EECS 144/244: Fundamental Algorithms for System Modeling, Analysis, and Optimization

Probabilistic Systems

Lecture: Markov Chains and Markov Decision Processes

Stavros Tripakis

University of California, Berkeley

Probabilistic systems

From non-deterministic systems:

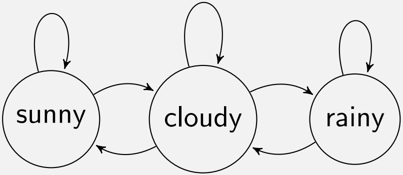

Probabilistic systems

To probabilistic systems:

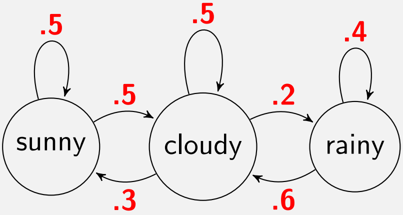

Markov chains

A finite Markov chain:

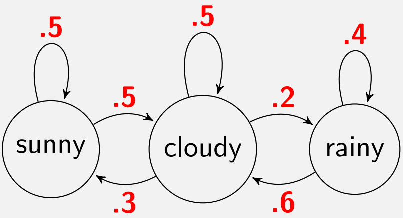

$$
\begin{aligned}
P(\text{next state is "sunny"} \mid \text{current state is "sunny"}) &= 0.5 \\
P(\text{next state is "sunny"} \mid \text{current state is "cloudy"}) &= 0.3 \\
P(\text{next state is "sunny"} \mid \text{current state is "rainy"}) &= 0 \\
\ldots
\end{aligned}
$$

Markov chains

A finite Markov chain:

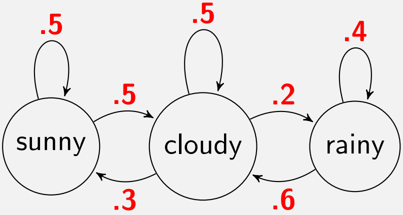

$$
\begin{aligned}
P(\text{next state is ``sunny'' } \mid \text{ current state is ``sunny''}) &= 0.5 \\
P(\text{next state is ``sunny'' } \mid \text{ current state is ``cloudy''}) &= 0.3 \\
P(\text{next state is ``sunny'' } \mid \text{ current state is ``rainy''}) &= 0 \\
\ldots
\end{aligned}
$$

The *Markov property*: only current state matters:

$$P(s_{k+1} = v_{k+1} \mid s_k = v_k, s_{k-1} = v_{k-1}, ..., s_0 = v_0) = P(s_{k+1} = v_{k+1} \mid s_k = v_k)$$

Markov chains

An infinite Markov chain:

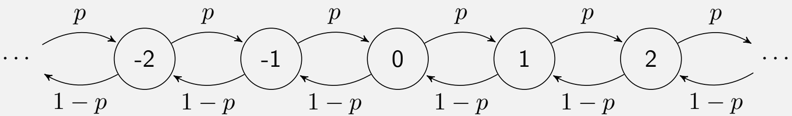

Markov chains

A finite Markov chain with n states can be represented by a square $n \times n$ *probability matrix* $\mathbf{P}$:

$$\mathbf{P} = \begin{bmatrix} p_{11} & p_{12} & \cdots & p_{1n} \\ p_{21} & p_{22} & \cdots & p_{2n} \\ \vdots & & \ddots & \vdots \\ p_{n1} & p_{n2} & \cdots & p_{nn} \end{bmatrix}$$

where

$$p_{ij} = P(\text{next state is } j \mid \text{current state is } i)$$

Markov chains

A finite Markov chain with n states can be represented by a square $n \times n$ *probability matrix* $\mathbf{P}$:

$$\mathbf{P} = \begin{bmatrix} p_{11} & p_{12} & \cdots & p_{1n} \\ p_{21} & p_{22} & \cdots & p_{2n} \\ \vdots & & \ddots & \vdots \\ p_{n1} & p_{n2} & \cdots & p_{nn} \end{bmatrix}$$

where

$$p_{ij} = P(\text{next state is } j \mid \text{current state is } i)$$

To be a valid probability matrix, $\mathbf{P}$ must satisfy:

$$\forall i, j : p_{ij} \geq 0 \qquad \text{and} \qquad \forall i : \sum_{j=1}^{n} p_{ij} = 1$$

Markov chains

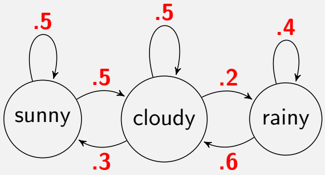

$$\mathbf{P} = \begin{bmatrix} 0.5 & 0.5 & 0 \\ 0.3 & 0.5 & 0.2 \\ 0 & 0.6 & 0.4 \end{bmatrix}$$

Transforming a process into a Markov chain

Homework:
Suppose that whether or not it rains today depends on the previous weather conditions during the last two days. Specifically:

- If it has rained for the past two days, then it will rain tomorrow with probability 0.7.
- If it rained today but not yesterday, then it will rain tomorrow with probability 0.5.
- If it rained yesterday but not today, then it will rain tomorrow with probability 0.4.
- If it has not rained in the past two days, then it will rain tomorrow with probability 0.2.

Is this process Markovian? If so build a Markov chain that models the process.

Discrete systems vs. Markov chains

Some discrete systems are Markov chains ...

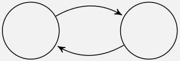

$$\mathbf{P} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

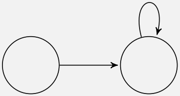

$$\mathbf{P} = \begin{bmatrix} 0 & 1 \\ 0 & 1 \end{bmatrix}$$

Discrete systems vs. Markov chains

Some discrete systems are Markov chains ...

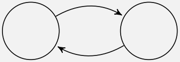

$$\mathbf{P} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

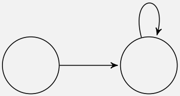

$$\mathbf{P} = \begin{bmatrix} 0 & 1 \\ 0 & 1 \end{bmatrix}$$

... but not all:

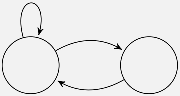

$$\mathbf{P} = \begin{bmatrix} ? & ? \\ 1 & 0 \end{bmatrix}$$

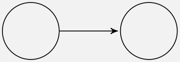

$$\mathbf{P} = \begin{bmatrix} 0 & 1 \\ \mathbf{0} & \mathbf{0} \end{bmatrix}$$

Discrete systems vs. Markov chains

In the other direction, Markov chains are extensions of discrete systems:

MCs contain more information (next-state probabilities).

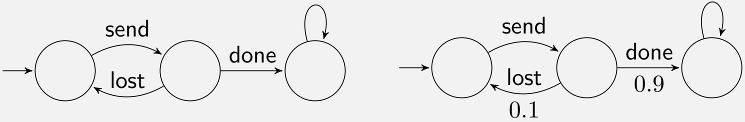

ANALYSIS OF MARKOV CHAINS

Analysis of Markov chains

Interesting questions:

- After k steps, what is the likelihood that the system is at state i?
- In the long run, how much time does the system spend at state i? (i.e., how often is i visited?)
- What is the probability that the system will ever reach a given state (or group of states)?
- What is the expected time until the system reaches a given state (or group of states)?

Computing state probability vectors

Let $\mathbf{x} = [p_1\ p_2\ \cdots\ p_n]$ be a *state probability vector*, where

$$p_i = P(\text{current state is } i)$$

Of course we must have: $\sum_{i=1}^{n} p_i = 1$.

Computing state probability vectors

Let $\mathbf{x} = [p_1 \ p_2 \ \cdots \ p_n]$ be a *state probability vector*, where

$$p_i = P(\text{current state is } i)$$

Of course we must have: $\sum_{i=1}^{n} p_i = 1$.

Let $\mathbf{x}' = [p'_1 \ p'_2 \ \cdots \ p'_n]$ be the **next state** probability vector. Then, for $i = 1, ..., n$:

$$p'_i = p_1 \cdot p_{1i} + p_2 \cdot p_{2i} + \cdots + p_n \cdot p_{ni}$$

Computing state probability vectors

Let $\mathbf{x} = [p_1\ p_2\ \cdots\ p_n]$ be a *state probability vector*, where

$$p_i = P(\text{current state is } i)$$

Of course we must have: $\sum_{i=1}^{n} p_i = 1$.

Let $\mathbf{x}' = [p_1'\ p_2'\ \cdots\ p_n']$ be the **next state** probability vector.
Then, for $i = 1, ..., n$:

$$p_i' = p_1 \cdot p_{1i} + p_2 \cdot p_{2i} + \cdots + p_n \cdot p_{ni}$$

So:

$$\mathbf{x}' = \mathbf{x} \cdot \mathbf{P}$$

Computing state probability vectors

Example:

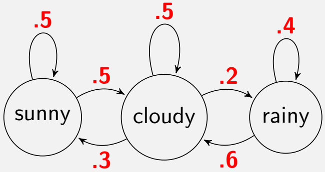

$$[1 \;\; 0 \;\; 0] \cdot \begin{bmatrix} 0.5 & 0.5 & 0 \\ 0.3 & 0.5 & 0.2 \\ 0 & 0.6 & 0.4 \end{bmatrix} = [0.5 \;\; 0.5 \;\; 0]$$

Computing state probability vectors

Example:

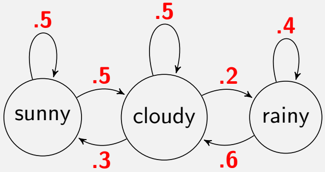

$$[1 \;\; 0 \;\; 0] \cdot \begin{bmatrix} 0.5 & 0.5 & 0 \\ 0.3 & 0.5 & 0.2 \\ 0 & 0.6 & 0.4 \end{bmatrix} = [0.5 \;\; 0.5 \;\; 0]$$

$$[0.5 \;\; 0.5 \;\; 0] \cdot \begin{bmatrix} 0.5 & 0.5 & 0 \\ 0.3 & 0.5 & 0.2 \\ 0 & 0.6 & 0.4 \end{bmatrix} = [0.4 \;\; 0.5 \;\; 0.1]$$

Computing state probability vectors

The probabilities for the **next state** are given by

$$\mathbf{x}' \;=\; \mathbf{x} \cdot \mathbf{P}$$

In general:

$$\mathbf{x}_{k+1} \;=\; \mathbf{x}_k \cdot \mathbf{P} \;=\; (\mathbf{x}_{k-1} \cdot \mathbf{P}) \cdot \mathbf{P} \;=\; \cdots \;=\; \mathbf{x}_0 \cdot \underbrace{\mathbf{P} \cdot \mathbf{P} \cdots \mathbf{P}}_{k+1 \text{ times}}$$

$$\mathbf{x}_{k+1} \;=\; \mathbf{x}_0 \cdot \mathbf{P}^{k+1}$$

Analysis of Markov chains

Interesting questions:

- After k steps, what is the likelihood that the system is at state i?
- In the long run, how much time does the system spend at state i? (i.e., how often is i visited?)
- What is the probability that the system will ever reach a given state (or group of states)?
- What is the expected time until the system reaches a given state (or group of states)?

Which of these questions do we have an answer to already?

Analysis of Markov chains

- After k steps, what is the likelihood that the system is at state i?

$$\mathbf{x}_k \quad = \quad \mathbf{x}_0 \cdot \mathbf{P}^k$$

Analysis of Markov chains

Interesting questions:

- ✓ After k steps, what is the likelihood that the system is at state i?
- In the long run, how much time does the system spend at state i? (i.e., how often is i visited?)
- What is the probability that the system will ever reach a given state (or group of states)?
- What is the expected time until the system reaches a given state (or group of states)?

Let's see how to answer the rest.

Markov chains and graphs

A Markov chain has a graph structure.

We can partly answer the questions simply by studying this structure, **completely ignoring the probability numbers**.

n-step transition probabilities

Let $\mathbf{P}^n_{ij}$ be the (i, j) element of $\mathbf{P}^n$.

What does $\mathbf{P}^n_{ij}$ represent?

n-step transition probabilities

Let $\mathbf{P}^n_{ij}$ be the (i, j) element of $\mathbf{P}^n$.

What does $\mathbf{P}^n_{ij}$ represent?

$\mathbf{P}^n_{ij} = P(s_{k+n} = j \mid s_k = i)$: probability that, starting from state i, after n steps the state will be j.

Classification of states in Markov chains

State i is *absorbing* if $\mathbf{P}_{ii} = 1$. This implies $\mathbf{P}_{ij} = 0$ for all $j \neq i$.

Classification of states in Markov chains

State i is *absorbing* if $\mathbf{P}_{ii} = 1$. This implies $\mathbf{P}_{ij} = 0$ for all $j \neq i$.
This says that i only has a self-loop transition.

Classification of states in Markov chains

State i is *absorbing* if $\mathbf{P}_{ii} = 1$. This implies $\mathbf{P}_{ij} = 0$ for all $j \neq i$.
This says that i only has a self-loop transition.

State j is *accessible* from state i if $\mathbf{P}^n_{ij} > 0$ for some n.

Classification of states in Markov chains

State i is *absorbing* if $\mathbf{P}_{ii} = 1$. This implies $\mathbf{P}_{ij} = 0$ for all $j \neq i$.
This says that i only has a self-loop transition.

State j is *accessible* from state i if $\mathbf{P}_{ij}^n > 0$ for some n.
This says that j is reachable from i.

Classification of states in Markov chains

State i is *absorbing* if $\mathbf{P}_{ii} = 1$. This implies $\mathbf{P}_{ij} = 0$ for all $j \neq i$.
This says that i only has a self-loop transition.

State j is *accessible* from state i if $\mathbf{P}^n_{ij} > 0$ for some n.
This says that j is reachable from i.

Two states i and j *communicate*, written $i \leftrightarrow j$, if i is accessible from j and vice-versa.

A set of states that communicate is called a *class*.

Classification of states in Markov chains

State i is *absorbing* if $\mathbf{P}_{ii} = 1$. This implies $\mathbf{P}_{ij} = 0$ for all $j \neq i$.
This says that i only has a self-loop transition.

State j is *accessible* from state i if $\mathbf{P}^n_{ij} > 0$ for some n.
This says that j is reachable from i.

Two states i and j *communicate*, written $i \leftrightarrow j$, if i is accessible from j and vice-versa.

A set of states that communicate is called a *class*.
A class is a strongly-connected component.

Strongly-connected components

In a directed graph $G = (V, \rightarrow)$, a *strongly-connected component* (SCC) is a subset of nodes $C \subseteq V$, such that every node in C is reachable from every other node in C.

C is called *maximal* if we cannot add more nodes to C and still preserve its SCC property, i.e., $\nexists C' \supset C$ s.t. C' is also a SCC.

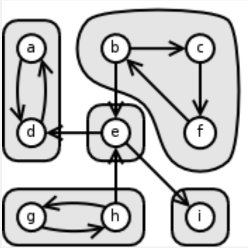

The acyclic graph of maximal SCCs

The set of all maximal SCCs of a graph defines a new graph, where nodes are maximal SCCs, $C_1, C_2, ..., C_m$.

An edge $C_i \rightsquigarrow C_j$ exists iff $C_i \neq C_j$ and there is a node in C_i that has a successor node in C_j.

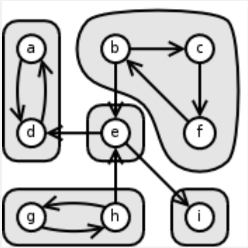

This graph of SCCs is by definition acyclic: why?

Irreducible Markov chains

The Markov chain is *irreducible* if it has only one class, i.e., all states communicate with each other.

Irreducible Markov chains

The Markov chain is *irreducible* if it has only one class, i.e., all states communicate with each other.

This says that the whole Markov chain is a SCC.

Examples

Weather model (irreducible):

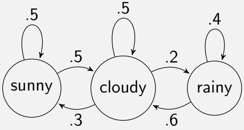

Gambling model (reducible):

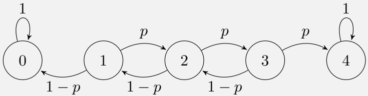

Learning model (reducible):

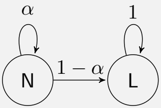

Examples

Weather model (irreducible):

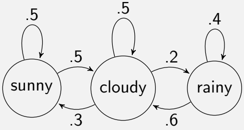

Gambling model (reducible):

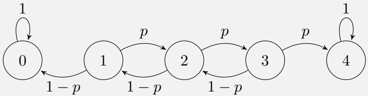

Learning model (reducible):

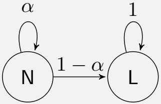

States 0, 4, and L are absorbing states.

Recurrent and transient states

Let i, j be two states. Define:

$f_{ij}^n =$ probability that, starting in i, the first transition into j happens after n steps

$$f_{ij} = \sum_{n=1}^{\infty} f_{ij}^n = \text{probability of reaching } j \text{ from } i \text{ in any \# steps}$$

Recurrent and transient states

Let i, j be two states. Define:

$f_{ij}^n =$ probability that, starting in i, the first transition into j happens after n steps

$$f_{ij} = \sum_{n=1}^{\infty} f_{ij}^n = \text{probability of reaching } j \text{ from } i \text{ in any \# steps}$$

State i is:

- *Recurrent* if $f_{ii} = 1$.
- *Transient* if $f_{ii} < 1$.

Recurrent and transient states

Let i, j be two states. Define:

$f_{ij}^n =$ probability that, starting in i, the first transition into j happens after n steps

$$f_{ij} = \sum_{n=1}^{\infty} f_{ij}^n = \text{probability of reaching } j \text{ from } i \text{ in any \# steps}$$

State i is:

- *Recurrent* if $f_{ii} = 1$.
- *Transient* if $f_{ii} < 1$.

It follows that:

- If i is recurrent then system visits i infinitely often (with probability 1).
- If i is transient then system visits i only finitely many times (with probability 1).

Structural characterization of transient and recurrent states

A SCC C is called *terminal* if there is no C' such that $C \rightsquigarrow C'$.

Otherwise C is called *transient*.

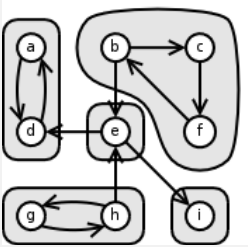

Structural characterization of transient and recurrent states

A SCC C is called *terminal* if there is no C' such that $C \rightsquigarrow C'$.

Otherwise C is called *transient*.

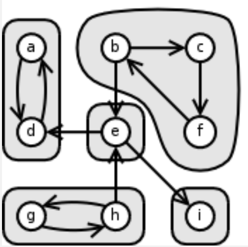

Terminal SCCs: $\{a, d\}$ and $\{i\}$.

Transient SCCs: $\{b, c, f\}$, $\{e\}$, and $\{g, h\}$.

Structural characterization of transient and recurrent states

A SCC C is called *terminal* if there is no C' such that $C \rightsquigarrow C'$.

Otherwise C is called *transient*.

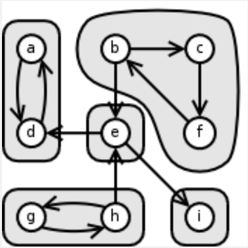

Terminal SCCs: $\{a, d\}$ and $\{i\}$.

Transient SCCs: $\{b, c, f\}$, $\{e\}$, and $\{g, h\}$.

- Recurrent states = states belonging to terminal SCCs.
- Transient states = states belonging to transient SCCs.

Structural characterization of transient and recurrent states

If i is recurrent and $i \leftrightarrow j$ then j is also recurrent.

Structural characterization of transient and recurrent states

If i is recurrent and $i \leftrightarrow j$ then j is also recurrent.

In a finite Markov chain M, some states will be recurrent. Why?

Structural characterization of transient and recurrent states

If i is recurrent and $i \leftrightarrow j$ then j is also recurrent.

In a finite Markov chain M, some states will be recurrent. Why?
Because M is finite and deadlock-free.

Structural characterization of transient and recurrent states

If i is recurrent and $i \leftrightarrow j$ then j is also recurrent.

In a finite Markov chain M, some states will be recurrent. Why?
Because M is finite and deadlock-free.

Does this hold also for infinite Markov chains?

Structural characterization of transient and recurrent states

Viewed as a graph, every finite Markov chain M has at least one terminal maximal SCC.

Theorem

In the long run the amount of time that M spends in transient SCCs is 0.

Therefore, the probability that after some time M will reach a terminal SCC and remain forever there is 1.

Examples

Weather model: all states visited infinitely often.

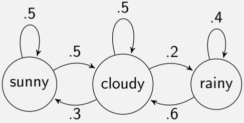

Gambling model: eventually system enters either 0 or 4 and then stays there forever.

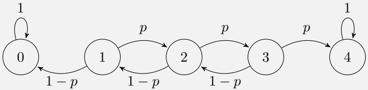

Learning model: eventually system enters L and never leaves.

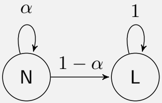

Probabilities vs. Nondeterminism and Fairness

In a probabilistic system, the behavior where the message keeps getting lost after being sent has probability 0, **independently of the value of** p (provided $p > 0$):

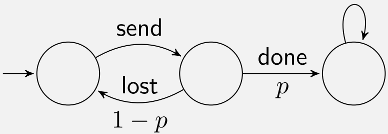

Probabilities vs. Nondeterminism and Fairness

In a non-deterministic system, it is possible that the message always gets lost:

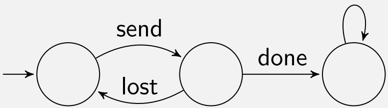

Probabilities vs. Nondeterminism and Fairness

In a non-deterministic system, it is possible that the message always gets lost:

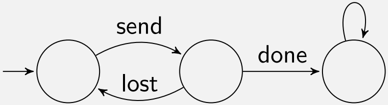

However, we can add fairness constraints to ensure that it does not, e.g.,:

$$\underbrace{\big((\Box\Diamond \text{send}) \rightarrow (\Diamond \text{done})\big)}_{\text{fairness constraint}} \rightarrow (\text{what we want})$$

Probabilities vs. Nondeterminism and Fairness

In a non-deterministic system, it is possible that the message always gets lost:

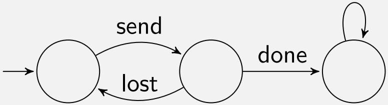

However, we can add fairness constraints to ensure that it does not, e.g.,:

$$\underbrace{\big((\Box\Diamond\text{send}) \rightarrow (\Diamond\text{done})\big)}_{\text{fairness constraint}} \rightarrow (\text{what we want})$$

If this is all we need, probabilities are an overkill.

Examples

Weather model: all states visited infinitely often.

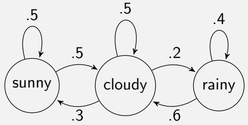

How much time does it rain on the average?

Gambling model: eventually system enters either 0 or 4 and then stays there forever.

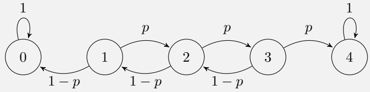

Learning model: eventually system enters L and never leaves.

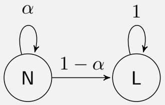

Examples

Weather model: all states visited infinitely often.

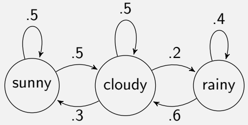

How much time does it rain on the average?

Gambling model: eventually system enters either 0 or 4 and then stays there forever.

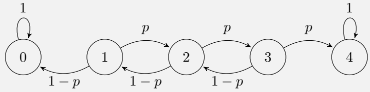

Learning model: eventually system enters L and never leaves.

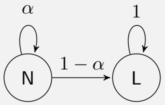

Examples

Weather model: all states visited infinitely often.

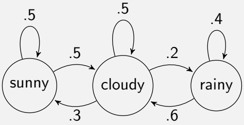

How much time does it rain on the average?

Gambling model: eventually system enters either 0 or 4 and then stays there forever.

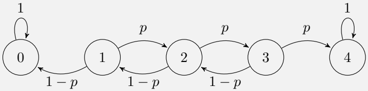

When does the gambler go bankrupt?

Learning model: eventually system enters L and never leaves.

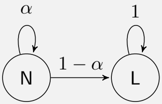

Periodic states

State i has period $d \in \mathbb{N}$ if $\mathbf{P}_{ii}^{n} = 0$ whenever n is not divisible by d, and d is the largest such number.

If $d = 1$ (i.e., $\mathbf{P}_{ii}^{n} > 0$ for all n) then state i is called *aperiodic*.

Periodic states

State i has period $d \in \mathbb{N}$ if $\mathbf{P}_{ii}^n = 0$ whenever n is not divisible by d, and d is the largest such number.

If $d = 1$ (i.e., $\mathbf{P}_{ii}^n > 0$ for all n) then state i is called *aperiodic*.

Examples:

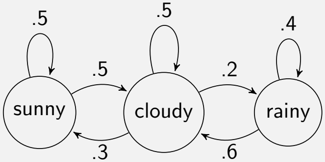

Periodic states

State i has period $d \in \mathbb{N}$ if $\mathbf{P}_{ii}^n = 0$ whenever n is not divisible by d, and d is the largest such number.

If $d = 1$ (i.e., $\mathbf{P}_{ii}^n > 0$ for all n) then state i is called *aperiodic*.

Examples:

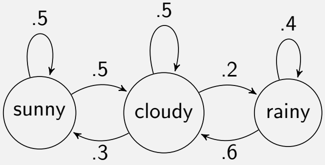

All states are aperiodic.

Periodic states

State i has period $d \in \mathbb{N}$ if $\mathbf{P}^n_{ii} = 0$ whenever n is not divisible by d, and d is the largest such number.

If $d = 1$ (i.e., $\mathbf{P}^n_{ii} > 0$ for all n) then state i is called *aperiodic*.

Examples:

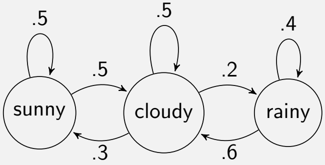

All states are aperiodic.

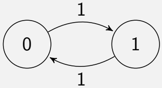

Both states have period 2.

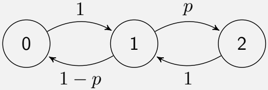

All states have period 2.

Ergodic states

An aperiodic recurrent state is called *ergodic*.

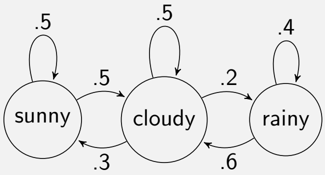

All states are ergodic.

Stationary distribution

Theorem

Let M be a finite, irreducible Markov chain where all states of M are ergodic. Then the limit

$$\lim_{k\to\infty} \mathbf{P}_{ij}^{k}$$

exists and is independent of i (i.e., $\forall i, i' : \lim_{k\to\infty} \mathbf{P}_{ij}^{k} = \lim_{k\to\infty} \mathbf{P}_{i'j}^{k}$).

Stationary distribution

Theorem

Let M be a finite, irreducible Markov chain where all states of M are ergodic. Then the limit

$$\lim_{k\to\infty} \mathbf{P}^k_{ij}$$

exists and is independent of i (i.e., $\forall i, i' : \lim_{k\to\infty} \mathbf{P}^k_{ij} = \lim_{k\to\infty} \mathbf{P}^k_{i'j}$).

Furthermore, letting

$$\pi_j = \lim_{k\to\infty} \mathbf{P}^k_{ij}$$

then $\pi = [\pi_1\ \pi_2\ \cdots\ \pi_n]$ is the unique non-negative solution of

$$\pi = \pi \cdot \mathbf{P}$$

π is called the *stationary distribution*.

Stationary distribution: example

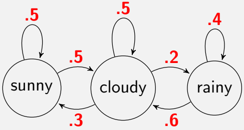

$$\mathbf{P} = \begin{bmatrix} 0.5 & 0.5 & 0 \\ 0.3 & 0.5 & 0.2 \\ 0 & 0.6 & 0.4 \end{bmatrix}$$

Stationary distribution: example

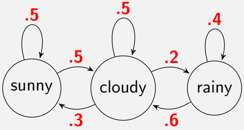

$$\mathbf{P} = \begin{bmatrix} 0.5 & 0.5 & 0 \\ 0.3 & 0.5 & 0.2 \\ 0 & 0.6 & 0.4 \end{bmatrix} \qquad \forall k > 15 : \mathbf{P}^k = \begin{bmatrix} 0.3103 & 0.5172 & 0.1724 \\ 0.3103 & 0.5172 & 0.1724 \\ 0.3103 & 0.5172 & 0.1724 \end{bmatrix}$$

What does this imply?

Stationary distribution: example

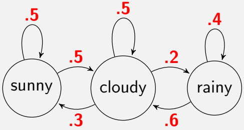

$$\mathbf{P} = \begin{bmatrix} 0.5 & 0.5 & 0 \\ 0.3 & 0.5 & 0.2 \\ 0 & 0.6 & 0.4 \end{bmatrix} \qquad \forall k > 15 : \mathbf{P}^k = \begin{bmatrix} 0.3103 & 0.5172 & 0.1724 \\ 0.3103 & 0.5172 & 0.1724 \\ 0.3103 & 0.5172 & 0.1724 \end{bmatrix}$$

What does this imply?

for any probability vector $\mathbf{x}_0, \forall k > 15 : \mathbf{x}_0 \cdot \mathbf{P}^k = \pi = [0.3103\ 0.5172\ 0.1724]$

Stationary distribution: example

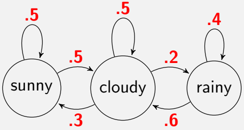

$$\mathbf{P} = \begin{bmatrix} 0.5 & 0.5 & 0 \\ 0.3 & 0.5 & 0.2 \\ 0 & 0.6 & 0.4 \end{bmatrix} \qquad \forall k > 15 : \mathbf{P}^k = \begin{bmatrix} 0.3103 & 0.5172 & 0.1724 \\ 0.3103 & 0.5172 & 0.1724 \\ 0.3103 & 0.5172 & 0.1724 \end{bmatrix}$$

What does this imply?

for any probability vector $\mathbf{x}_0, \forall k > 15 : \mathbf{x}_0 \cdot \mathbf{P}^k = \pi = [0.3103\ 0.5172\ 0.1724]$

i.e., stationary distribution π independent from the initial state distribution $\mathbf{x}_0$.

Stationary distribution: example

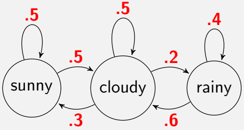

$$\mathbf{P} = \begin{bmatrix} 0.5 & 0.5 & 0 \\ 0.3 & 0.5 & 0.2 \\ 0 & 0.6 & 0.4 \end{bmatrix} \qquad \forall k > 15 : \mathbf{P}^k = \begin{bmatrix} 0.3103 & 0.5172 & 0.1724 \\ 0.3103 & 0.5172 & 0.1724 \\ 0.3103 & 0.5172 & 0.1724 \end{bmatrix}$$

What does this imply?

for any probability vector $\mathbf{x}_0, \forall k > 15 : \mathbf{x}_0 \cdot \mathbf{P}^k = \pi = [0.3103\ 0.5172\ 0.1724]$

i.e., stationary distribution π independent from the initial state distribution $\mathbf{x}_0$.

So how much time does it rain on the average?

Stationary distribution

If the chain is not ergodic, the limit may not exist, e.g.,

$$\mathbf{P} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

Analysis of Markov chains

Interesting questions:

✓ After k steps, what is the likelihood that the system is at state i?

- In the long run, how much time does the system spend at state i? (i.e., how often is i visited?)
- What is the probability that the system will ever reach a given state (or group of states)?
- What is the expected time until the system reaches a given state (or group of states)?

Which of these questions do we have an answer to already?

Analysis of Markov chains

Interesting questions:

- ✓ After k steps, what is the likelihood that the system is at state i?
- ✓ In the long run, how much time does the system spend at state i? (i.e., how often is i visited?)
- What is the probability that the system will ever reach a given state (or group of states)?
- What is the expected time until the system reaches a given state (or group of states)?

Examples

Weather model:

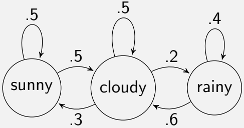

Gambling model:

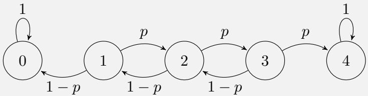

How much time does an average game last?

Learning model:

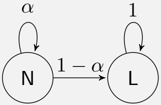

How long until we learn something?

Transient analysis

Order the states of a Markov chain M so that $\{1, 2, ..., t\}$ is the set of transient states.

Let

$$\mathbf{P}_T = \begin{bmatrix} p_{11} & p_{12} & \cdots & p_{1t} \\ p_{21} & p_{22} & \cdots & p_{2t} \\ \vdots & & \ddots & \vdots \\ p_{t1} & p_{t2} & \cdots & p_{tt} \end{bmatrix}$$

Observation: some rows of $\mathbf{P}_T$ sum to < 1 (otherwise this would be a SCC).

Transient analysis

Let

$q_{ij} =$ mean time spent in j, given that the system starts in i

Then

$$q_{ij} = \delta_{ij} + \sum_k p_{ik} \cdot q_{kj}$$

where

$$\delta_{ij} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{otherwise.} \end{cases}$$

Transient analysis

Let

$$q_{ij} = \text{mean time spent in } j\text{, given that the system starts in } i$$

Then

$$\begin{aligned} q_{ij} &= \delta_{ij} + \sum_{k} p_{ik} \cdot q_{kj} \\ &= \delta_{jk} + \sum_{k=1}^{t} p_{ik} \cdot q_{kj} \qquad \textit{why?} \end{aligned}$$

where

$$\delta_{ij} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{otherwise.} \end{cases}$$

Transient analysis

Let

$$q_{ij} = \text{mean time spent in } j\text{, given that the system starts in } i$$

Then

$$\begin{aligned} q_{ij} &= \delta_{ij} + \sum_k p_{ik} \cdot q_{kj} \\ &= \delta_{jk} + \sum_{k=1}^{t} p_{ik} \cdot q_{kj} \qquad why? \end{aligned}$$

Because $q_{kj} = 0$ when k is recurrent (cannot move from recurrent state to transient state).

where

$$\delta_{ij} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{otherwise.} \end{cases}$$

Transient analysis

Let

$$\mathbf{Q} = \begin{bmatrix} q_{11} & q_{12} & \cdots & q_{1t} \\ q_{21} & q_{22} & \cdots & q_{2t} \\ \vdots & & \ddots & \vdots \\ q_{t1} & q_{t2} & \cdots & q_{tt} \end{bmatrix}$$

Then

$$\mathbf{Q} = \mathbf{I} + \mathbf{P}_T \cdot \mathbf{Q}$$

where $\mathbf{I}$ is the identity matrix of size t.

Transient analysis

Let

$$\mathbf{Q} = \begin{bmatrix} q_{11} & q_{12} & \cdots & q_{1t} \\ q_{21} & q_{22} & \cdots & q_{2t} \\ \vdots & & \ddots & \vdots \\ q_{t1} & q_{t2} & \cdots & q_{tt} \end{bmatrix}$$

Then

$$\mathbf{Q} = \mathbf{I} + \mathbf{P}_T \cdot \mathbf{Q}$$

where $\mathbf{I}$ is the identity matrix of size t.

It can be shown that $\mathbf{I} - \mathbf{P}_T$ is invertible. Therefore:

$$\mathbf{Q} = (\mathbf{I} - \mathbf{P}_T)^{-1}$$

Example: gambler's ruin

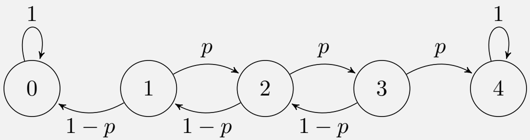

$$\mathbf{Q} = (\mathbf{I} - \mathbf{P}_T)^{-1} = \begin{bmatrix} 1 & -p & 0 \\ p-1 & 1 & -p \\ 0 & p-1 & 1 \end{bmatrix}^{-1} = \cdots$$

$$\mathbf{Q}_{p=0} = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 1 & 1 \end{bmatrix} \quad \mathbf{Q}_{p=1} = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix} \quad \mathbf{Q}_{p=\frac{1}{2}} = \begin{bmatrix} 1.5 & 1 & 0.5 \\ 1 & 2 & 1 \\ 0.5 & 1 & 1.5 \end{bmatrix}$$

Example: gambler's ruin

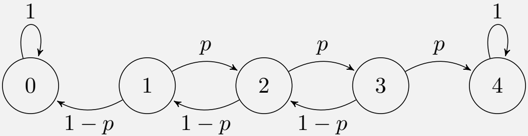

$$\mathbf{Q} = (\mathbf{I} - \mathbf{P}_T)^{-1} = \begin{bmatrix} 1 & -p & 0 \\ p-1 & 1 & -p \\ 0 & p-1 & 1 \end{bmatrix}^{-1} = \cdots$$

$$\mathbf{Q}_{p=0} = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 1 & 1 \end{bmatrix} \quad \mathbf{Q}_{p=1} = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix} \quad \mathbf{Q}_{p=\frac{1}{2}} = \begin{bmatrix} 1.5 & 1 & 0.5 \\ 1 & 2 & 1 \\ 0.5 & 1 & 1.5 \end{bmatrix}$$

What is the average playing time with $p = \frac{1}{2}$?

Example: gambler's ruin

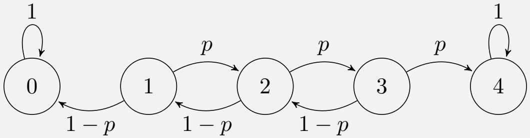

$$\mathbf{Q} = (\mathbf{I} - \mathbf{P}_T)^{-1} = \begin{bmatrix} 1 & -p & 0 \\ p-1 & 1 & -p \\ 0 & p-1 & 1 \end{bmatrix}^{-1} = \cdots$$

$$\mathbf{Q}_{p=0} = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 1 & 1 \end{bmatrix} \quad \mathbf{Q}_{p=1} = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix} \quad \mathbf{Q}_{p=\frac{1}{2}} = \begin{bmatrix} 1.5 & 1 & 0.5 \\ 1 & 2 & 1 \\ 0.5 & 1 & 1.5 \end{bmatrix}$$

What is the average playing time with $p = \frac{1}{2}$?

3 if I start with \$1 or \$3. 4 if I start with \$2.

Analysis of Markov chains

Interesting questions:

- ✓ After k steps, what is the likelihood that the system is at state i?
- ✓ In the long run, how much time does the system spend at state i? (i.e., how often is i visited?)
- ▶ What is the probability that the system will ever reach a given state (or group of states)?
- ✓ What is the expected time until the system reaches a given state (or group of states)?

REACHABILITY PROBABILITIES IN MARKOV CHAINS

Material taken mainly from [BK08], Chapter 10.

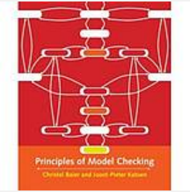

Reachability question for Markov chains

Let M be a Markov chain and B a set of states of M.

Reachability question for Markov chains: what is the probability of reaching B?

Note: this is not a yes/no question, as in standard model-checking. Here, we want to compute a probability $p \in [0, 1]$.

Example: gambler's ruin

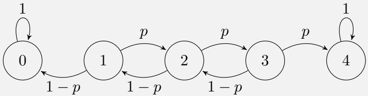

When does the gambler go bankrupt?

Model in PRISM:

```
dtmc  // this model is a Markov chain

const double p = 0.7;

module M
    s : [0..4] init 2;

    [] s=0 -> (s'=0);
    [] s=1 -> p:(s'=2) + (1-p):(s'=0);
    [] s=2 -> p:(s'=3) + (1-p):(s'=1);
    [] s=3 -> p:(s'=4) + (1-p):(s'=2);
    [] s=4 -> (s'=4);
endmodule
```

PCTL formula in PRISM:

`P=? [ F s=4 ]`

PRISM answers:

p	answer
0	0
0.5	$0.499999\cdots$
0.7	$0.844827\cdots$
1	1

Computing reachability probabilities

Let x_i be the probability that the target set B is reached starting from state i.

Then:

- If $i \in B$ then $x_i = 1$.
- If i cannot reach B in the graph sense, then $x_i = 0$.
- Otherwise

$$x_i = \sum_j p_{ij} \cdot x_j$$

This forms a set of linear equations. For finite chains it is finite and is guaranteed to have a unique solution.

Example: gambler's ruin

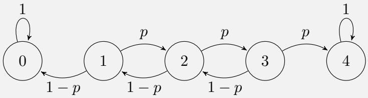

When does the gambler go bankrupt?

$$
\begin{aligned}
x_0 &= 0 \\
x_4 &= 1 \\
x_1 &= p \cdot x_2 + (1-p) \cdot x_0 = p \cdot x_2 \\
x_2 &= p \cdot x_3 + (1-p) \cdot x_1 \\
x_3 &= p \cdot x_4 + (1-p) \cdot x_2 = p + (1-p) \cdot x_2
\end{aligned}
$$

For $p = \frac{1}{2}$, $x_2 = \frac{1}{2}$. For $p = 0.7$, $x_2 = 0.8448$.

Example: gambler's ruin

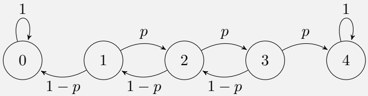

When does the gambler go bankrupt?

$$
\begin{aligned}
x_0 &= 0 \quad \text{This is necessary for uniqueness of solution. Why?} \\
x_4 &= 1 \\
x_1 &= p \cdot x_2 + (1-p) \cdot x_0 = p \cdot x_2 \\
x_2 &= p \cdot x_3 + (1-p) \cdot x_1 \\
x_3 &= p \cdot x_4 + (1-p) \cdot x_2 = p + (1-p) \cdot x_2
\end{aligned}
$$

For $p = \frac{1}{2}$, $x_2 = \frac{1}{2}$. For $p = 0.7$, $x_2 = 0.8448$.

COMPOSITION OF MARKOV CHAINS

and

MARKOV DECISION PROCESSES

Composition of Markov chains

Suppose we want to compose the following two Markov chains:

PRISM code:

```
module M1
    s : [0..2] init 0;
    [] s=0 -> (s'=1);
    [] s=1 -> 0.5:(s'=0) + 0.5:(s'=2);
    [] s=2 -> (s'=2);
endmodule
```

```
module M2 = M1 [ s=t ] endmodule
```

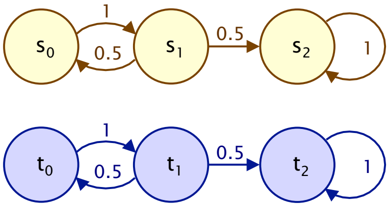

Several figures due to Dave Parker.

Synchronous composition of Markov chains

What does the synchronous composition of these processes look like?

PRISM code:

```
module M1
    s : [0..2] init 0;
    [] s=0 -> (s'=1);
    [] s=1 -> 0.5:(s'=0) + 0.5:(s'=2);
    [] s=2 -> (s'=2);
endmodule
```

```
module M2 = M1 [ s=t ] endmodule
```

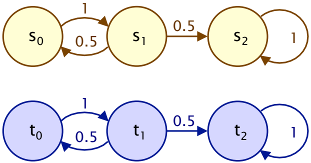

Synchronous composition of Markov chains

What does the synchronous composition of these processes look like?

PRISM code:

```
module M1
   s : [0..2] init 0;
   [] s=0 -> (s'=1);
   [] s=1 -> 0.5:(s'=0) + 0.5:(s'=2);
   [] s=2 -> (s'=2);
endmodule
```

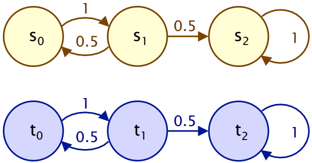


```
module M2 = M1 [ s=t ] endmodule
```

$$P((s_1, t_1) \mid (s_0, t_0)) \quad =$$

Synchronous composition of Markov chains

What does the synchronous composition of these processes look like?

PRISM code:

```
module M1
    s : [0..2] init 0;
    [] s=0 -> (s'=1);
    [] s=1 -> 0.5:(s'=0) + 0.5:(s'=2);
    [] s=2 -> (s'=2);
endmodule
```

```
module M2 = M1 [ s=t ] endmodule
```

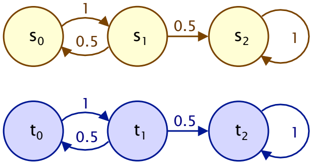

$$P((s_1, t_1) \mid (s_0, t_0)) \quad = \quad 1 \cdot 1 = 1$$

Synchronous composition of Markov chains

What does the synchronous composition of these processes look like?

PRISM code:

```
module M1
    s : [0..2] init 0;
    [] s=0 -> (s'=1);
    [] s=1 -> 0.5:(s'=0) + 0.5:(s'=2);
    [] s=2 -> (s'=2);
endmodule
```

```
module M2 = M1 [ s=t ] endmodule
```

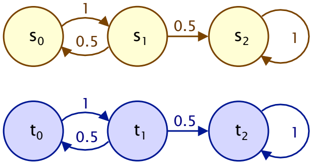

$$P((s_1,t_1) \mid (s_0,t_0)) = 1 \cdot 1 = 1$$
$$P((s_2,t_2) \mid (s_1,t_1)) =$$

Synchronous composition of Markov chains

What does the synchronous composition of these processes look like?

PRISM code:

```
module M1
    s : [0..2] init 0;
    [] s=0 -> (s'=1);
    [] s=1 -> 0.5:(s'=0) + 0.5:(s'=2);
    [] s=2 -> (s'=2);
endmodule
```

```
module M2 = M1 [ s=t ] endmodule
```

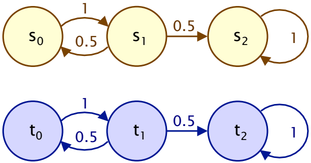

$$
\begin{aligned}
P((s_1, t_1) \mid (s_0, t_0)) &= 1 \cdot 1 = 1 \\
P((s_2, t_2) \mid (s_1, t_1)) &= \frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4}
\end{aligned}
$$

Synchronous composition of Markov chains

What does the synchronous composition of these processes look like?

PRISM code:

```
module M1
    s : [0..2] init 0;
    [] s=0 -> (s'=1);
    [] s=1 -> 0.5:(s'=0) + 0.5:(s'=2);
    [] s=2 -> (s'=2);
endmodule
```

```
module M2 = M1 [ s=t ] endmodule
```

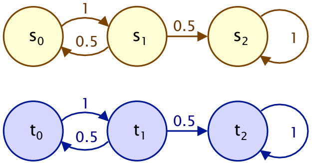

$$P((s_1,t_1) \mid (s_0,t_0)) = 1 \cdot 1 = 1$$
$$P((s_2,t_2) \mid (s_1,t_1)) = \frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4}$$
$$\ldots$$

Synchronous composition of Markov chains

What does the synchronous composition of these processes look like?

PRISM code:

```
module M1
    s : [0..2] init 0;
    [] s=0 -> (s'=1);
    [] s=1 -> 0.5:(s'=0) + 0.5:(s'=2);
    [] s=2 -> (s'=2);
endmodule
```

```
module M2 = M1 [ s=t ] endmodule
```

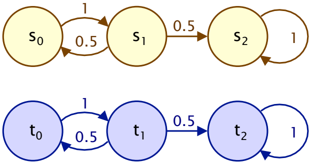

$$
\begin{aligned}
P((s_1,t_1) \mid (s_0,t_0)) &= 1 \cdot 1 = 1 \\
P((s_2,t_2) \mid (s_1,t_1)) &= \frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4} \\
&\cdots
\end{aligned}
$$

Is the synchronous composition of two Markov chains a Markov chain?

Synchronous composition of Markov chains

What does the synchronous composition of these processes look like?

PRISM code:

```
module M1
    s : [0..2] init 0;
    [] s=0 -> (s'=1);
    [] s=1 -> 0.5:(s'=0) + 0.5:(s'=2);
    [] s=2 -> (s'=2);
endmodule
```

```
module M2 = M1 [ s=t ] endmodule
```

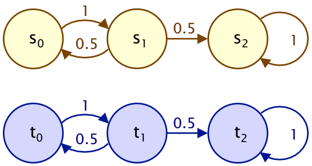

$$P((s_1, t_1) \mid (s_0, t_0)) = 1 \cdot 1 = 1$$
$$P((s_2, t_2) \mid (s_1, t_1)) = \frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4}$$
$$\cdots$$

Is the synchronous composition of two Markov chains a Markov chain?

Yes!

Asynchronous composition of Markov chains

What would the asynchronous composition of these two processes look like?

PRISM code:

```
module M1
    s : [0..2] init 0;
    [] s=0 -> (s'=1);
    [] s=1 -> 0.5:(s'=0) + 0.5:(s'=2);
    [] s=2 -> (s'=2);
endmodule
```

```
module M2 = M1 [ s=t ] endmodule
```

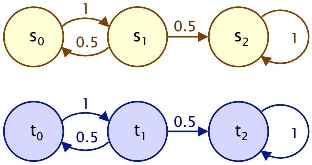

Asynchronous composition of Markov chains

Asynchronous parallel composition of two 3-state DTMCs

Action labels omitted here

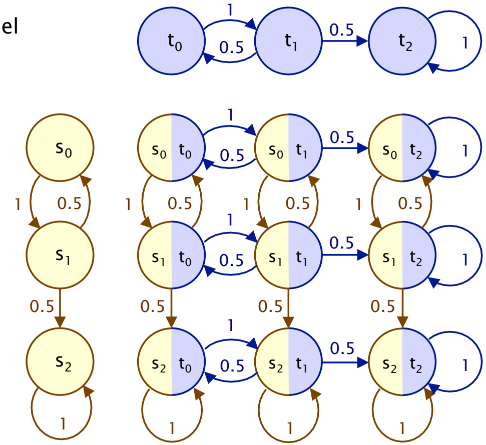

Asynchronous composition of Markov chains

Asynchronous parallel composition of two 3-state DTMCs

Action labels omitted here

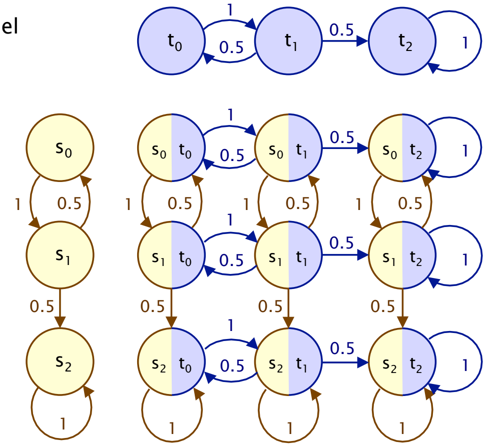

Is the asynchronous composition of two Markov chains a Markov chain?

Asynchronous composition of Markov chains

Asynchronous parallel composition of two 3-state DTMCs

Action labels omitted here

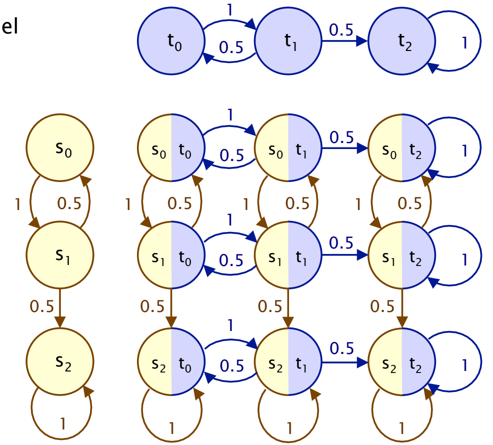

Is the asynchronous composition of two Markov chains a Markov chain? No! It is a Markov Decision Process.

Markov Decision Processes (MDPs)

Combine non-deterministic and probabilistic choice.

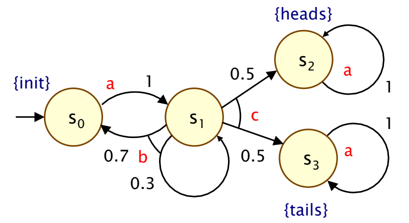

Intuitive semantics:

- First choose action non-deterministically among possible actions.
  - In state s_0, only one possible action, a.
  - In state s_1, two possible actions, b and c.
- Then, given chosen action, throw a dice a pick successor state w.r.t. the specified probability distribution for that action.

Markov Decision Processes (MDPs)

Non-determinism has multiple uses, as in discrete systems.
E.g., useful to model abstraction:

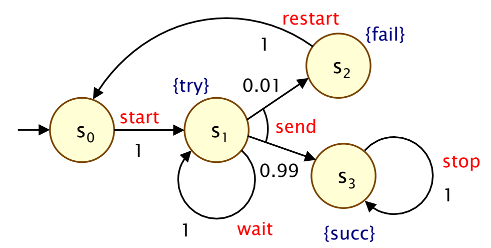

At state s_1, if channel is ready attempt to send, otherwise wait.
Details of when channel is ready are not modeled.

Model-checking MDPs

Answers queries like:

- Byzantine agreement protocol
  - $P_{min=?}$ [F (agreement $\wedge$ rounds$\leq$2)]
  - "what is the minimum probability that agreement is reached within two rounds?"
- CSMA/CD communication protocol
  - $P_{max=?}$ [F collisions=k]
  - "what is the maximum probability of k collisions?"
- Self-stabilisation protocols
  - $P_{min=?}$ [$F^{\leq t}$ stable]
  - "what is the minimum probability of reaching a stable state within k steps?"

See PRISM web site and literature for details:
http://www.prismmodelchecker.org/

Bibliography

C. Baier and J.-P. Katoen.
Principles of Model Checking.
MIT Press, 2008.

PRISM lecture material available online:
`http://www.prismmodelchecker.org/lectures/`.

S. M. Ross.
Introduction to Probability Models.
Academic Press, 2006.