Formal Approaches to Decision-Making under Uncertainty

Lecture 1-2: Discrete-Time Markov Chains

Arnd Hartmanns
Formal Methods and Tools
UNIVERSITY OF TWENTE

Probability Theory Recap

A **discrete probability distribution** over a set of outcomes Ω

is a function $\mu\colon \Omega \to [0,1]$ such that

1. $\sum_{\omega \in \Omega} \mu(\omega) = 1$ and
2. $\mathrm{spt}(\mu) \overset{\mathrm{def}}{=} \{\omega \in \Omega \mid \mu(\omega) > 0\}$ is countable.

Ex: coin

$\Omega = \{h, t\}$

$\mu(h) = 0.5, \ \mu(t) = 0.5$

$\mathrm{Dist}(\Omega)$ is the set of all discrete probability distributions over Ω.

Example: 6-sided die

$\Omega = \{1, 2, \ldots, 6\}$

$\mu(i) = \frac{1}{6} \quad \forall i \in \Omega$

Geometric distr.

$\Omega = \mathbb{N}$

$\mu(0) = 0.5$

$\mu(1) = 0.25$

$\mu(2) = \frac{1}{8}$

...

Probability Theory Recap

Ω μ

A **probability space** is a triple $\langle \Omega, \mathfrak{F}, P \rangle$ consisting of

a sample space Ω containing all possible **outcomes**,

a σ-algebra $\mathfrak{F} \subseteq 2^{\Omega}$ of measurable **events**, and

a probability measure $P\colon \mathfrak{F} \to [0,1]$

where $P(\Omega) = 1$ and P is σ-additive, i.e.

for any pairwise disjoint sets $S_1, S_2, \ldots \in \mathfrak{F}$,

we have $P(\bigcup_i S_i) = \sum_i P(S_i)$.

$\Omega = [0,10]$

$\mathfrak{F} = \{ \emptyset, [0,10], [0,1], [0,0.1], [1,2], \ldots, (0,9), \ldots \}$

$P(\emptyset) = 0, \; P([0,10]) = 1, \; P([0,5]) = \frac{1}{2}, \; P([2,3)) = \frac{1}{10}$

$P([0,1] \cup [9,10]) = P([0,1]) + P([9,10]) = \frac{1}{10} + \frac{1}{10}$

Probability Theory Recap

A probability space is a triple $\langle \Omega, \mathfrak{F}, P \rangle$ …

Example: 6-sided die

$\Omega = \{1, \ldots, 6\}$

$\mathcal{F} = 2^{\Omega}$

$P(S) = \frac{1}{6}|S|$

$P(\{1,2,4\}) = \frac{1}{2}$

Example: Uniform distribution over $[1, 3]$

(see previous slide for $[0,10]$)

$P([a,b]) = \frac{b-a}{2}$

$P([1.5, 2.5]) = \frac{1}{2}$

Probability Theory Recap

Given a probability space $\langle \Omega, \mathfrak{F}, P \rangle$ and a measurable space $\langle S, \Sigma \rangle$, a **random variable** is a measurable function $X\colon \Omega \to S$.

If $S \subseteq \mathbb{R}$, then its expected value is $\mathbb{E}(X) \stackrel{\text{def}}{=} \int_{\omega} X(\omega)\, \mathrm{d}P(\omega)$

discrete: $\sum_{\omega} P(\{\omega\}) \cdot X(\omega)$

Example: Value of coin

$\Omega = \{h, t\}$ $\quad X(h) = 0$, $X(t) = 1$

Example: Distance of thrown ball

$\Omega = \{ \langle x, y \rangle \in [0, 10]^2 \}$

$X_{dist}(\langle x, y \rangle) = \sqrt{x^2 + y^2}$

A Markov chain from Córdoba airport to Rio Cuarto:

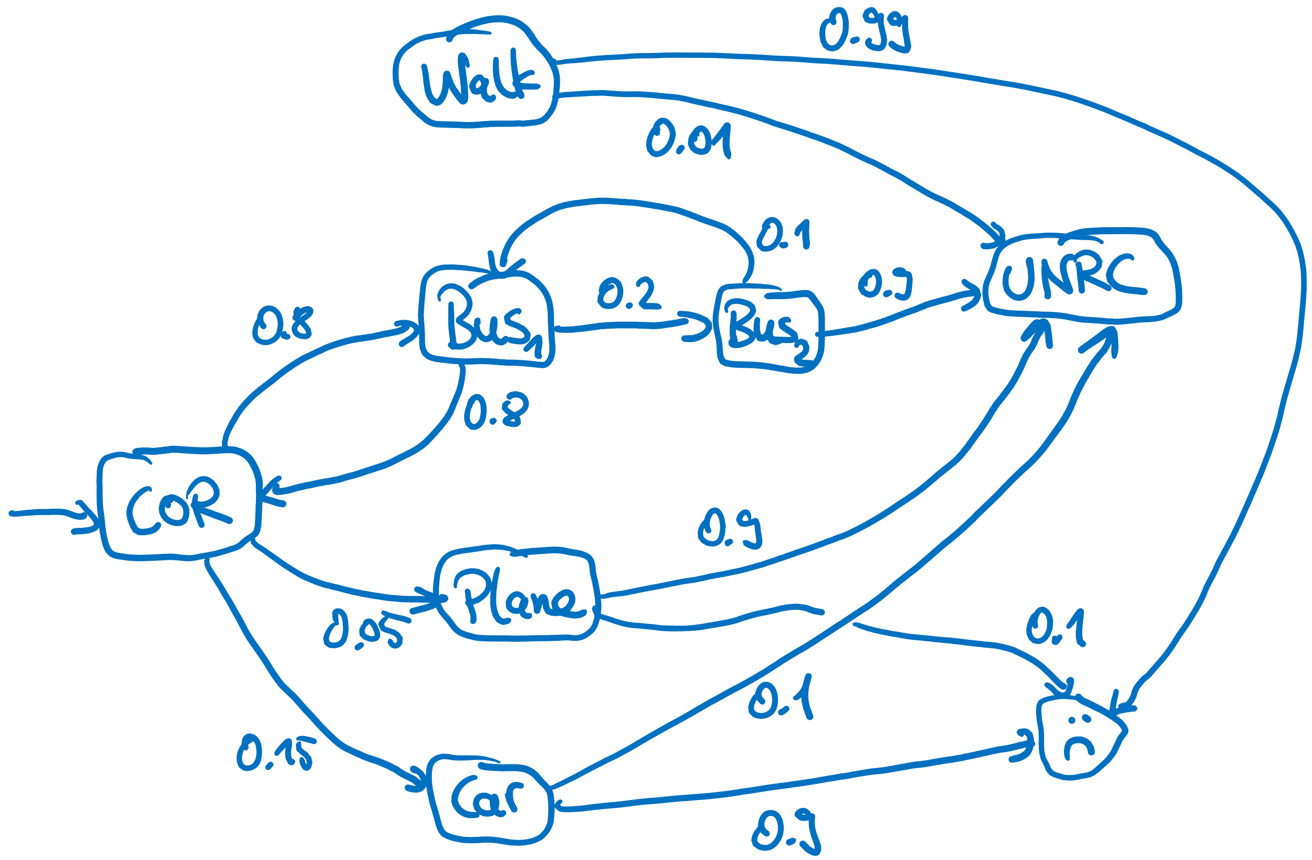

A Markov chain from Córdoba airport to Rio Cuarto:

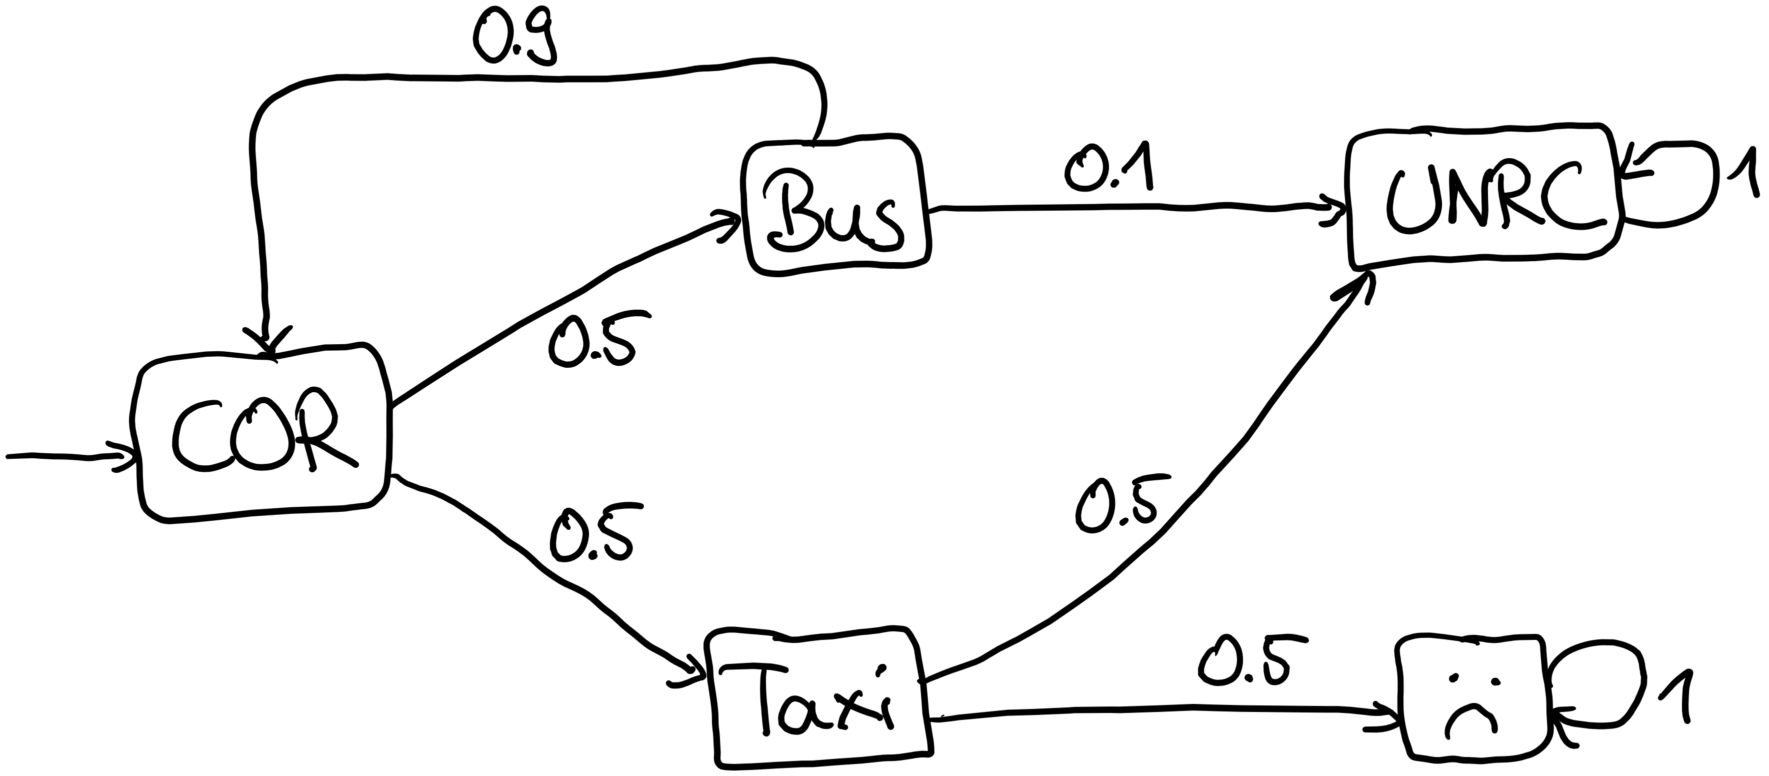

Markov Chains

DTMC

A **discrete-time Markov chain** M is a triple

$$\langle S, T, s_I \rangle$$

where S is the finite set of states,

$T\colon S \to \mathrm{Dist}(S)$ is the transition function, and

$s_I \in S$ is its initial state.

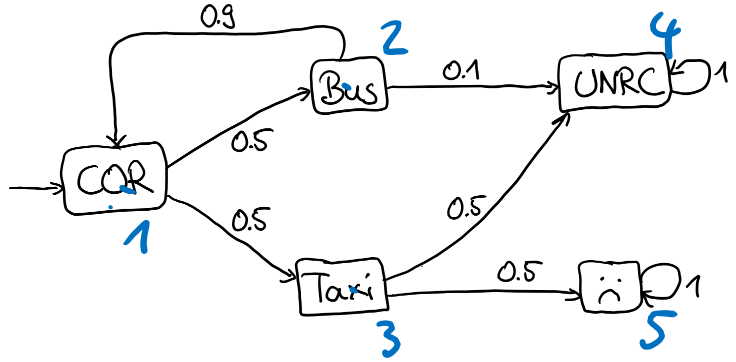

T as probability matrix:

$$\begin{array}{cc} & \begin{array}{ccccc} 1 & 2 & & & \end{array} \\ 1 & \begin{pmatrix} 0 & 0.5 & 0.5 & 0 & 0 \\ 0.9 & 0 & 0 & 0.1 & 0 \\ 0 & 0 & 0 & 0.5 & 0.5 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix} \end{array}$$

A path $\pi \in \mathrm{Paths}(M)$ is a sequence

$$s_0\, s_1 \ldots$$

of states such that $\forall i\colon T(s_i)(s_{i+1}) > 0$.

$\pi_{\leq n}$ is the prefix $s_0 \ldots s_n$ of π of length n.

Markov Chain Probability Space

Given a path prefix π_{fin},
the cylinder set $\mathrm{cyl}(\pi_{fin})$ is
the set of all paths starting with π_{fin}.

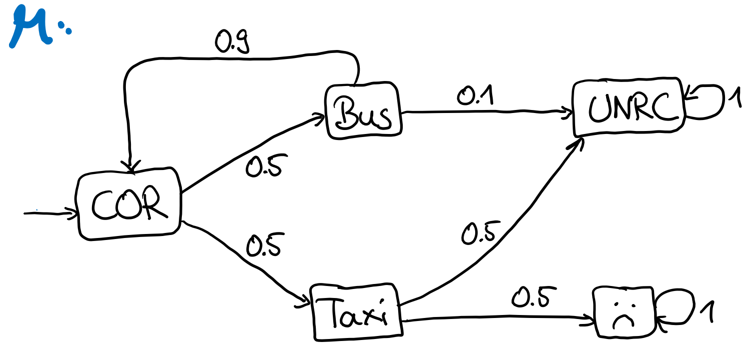

Examples:

cyl(COR Bus UNRC) = {COR Bus UNRC UNRC ...}

cyl((COR Taxi) = {COR Taxi UNRC ..., COR Taxi ☹ ☹ ...}

cyl((COR) = Paths(M)

Markov Chain Probability Space

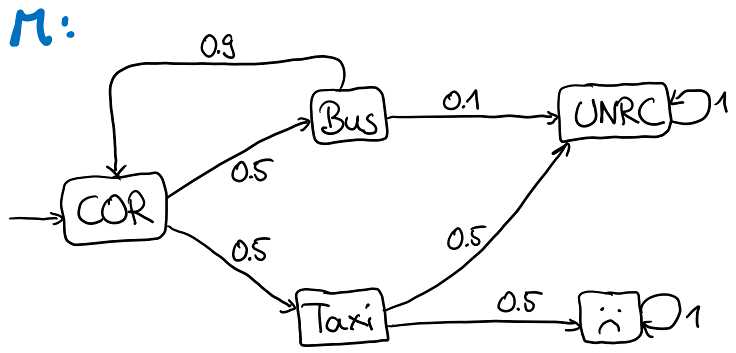

Let $\mathfrak{F}_{Cyl}$ be the smallest σ-algebra containing all cylinder sets of M.

Then **the probability space of** M is

$$\langle \text{Paths}(M), \mathfrak{F}_{Cyl}, P_M \rangle$$

$P_M(cyl(COR\ Bus)) = 0.5$

where P_M is defined by $P_M(\emptyset) = 0$ and

$$P_M(cyl(\underbrace{s_0\ s_1 \ldots s_n})) = \begin{cases} \prod_{i=0}^{n-1} T(\overset{COR}{s_i})(\overset{Bus}{s_{i+1}}) & \text{if } s_0 = s_I \wedge n > 0 \\ 1 & \text{if } s_0 = s_I \wedge n = 0 \\ 0 & \text{otherwise} \end{cases}$$

Reachability Probabilities

$G = \{UNRC\}$

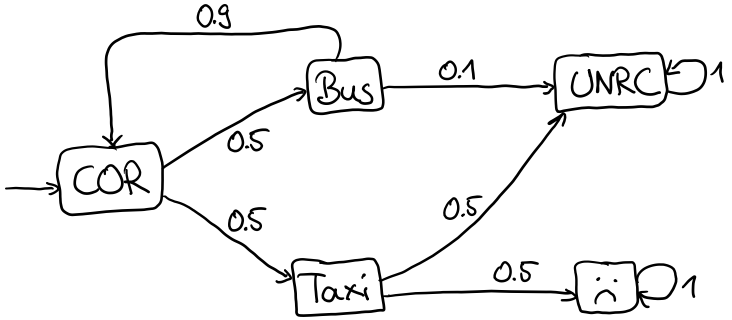

The **probability to reach** a goal state in $G \subseteq S$

$$P(\Diamond\, G) \qquad \text{or} \qquad P(\Diamond^{\leq b}\, G)$$

is now easily defined as

$$P(\Diamond\, G) = P_M(\{s_0\, s_1 \ldots \in \text{Paths}(M) \mid \exists i{:}\, s_i \in G\})$$

$P(\Diamond \{UNRC\}) = 0.5 \cdot 0.5 + 0.5 \cdot 0.1 + 0.5 \cdot 0.9 \cdot 0.5 \cdot 0.5 + \ldots$ $\quad P(\Diamond \{Bus\}) = 0.5$

and $P(\Diamond^{\leq b}\, G) = P_M(\{s_0\, s_1 \ldots \in \text{Paths}(M) \mid \exists i \leq b{:}\, s_i \in G\})$

Rewards and Costs

Buses and taxis aren't free.

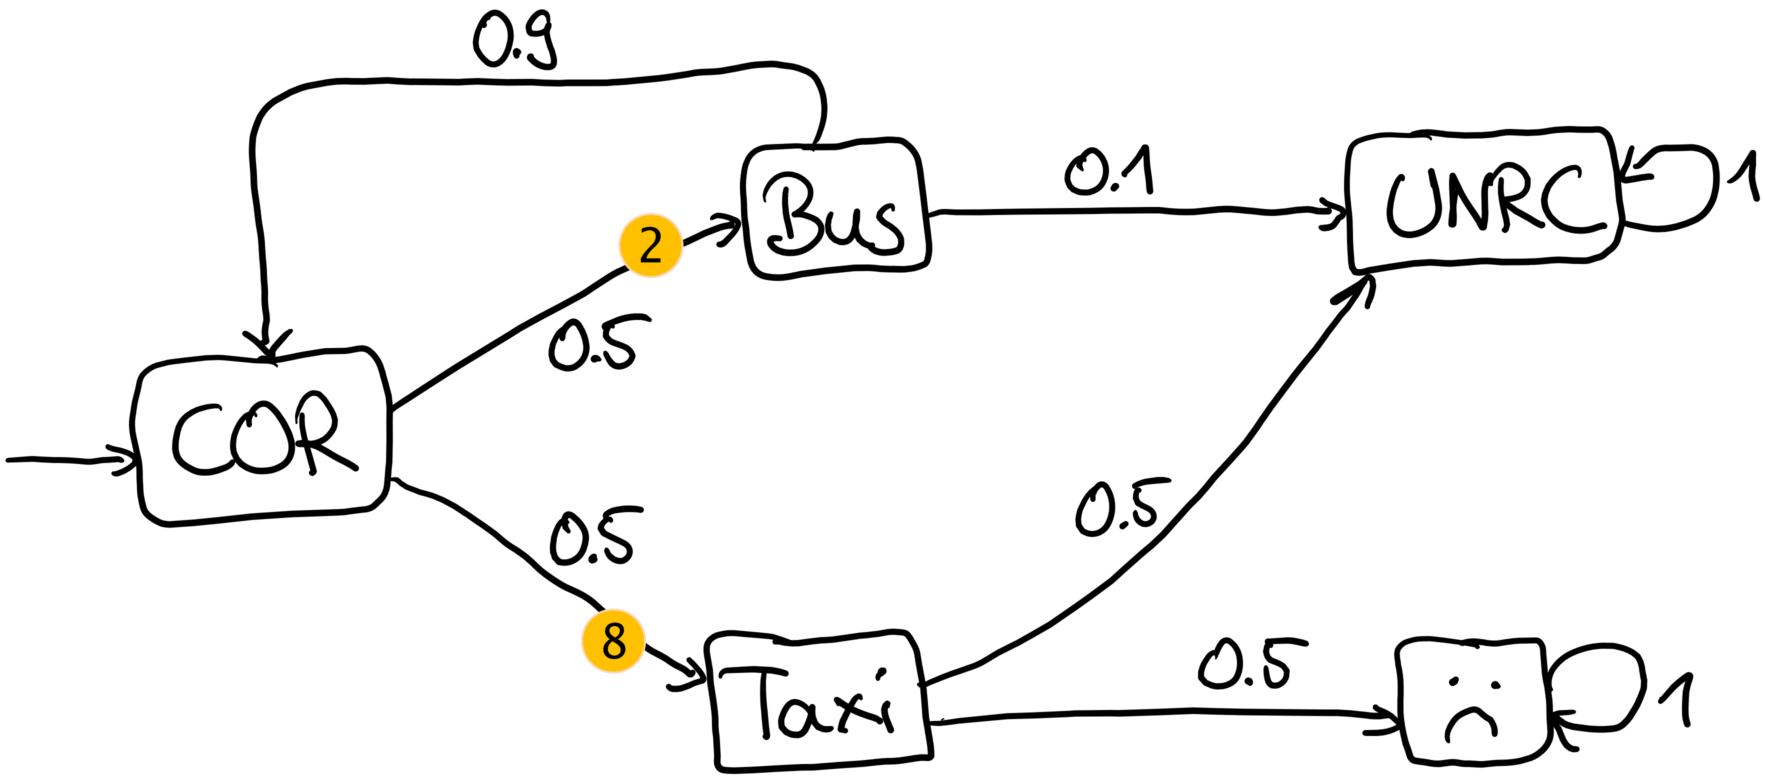

→ reward function $R\colon S \times S \to \mathbb{R}$

Rewards and Costs

The accumulated reward to reach a goal state in $G \subseteq S$ is

$$rew^{M}_{\diamond G}(s_0\ s_1 \ldots) \stackrel{\text{def}}{=} \begin{cases} \sum_{i=0}^{j} R(s_i, s_{i+1}) & \text{if } \exists j: s_j \in G \wedge \forall k < j: s_k \notin G \\ \infty & \text{otherwise} \end{cases}$$

$rew_{\diamond\{Bus\}}(COR\ Taxi\ ☹ \ldots) = \infty$

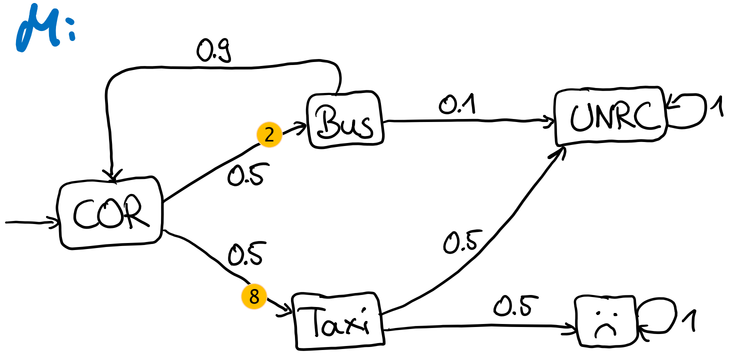

Example:

$rew_{\diamond\{UNRC\}}(COR\ Bus\ COR\ Bus\ UNRC \ldots) = 4$

→ *this is (almost) a random variable!*

Expected Rewards

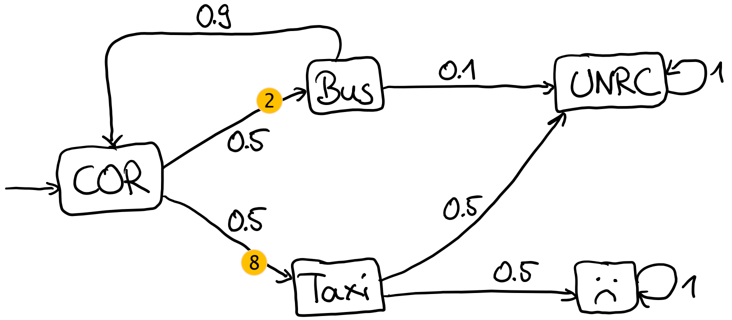

The **expected** accumulated **reward** to reach a goal state in $G \subseteq S$ is

$$R(\lozenge G) \stackrel{\text{def}}{=} \mathbb{E}\left(rew^{M}_{\lozenge G}\right)$$

Examples:

$R(\lozenge\{Bus, Taxi\}) = 0.5 \cdot 2 + 0.5 \cdot 8 = 5$

$R(\lozenge\{Bus\}) = \infty$

Markov Chain Modelling

Have: Coin that we can flip

Want: Play a game that requires 6-sided die

Claim: Can simulate die with coin as follows:

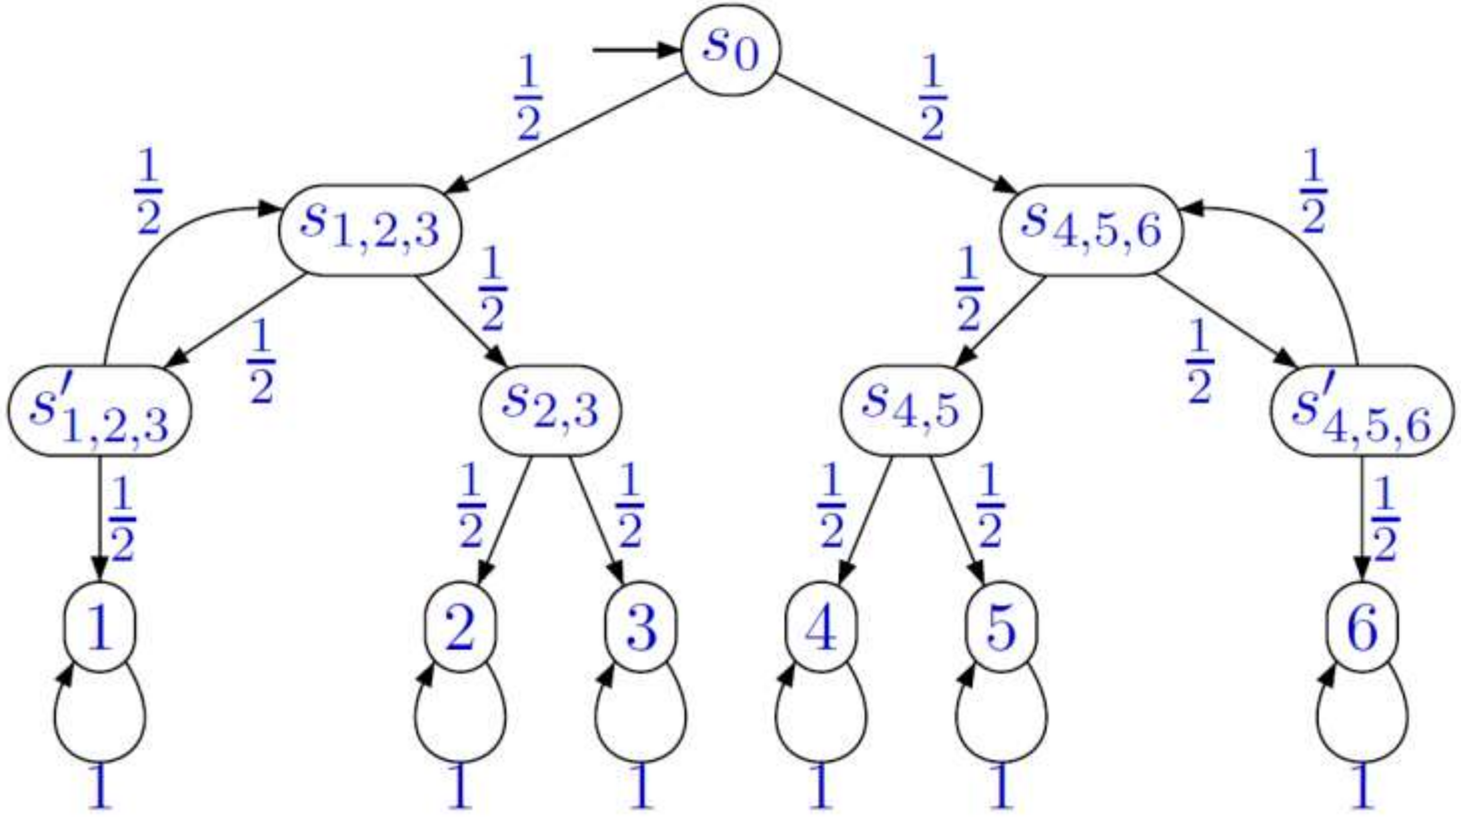

$P(\Diamond\{1\}) = \frac{1}{6}$

$\vdots$