of coin flips

> if s is in T, then $y_s = 0$
> if s does not reach T with probability 1, then $y_s = \infty$
> otherwise $y_s = r(s) + \Sigma_{s' \in S} \mathbf{P}(s,s') \cdot \mathbf{R}(s,s') + \Sigma_{s' \in S} \mathbf{P}(s,s') \cdot y_{s'}$

– reward structure $\mathbf{R}(s,s')=1$ for all states labelled done

- $y_s = 0$ for all states labelled done
- $y_{s0} = 11/3$
- $y_{s1} = 8/3$
- $y_{s2} = 8/3$
- $y_{s3} = 7/3$
- $y_{s4} = 1$
- $y_{s5} = 1$
- $y_{s6} = 7/3$

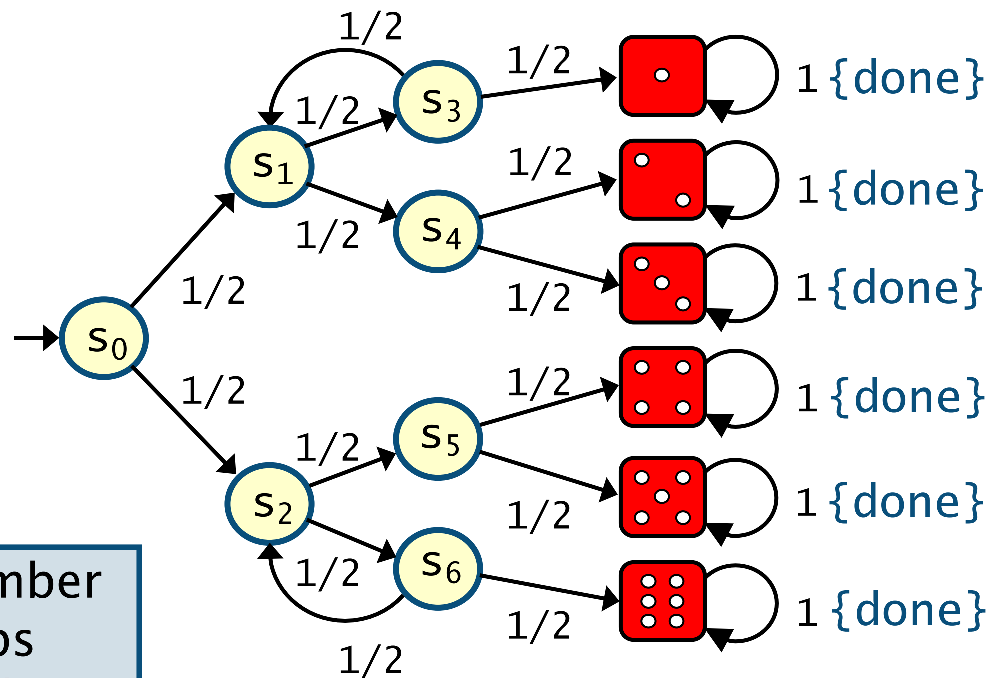

expected number of coin flips equals 11/3

Expected reachability – Complexity

Computing expected reachability values for DTMCs reduces to

- graph-based analysis (find states that reach the target with probability 1)
- solving a linear equation system

Graph based analysis

- linear in the size of the DTMC (simple backwards traversal to find the states that can reach the target)

Solving a system of linear equations

- polynomial (cubic) in the size of the DTMC (Gaussian elimination)
- again in practice use iterative methods
- as for probabilistic reachability can express as a least fixed point

Additional Reward Properties

Instantaneous

- "the expected value of the state reward at time-step k"
- e.g. "the expected queue size after exactly 90 seconds"

Cumulative (time-bounded)

- "the expected reward cumulated up to time-step k"
- e.g. "the expected power consumption over one hour"

Also long run average and multi-objective properties

In the next video

Markov decision processes

- extend DTMCs to allow the modelling of non-deterministic behaviour