Markov Chains

Principles of Online Decision-Making (CS-303)

Prof. Matthias Grossglauser

Information and Network Dynamics (INDY) lab
School of Computer and Communication Sciences (I&C)
EPFL

Andrej Markov (1856-1922)

- Attended Saint Petersburg Imperial University
- Doctoral advisor: Pafnuty Chebyshev (we've seen inequalities named after both)
- 1894: Professorship at St. Petersburg University
- 1908: Removed from teaching after refusing to spy on students during a period of unrest
- 1917: reinstated, resumed teaching and research until his death

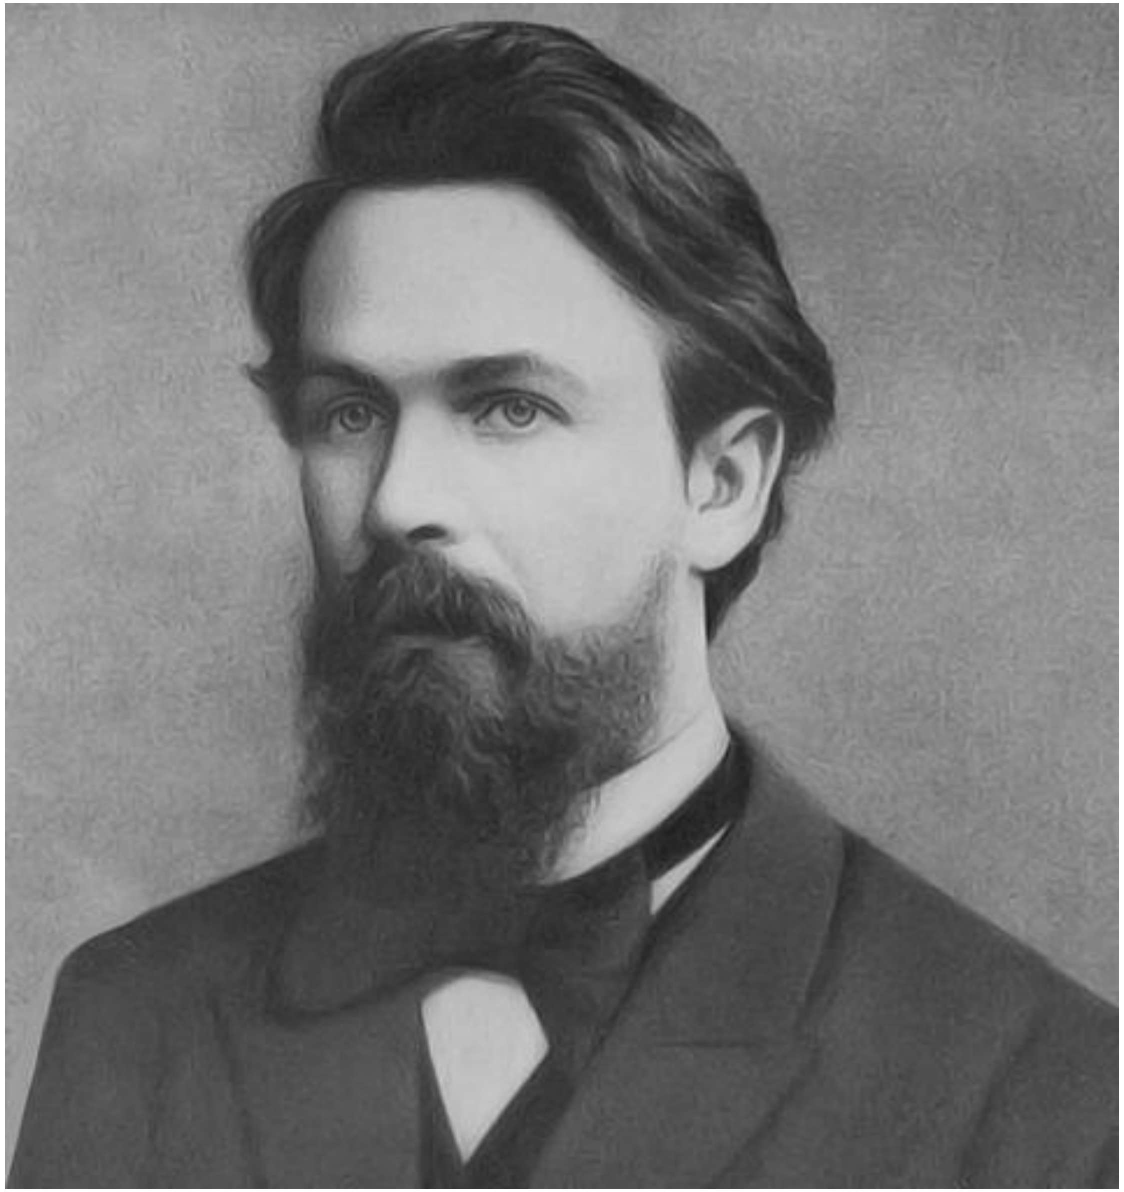

Markov condition

- Definition: the process $(X_0, X_1, X_2, \ldots)$ is a Markov chain if it satisfies the Markov condition:

$$\mathbb{P}[X_n = x | X_0 = x_0, X_1 = x_1, \ldots, X_{n-1} = x_{n-1}] = \mathbb{P}[X_n = x | X_{n-1} = x_{n-1}]$$

$X_0 \quad X_1 \quad X_2 \quad X_3 \quad X_4 \quad \ldots \quad X_{n-2} \quad X_{n-1} \quad X_n$ →

for all $n \geq 1$ and all $x, x_0, x_1, \ldots, x_{n-1} \in \mathcal{S}$

- Intuitively: to predict the future, no knowledge of the past is useful beyond the current state
- Note: we study discrete time, discrete space MC only
  - Discrete time: $t \in \mathbb{N}_0$
  - Discrete space: countable set $\mathcal{S}$ = state space
    - Later we will restrict $\mathcal{S}$ to finite

Equivalent Markov conditions

- $\mathbb{P}\left[X_{n+1} = x \middle| X_{n_1} = x_{n_1}, X_{n_2} = x_{n_2}, \dots, X_{n_k} = x_{n_k}\right] = \mathbb{P}\left[X_{n+1} = x \middle| X_{n_k} = x_{n_k}\right]$

  for all $n_1 < n_2 < \cdots < n_k \le n$

- $\mathbb{P}[X_{n+m} = x | X_0 = x_0, X_1 = x_1, \dots, X_n = x_n] = \mathbb{P}[X_{n+m} = x | X_n = x_n]$

  for all $m, n \ge 0$

Homogeneous chain

- The chain X is homogeneous (time-shift-invariant) if
$$\mathbb{P}[X_{n+1} = j | X_n = i] = \mathbb{P}[X_1 = j | X_0 = i]$$
for all n, i, j
  - Corresponds to "identically distributed" for i.id. RVs
- Transition matrix P: $|\mathcal{S}| \times |\mathcal{S}|$ matrix of transition probabilities:
  - $p_{ij} = \mathbb{P}[X_{n+1} = j | X_n = i]$
- From now on, unless otherwise specified, we assume homogeneous chains
- P is a **stochastic matrix**:
  - (a) Entries are non-negative: $p_{ij} \geq 0$
  - (b) Row sums equal to one: $\sum_j p_{ij} = 1$

P_n: n-step transition matrix

- P_n: $p_{ij}(n) = \mathbb{P}[X_{m+n} = j | X_m = i]$
- $P_1 = P$
- Example:
  - $P_1 = \begin{bmatrix} 1/2 & 1/2 \\ 3/4 & 1/4 \end{bmatrix}$
  - $P_2 = \begin{bmatrix} 0.625 & 0.375 \\ 0.563 & 0.438 \end{bmatrix}$
  - $P_3 = \begin{bmatrix} 0.594 & 0.406 \\ 0.609 & 0.391 \end{bmatrix}$
  - $P_{10} = \begin{bmatrix} 0.600 & 0.400 \\ 0.600 & 0.400 \end{bmatrix}$
  - After 10 steps, we end up in state 1 or 2 pretty much with the same probability regardless of where we started!
    - Ergodicity→ we shall return to this topic!

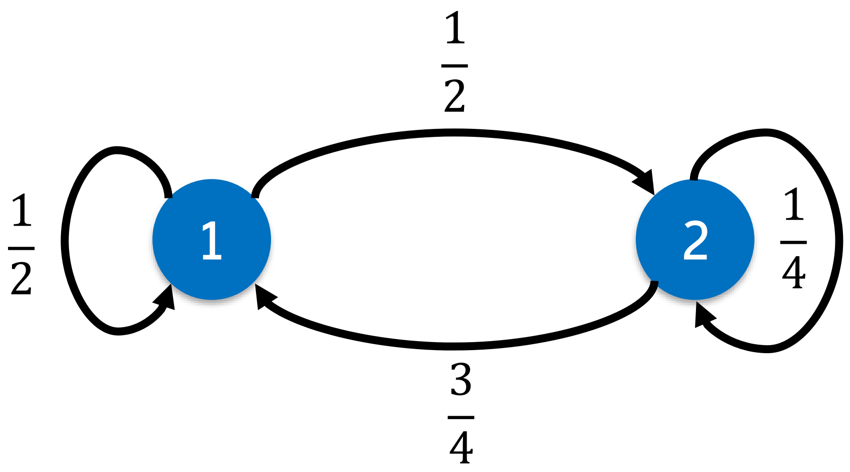

Chapman-Kolmogorov equations

- $p_{ij}(m+n) = \sum_k p_{ik}(m) p_{kj}(n)$

- Proof:
  - $p_{ij}(m+n) = \mathbb{P}[X_{m+n} = j | X_0 = i]$

$= \sum_k \mathbb{P}[X_{m+n} = j, X_m = k | X_0 = i]$ Law of total probability

$= \sum_k \mathbb{P}[X_{m+n} = j | X_m = k, X_0 = i] \ \mathbb{P}[X_m = k | X_0 = i]$ $\mathbb{P}[A, B | C] = \mathbb{P}[A | B, C] \mathbb{P}[B | C]$

$= \sum_k \mathbb{P}[X_{m+n} = j | X_m = k] \ \mathbb{P}[X_m = k | X_0 = i]$ Markov property

- In matrix form: $P_{m+n} = P_m \cdot P_n$
  - Hence $\boldsymbol{P_n} = \boldsymbol{P^n}$

Evolution of marginal distribution of $\boldsymbol{X_n}$

- $\mu_i(n) = \mathbb{P}[X_n = i]$, $\boldsymbol{\mu}(n) = [\mu_1(n), \mu_2(n), \dots]$
- Lemma: $\boldsymbol{\mu}(m+n) = \boldsymbol{\mu}(m)P_n$
- Proof:
  - $\mu_j(m+n) = \mathbb{P}[X_{m+n} = j]$

$$= \sum_i \mathbb{P}[X_{m+n} = j | X_m = i]\, \mathbb{P}[X_m = i]$$

Law of total probability

$$= \sum_i \mu_i(m) p_{ij}(n) = [\boldsymbol{\mu}(m)P_n]_j$$

- Hence $\boldsymbol{\mu}(n) = \boldsymbol{\mu}(0)P^n$
- Key point: random evolution of MC determined by $\boldsymbol{\mu}(0)$ (initial conditions) and P (dynamics)

Examples

- The simple random walk:
  - $\mathcal{S} = \{0, \pm 1, \pm 2, \pm 3, \pm 4, \pm 5, \dots\}$
  - $p_{ij} = \begin{cases} p & j = i + 1 \\ q = 1 - p & j = i - 1 \\ 0 & \text{otherwise} \end{cases}$
  - To go from i to j, we need to do $j - i$ more steps to the right than to the left
  - So $r + l = n$, and $r - l = j - i$
  - There are $\binom{n}{r}$ sequences of length n with r right steps and $l = n - r$ left steps
  - $r = \frac{1}{2}(n + j - i)$ (needs to be integer)
  - $p_{ij(n)} = \begin{cases} \binom{n}{\frac{1}{2}(n+j-i)} p^{\frac{1}{2}(n+j-i)} q^{1/2(n-j+i)} & n + j - i \text{ even} \\ 0 & \text{otherwise} \end{cases}$

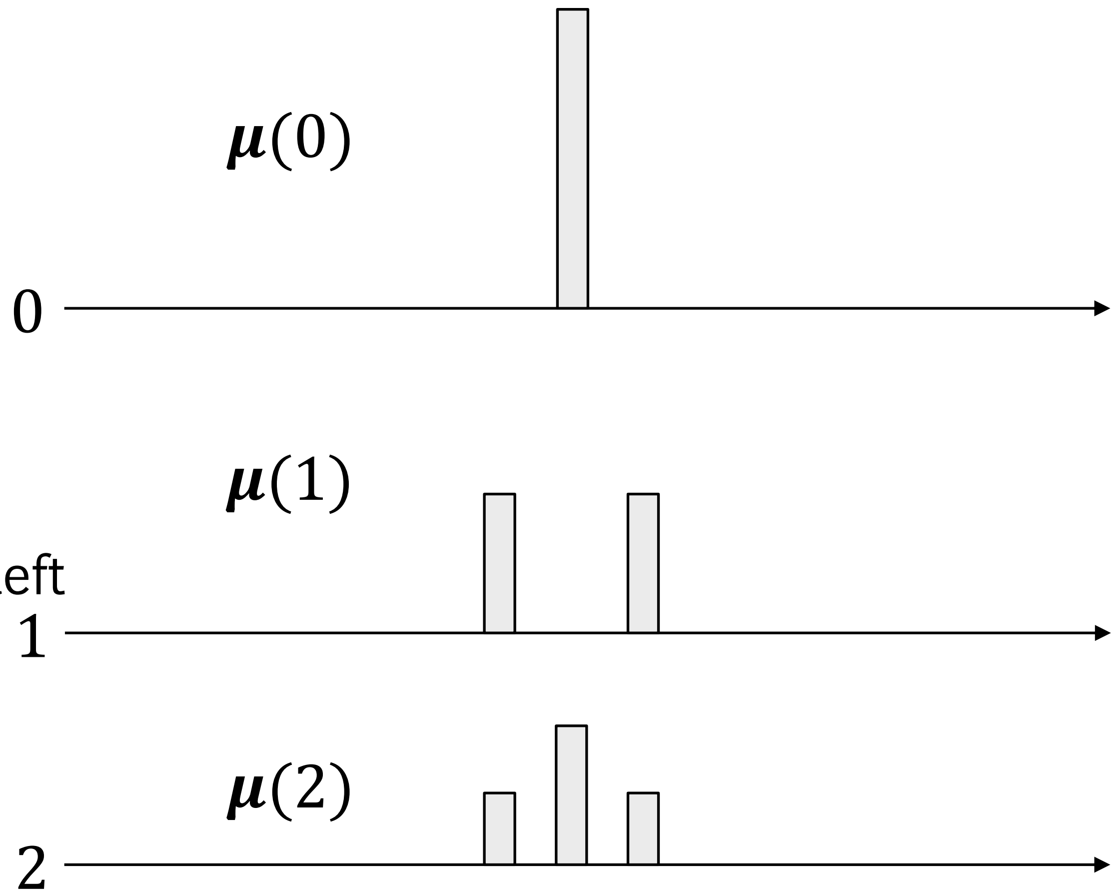

Chapman-Kolmogorov $\nRightarrow$ Markov

- Let $Y_1, Y_3, Y_5, \ldots$ an iid sequence of Bernoulli(1/2) RV with values ± 1
  - $(\mathbb{P}[Y_{2k+1} = -1] = \mathbb{P}[Y_{2k+1} = +1] = 1/2)$
- Let $Y_{2k} = Y_{2k-1} Y_{2k+1}$
  - The sequence $Y_2, Y_4, Y_6, \ldots$ is also iid with same distribution
- Also, even adjacent pairs Y_{2k} and Y_{2k+1} are independent
- Therefore, the whole sequence of RVs is pairwise independent
- Hence $p_{ij}(n) = \mathbb{P}[Y_{m+n} = j | Y_m = i] = \frac{1}{2}$ → Chapman-Kolmogorov condition satisfied

- But is Y a Markov chain? No!
  - $\mathbb{P}[Y_{2k+1} = 1 | Y_{2k} = -1] = \frac{1}{2}$
  - But $\mathbb{P}[Y_{2k+1} = 1 | Y_{2k} = -1, Y_{2k-1} = 1] = 0$

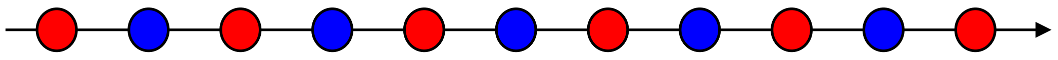

- CK necessary but not sufficient for Markov

Classification of states

- Def: State i is **recurrent** (or **persistent**) if
$$\mathbb{P}[X_n = i \text{ for some } n \geq 1 | X_0 = i] = 1$$
Note: starting at i is important in the definition – we may never visit a state starting from some $j \neq i$, even though i **is** recurrent!
- Def: State i is **transient** otherwise, ie, if
$$\mathbb{P}[X_n = i \text{ for some } n \geq 1 | X_0 = i] < 1$$
- Def: first passage time:
Probability of **first** visit to state j starting at state i **in** $\boldsymbol{n}$ **steps**:

$$f_{ij}(n) = \mathbb{P}[X_1 \neq j, X_2 \neq j, \dots, X_{n-1} \neq j, X_n = j | X_0 = i]$$

Probability of **ever** visiting j after i:
$$f_{ij} = \sum_{n=1}^{\infty} f_{ij}(n)$$
- Note: j is recurrent $\Leftrightarrow$ $f_{jj} = 1$

Linking transition probabilities and first passage times

- Consider the probability generating functions of $p_{ij}(n)$ and of $f_{ij}(n)$:
  - $P_{ij}(s) = \sum_n s^n p_{ij}(n)$ $(p_{ij}(0) = \delta_{ij})$
  - $F_{ij}(s) = \sum_n s^n f_{ij}(n)$ $(f_{ij}(0) = 0)$
- Theorem:
  - (a) $P_{ii}(s) = 1 + F_{ii}(s)P_{ii}(s)$
  - (b) $P_{ij}(s) = F_{ij}(s)P_{jj}(s)$ (for $i \neq j$)
- Proof:
  - Fix two states i, j
  - $A_m = \{X_m = j\}$
  - $B_m = \{X_r \neq j$ for $1 \leq r < m, X_m = j\}$ (first visit to j happens at time m)
    - Note: B_m are disjoint, and $B_m \Longrightarrow A_m$
  - $p_{ij}(m) = \mathbb{P}[A_m | X_0 = i] = \sum_{r=1}^{m} \mathbb{P}[A_m, B_r | X_0 = i]$
  - $\mathbb{P}[A_m, B_r | X_0 = i] = \mathbb{P}[A_m | B_r, X_0 = i]\mathbb{P}[B_r | X_0 = i]$

$$= \underbrace{\mathbb{P}[A_m | X_r = j\,]}_{p_{jj}(m-r)}\underbrace{\mathbb{P}[B_r | X_0 = i]}_{f_{ij}(r)}$$

$\mathbb{P}[A_m | B_r, X_0 = i] = \mathbb{P}[A_m | X_r = j\,]$:
From Markov property: A_m depends only on $X_r = j$, not on what happened before!

Linking transition probabilities and first passage times

- Theorem:
  - $P_{ij}(s) = F_{ij}(s)P_{jj}(s) + \delta_{ij}$
- Proof (cont):
  - $p_{ij}(m) = \sum_{r=1}^{m} f_{ij}(r)p_{jj}(m-r) \qquad (m = 1,2,\ldots)$
  - This is a convolution → multiplication in the transform domain
  - $P_{ij}(s) = \sum_{m=1}^{\infty} s^m p_{ij}(m) + \delta_{ij}$

$$= \sum_{m=1}^{\infty} s^m \sum_{r=1}^{m} f_{ij}(r)p_{jj}(m-r) + \delta_{ij}$$

$$= \sum_{m=1}^{\infty} \sum_{r=1}^{m} s^r f_{ij}(r) s^{m-r} p_{jj}(m-r) + \delta_{ij}$$

$$= \underbrace{\sum_{r=1}^{\infty} s^r f_{ij}(r)}_{F_{ij}(s)} \underbrace{\sum_{m=r}^{\infty} s^{m-r} p_{jj}(m-r)}_{P_{jj}(s)} + \delta_{ij}$$

Recurrence/transience in terms of n-step transition probs

- $\sum_n p_{jj}(n) = \infty \implies j$ is recurrent $\implies \sum_n p_{ij}(n) = \infty$
- $\sum_n p_{jj}(n) < \infty \implies j$ is transient $\implies \sum_n p_{ij}(n) < \infty$

- Proof:
  - First, show that j recurrent $\Leftrightarrow \sum_n p_{jj}(n) = \infty$: from $P_{ii}(s) = 1 + F_{ii}(s)P_{ii}(s)$, we have $P_{jj}(s) = \left[1 - F_{jj}(s)\right]^{-1}$ $(|s| < 1)$
  - As $s \uparrow 1$, $P_{jj}(s) \to \infty$ iff $f_{jj} = F_{jj}(1) = 1$
  - Abel's theorem for power series: $\lim_{s \uparrow 1} P_{jj}(s) = \sum_n p_{jj}(n)$

Classification of states as function of recurrence time

- Def: first visit to state j:

  $T_j = \min\{n \geq 1: X_n = j\}$
  $T_j = \infty$ if chain never visits j

- Note: $\mathbb{P}[T_i = \infty | X_0 = i] > 0$ iff i is transient $\Rightarrow$ $\mathbb{E}[T_i | X_0 = i] = \infty$

- Def: mean recurrence time $\mu_i = \mathbb{E}[T_i | X_0 = i] = \begin{cases} \sum_n n f_{ii}(n) & \text{if } i \text{ is recurrent} \\ \infty & \text{if } i \text{ is transient} \end{cases}$

- Def: a recurrent state i is called **null-recurrent** if $\mu_i = \infty$
  positive-recurrent if $\mu_i < \infty$

- Theorem: a recurrent state i is null $\Leftrightarrow$ $p_{ii}(n) \rightarrow 0$ as $n \rightarrow \infty$

Periodicity

- Def: the period d_i of state i is

$$d_i = GCD\{n: p_{ii}(n) > 0\}$$

- State i is **periodic** if $d_i > 1$, otherwise **aperiodic**
- Def: a state i is **ergodic** if it is positive-recurrent and aperiodic

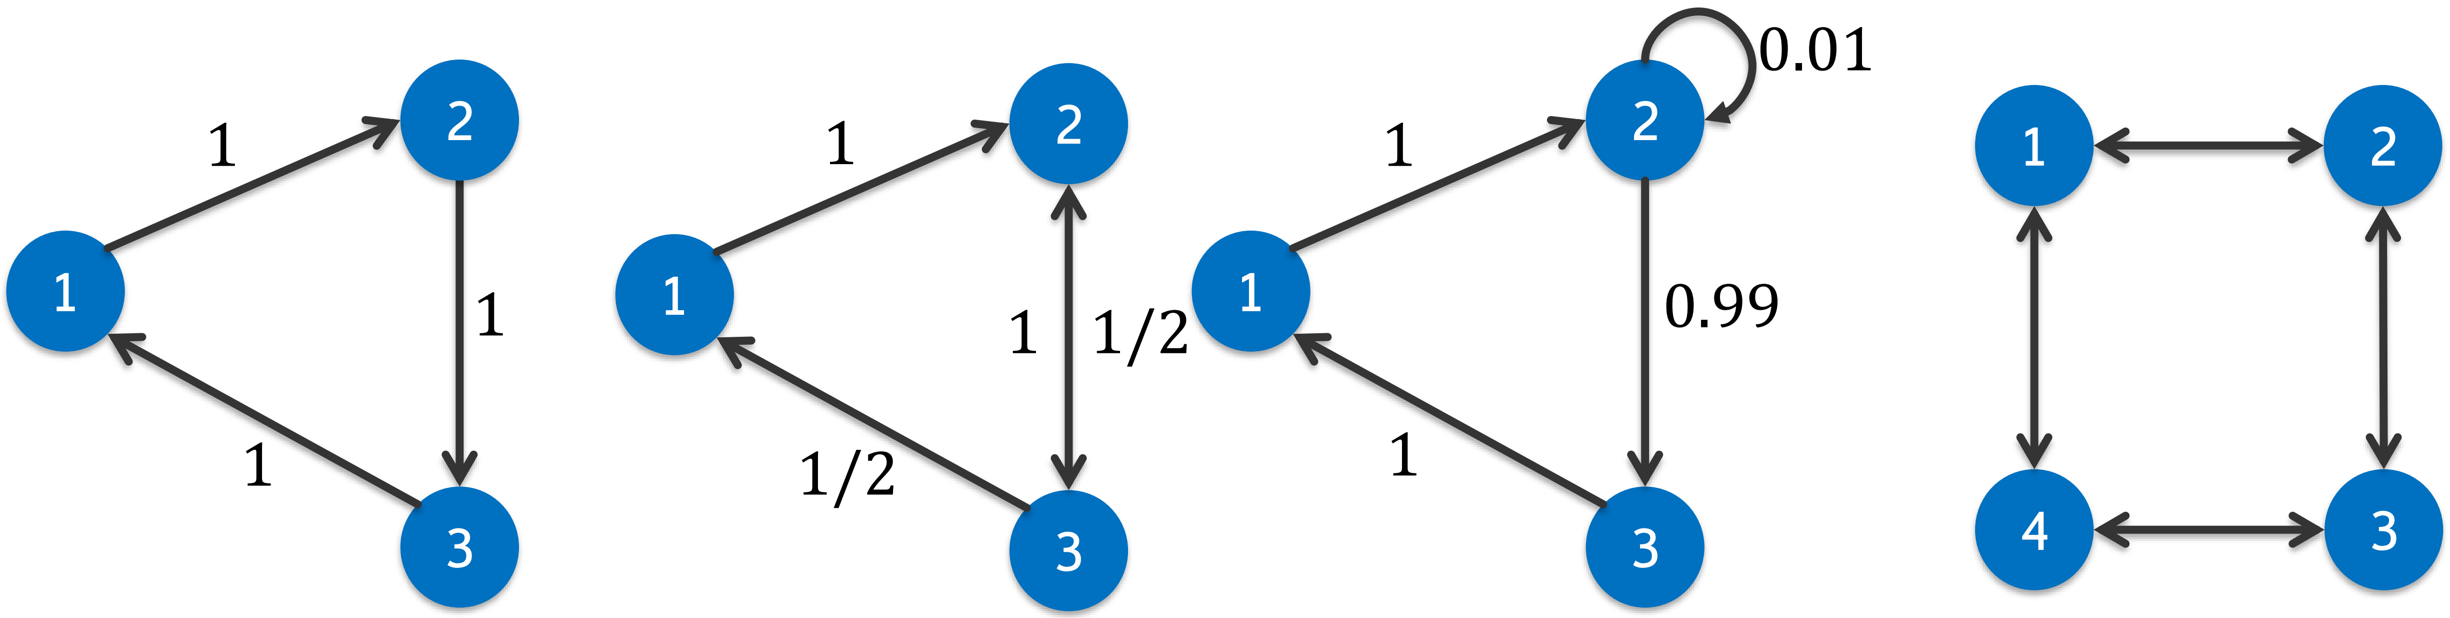

Periodicity

- Def: the period d_i of state i is
$$d_i = GCD\{n: p_{ii}(n) > 0\}$$
- State i is **periodic** if $d_i > 1$, otherwise **aperiodic**
- Def: a state i is **ergodic** if it is positive-recurrent and aperiodic

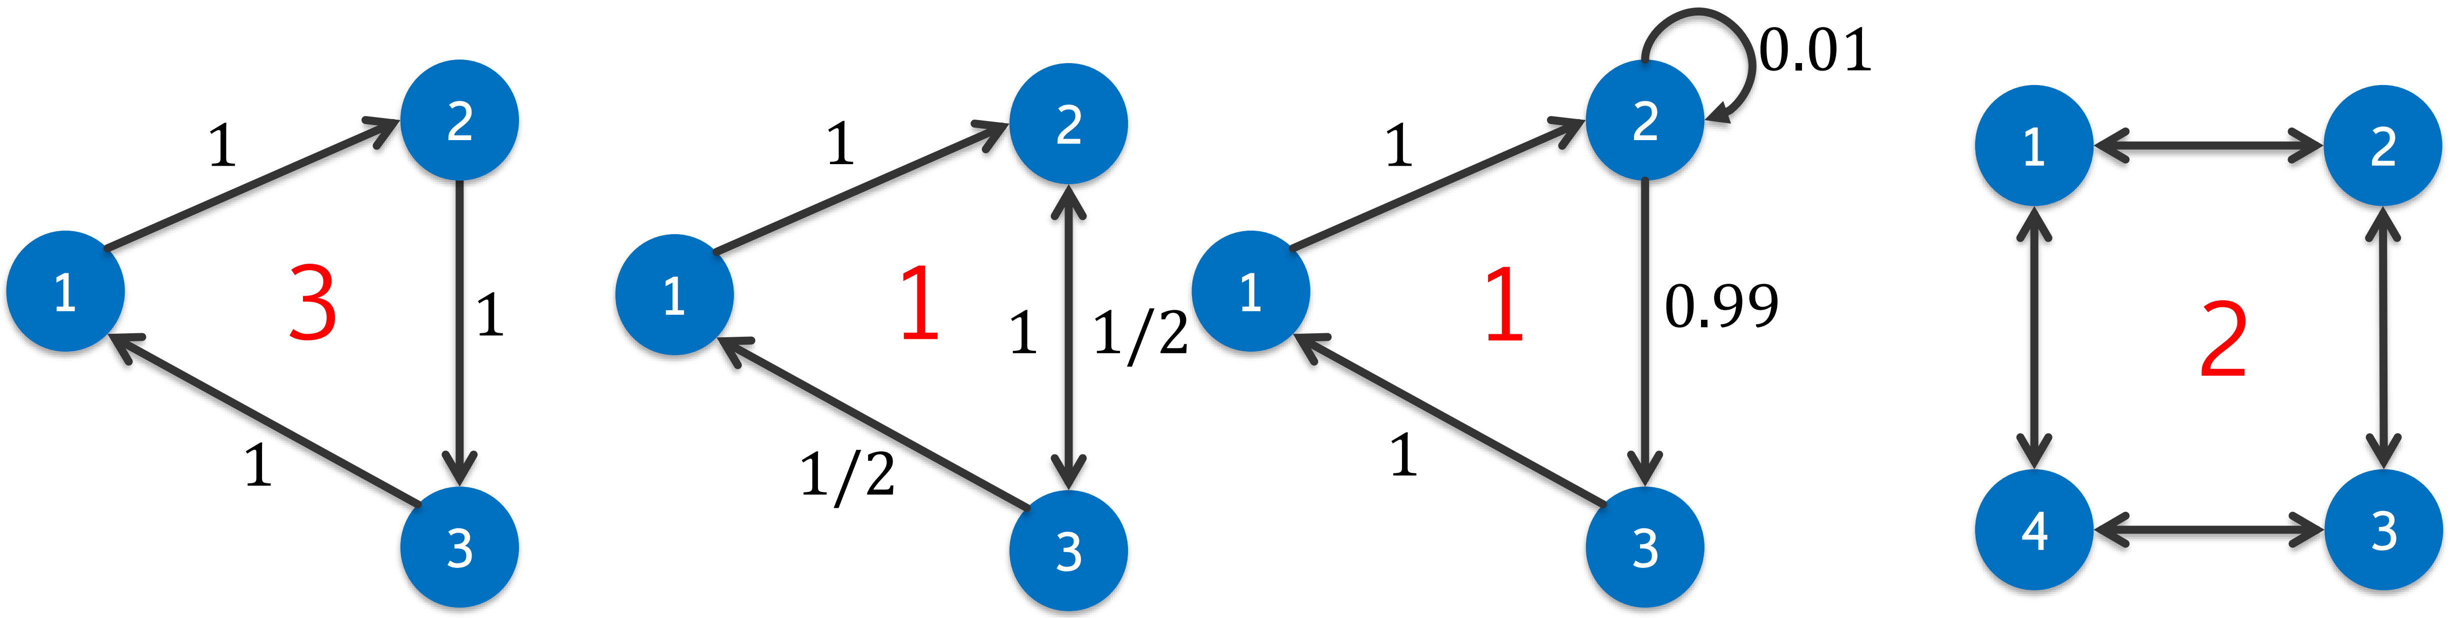

Classification of chains

- Relationships of states to each other
- Def: i **communicates** with j $(i \to j)$ if $p_{ij}(m) > 0$ for some m
- Def: i **intercommunicates** with j $(i \leftrightarrow j)$ if $i \to j$ and $j \to i$
- Note the following properties:
  - (a) For $i \neq j$: $i \to j \Leftrightarrow f_{ij} > 0$
  - (b) For $i = j$: $i \to i$ always because $p_{ii}(0) = 1$
  - (c) $i \leftrightarrow j$ and $j \leftrightarrow k \implies i \leftrightarrow k$

- "$\leftrightarrow$" is an equivalence relationship, partitions $\mathcal{S}$ into equivalence classes of states

Properties of classes of states

- Theorem: if $i \leftrightarrow j$
  - (a) i and j have the same period
  - (b) i is transient $\Leftrightarrow$ j is transient
  - (c) i is null-recurrent $\Leftrightarrow$ j is null-recurrent
  - (d) i is pos-recurrent $\Leftrightarrow$ j is pos-recurrent

- Proof of (b):
  - If $i \leftrightarrow j$, there exist $m, n \geq 0$ such that
    $$\alpha = p_{ij}(m)p_{ji}(n) > 0$$
  - By Chapman-Kolmogorov (using only one term from the sum),
    $p_{ii}(m + r + n) \geq p_{ij}(m)p_{jj}(r)p_{ji}(n) = \alpha p_{jj}(r)$ (for any $r \geq 0$)
  - Summing over r:
    $$\sum_r p_{ii}(r) < \infty \quad \Rightarrow \sum_r p_{jj}(r) < \infty$$
  - In other words: i transient implies j transient

Close and irreducible sets of states

- Def: a set $C \subset \mathcal{S}$ of states is called
  - (a) closed/absorbing if $p_{ij} = 0$ for all $i \in C, j \notin C$ ("we can never leave")
  - (b) irreducible if $i \leftrightarrow j$ for all $i, j \in C$ ("we can move freely inside C")

- Because of the preceding theorem: all states in an irreducible set share the properties of periodic/aperiodic and of transient/null-recurrent/pos-recurrent

- If the whole $\mathcal{S}$ is irreducible, we say that the chain has the shared properties

Example

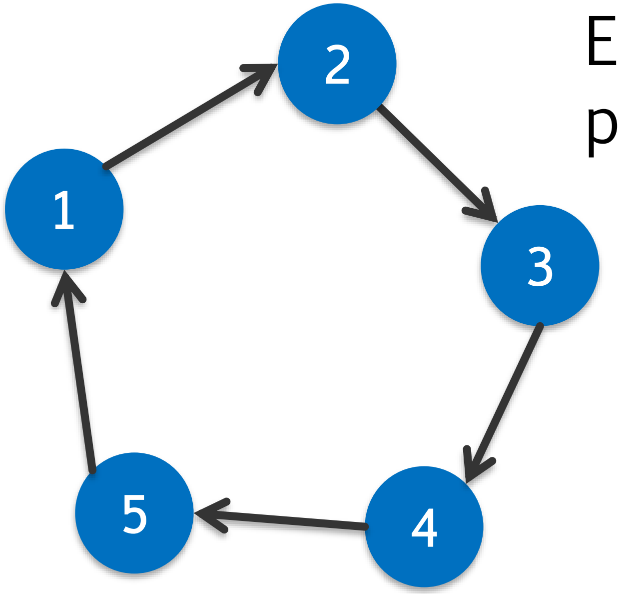

Entire $\mathcal{S}$ is pos-recurrent, periodic 5

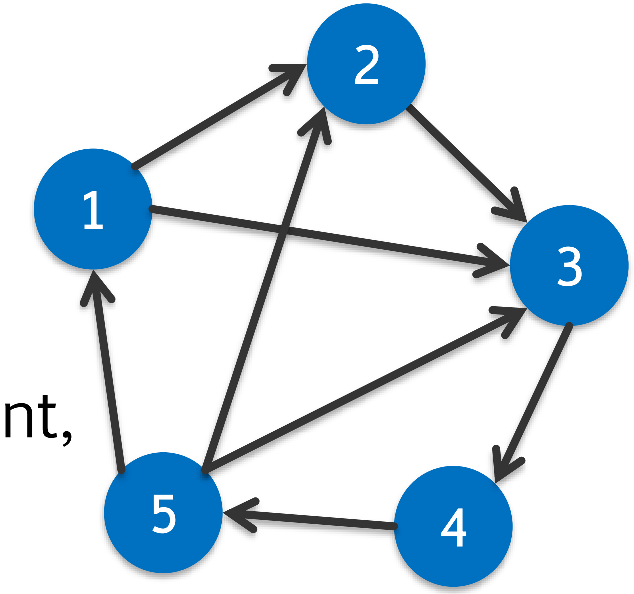

Entire $\mathcal{S}$ is pos-recurrent, aperiodic

Example: what are the irreducible classes?

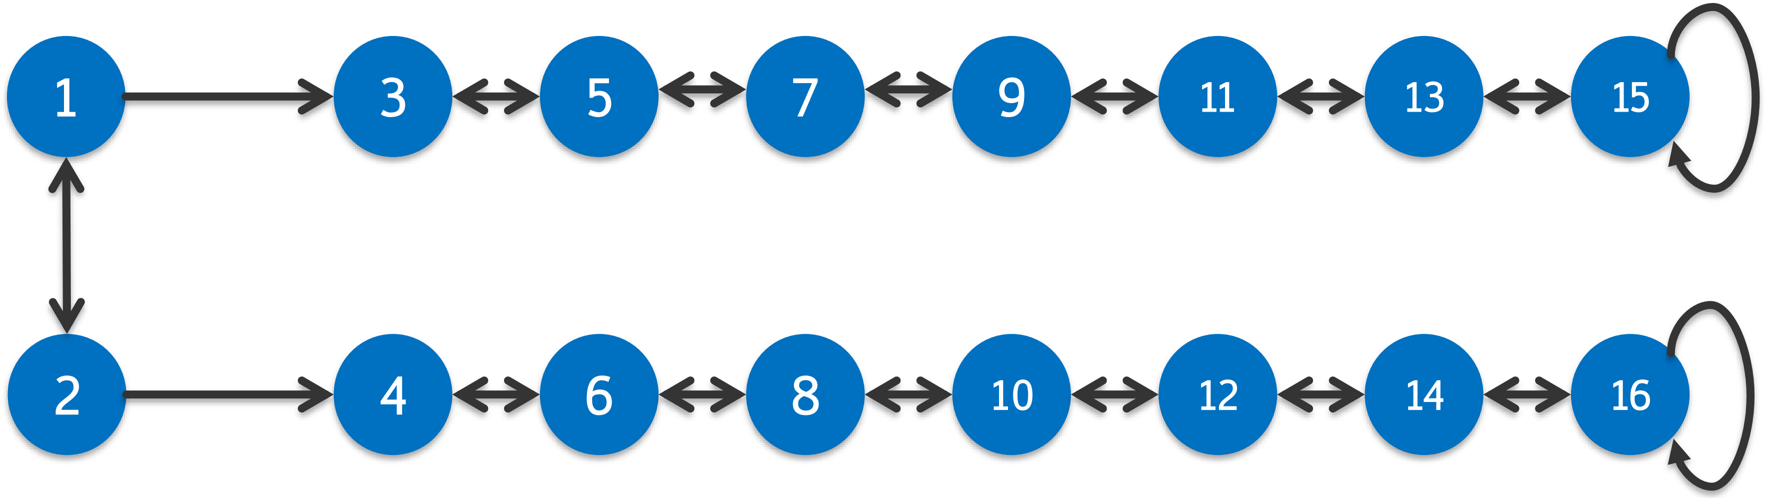

Example

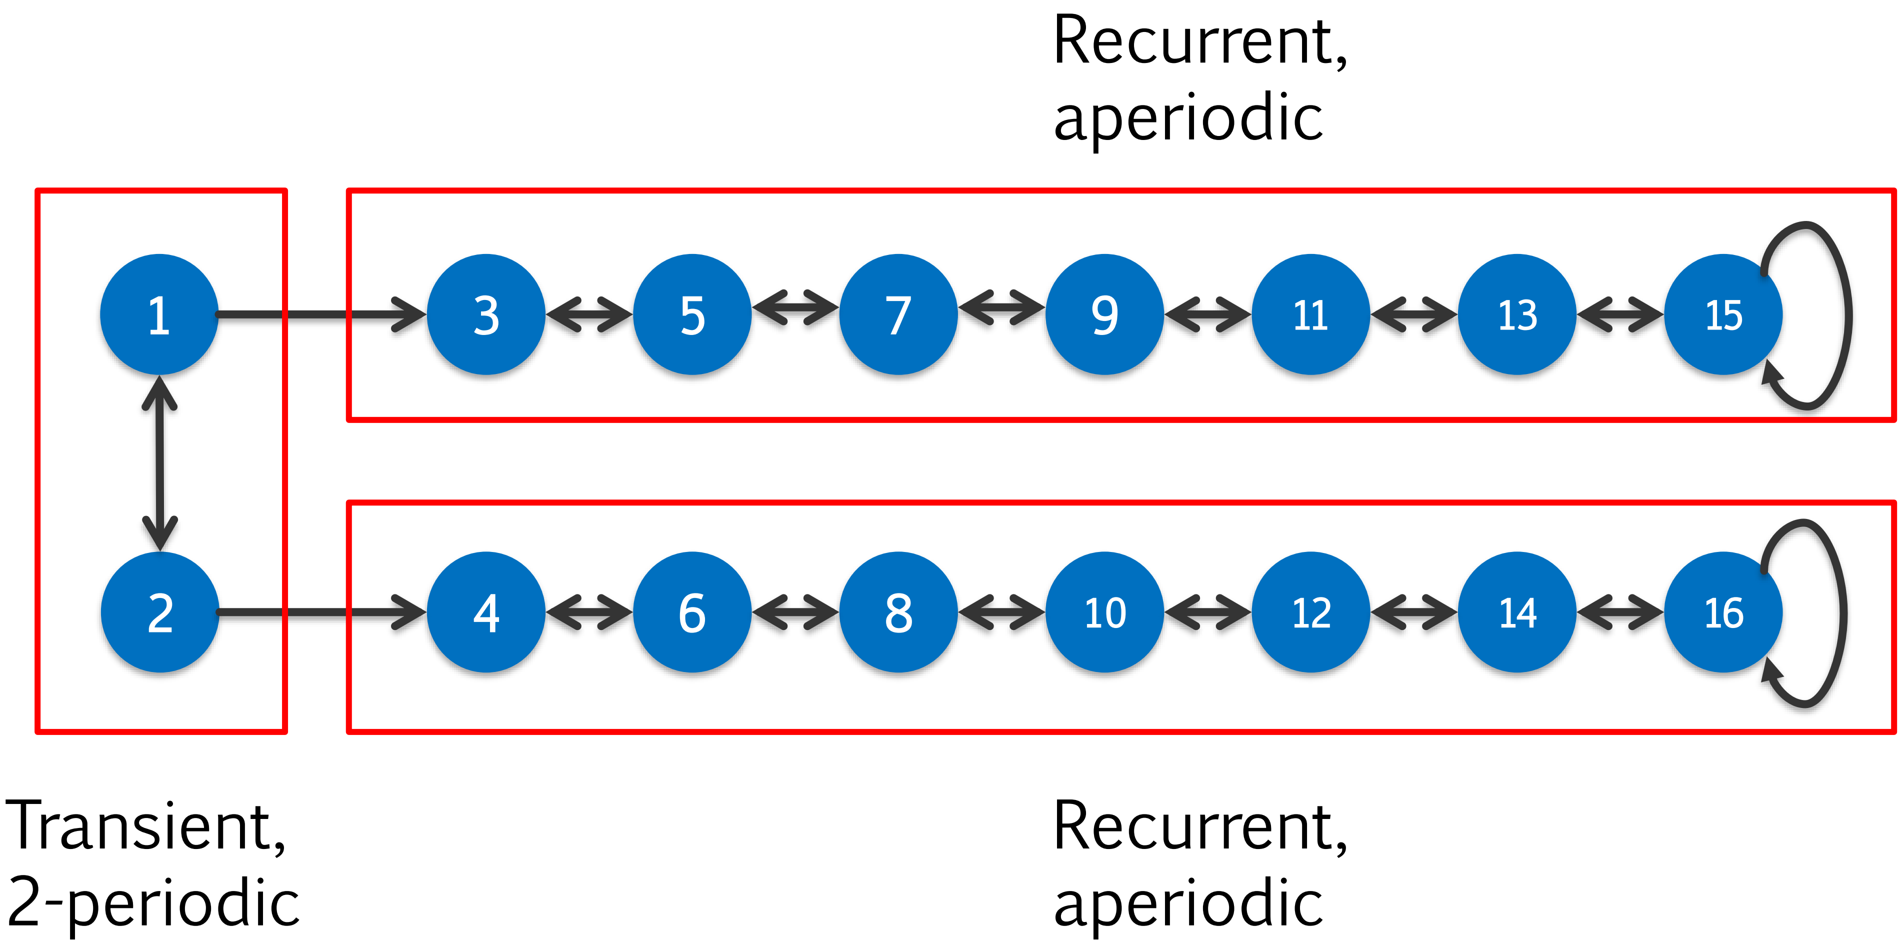
Recurrent,
aperiodic
1
2
3
5
7
9
11
13
15
4
6
8
10
12
14
16
Transient,
2-periodic
Recurrent,
aperiodic

Examples: random walks

- Simple random walk on $\mathbb{Z}$:

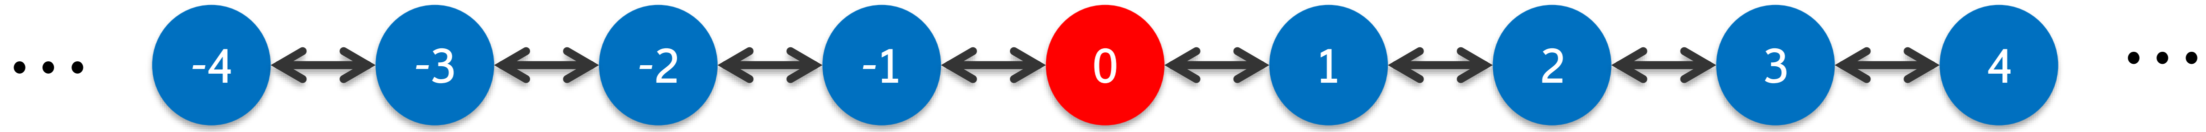

  - For $p = q = \frac{1}{2}$ (unbiased): null-recurrent (even though the mode of $\mathbb{P}[X_n = i]$ is always at 0!)
  - For $p \neq q$: transient
- Simple random walk with barrier at 0:

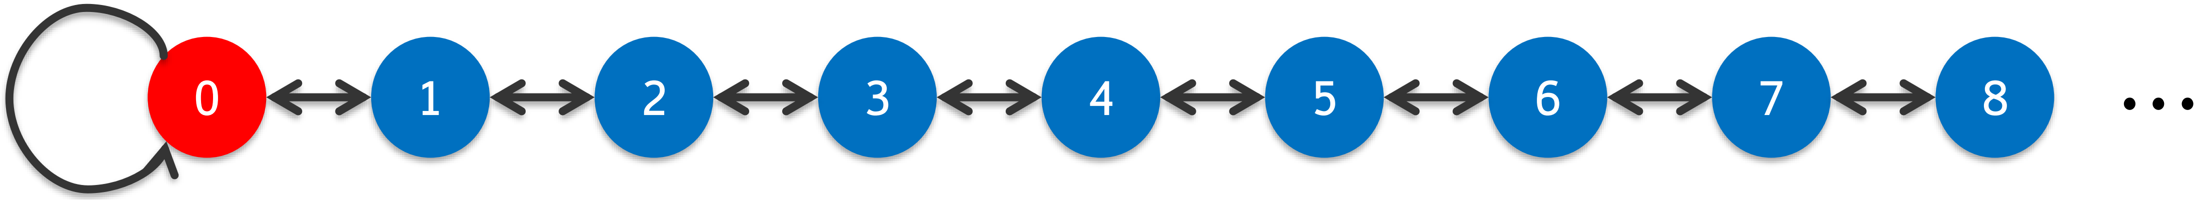

  - For $p = q = \frac{1}{2}$ (unbiased): null-recurrent
  - For $p > q$: transient
  - For $p < q$: positive-recurrent

Examples: more random walks

- Unbiased on $\mathbb{Z}^2$: null-recurrent

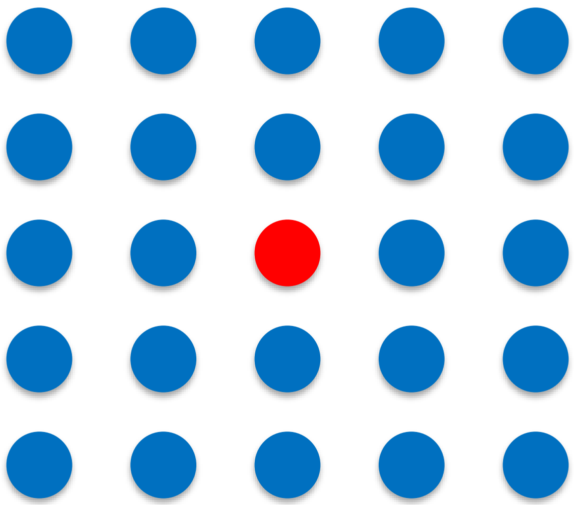

- Unbiased on $\mathbb{Z}^3$: transient

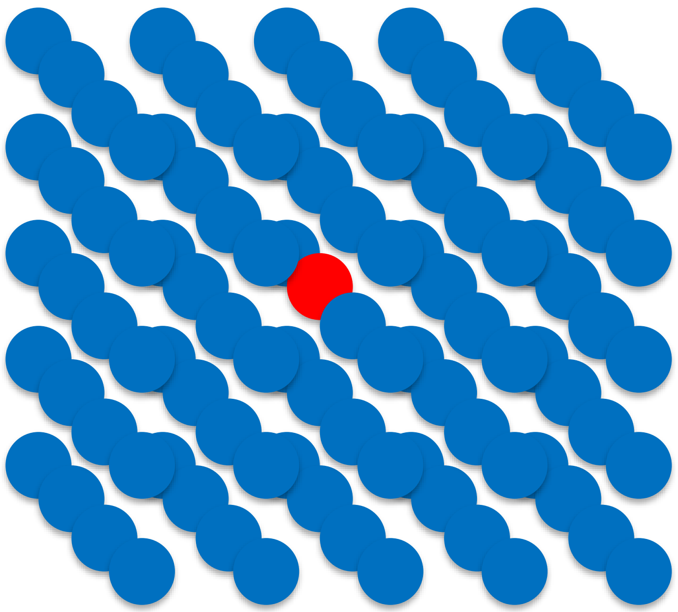

Decomposition theorem

- The state space $\mathcal{S}$ can be partitioned uniquely as follows:
$$S = T \cup C_1 \cup C_2 \cup \cdots$$
where T is the set of transient states, and C_k are irreducible closed sets of recurrent states
- Note: we do not assert that T is irreducible!
- Proof:
  - Let $C_1, C_2, \ldots$ be the recurrent equivalence classes of $\leftrightarrow$
  - We only need to show that they are all closed
  - Suppose the contrary: that there is a $i \in C_r, j \notin C_r$ such that $p_{ij} > 0$
  - Because $j \nrightarrow i$ (otherwise they intercommunicate and are therefore in the same class),
$$\mathbb{P}[X_n \neq i \text{ for all } n \geq 1 | X_0 = i] \geq \mathbb{P}[X_1 = j | X_0 = i] > 0$$
which contradicts the assumption that i is recurrent
- So there are two possibilities:
  - Either the chain stays in T forever…
  - Or it is (eventually or from the beginning) in one of the C_r's and stays there forever
- For finite $\mathcal{S}$, only the second scenario is possible

Decomposition for finite $\mathcal{S}$

- At least one state is recurrent
- There are no null-recurrent states
- A class of states is positive-recurrent $\Leftrightarrow$ the class is absorbing
- A class of states is transient $\Leftrightarrow$ the class is non-absorbing
- There exists at least one class of absorbing states
- If the chain is irreducible, it is positive-recurrent

Limit behavior: ergodic theorem

- Note that periodicity is problematic:
  - Example: $\mathcal{S} = \{1,2\}$, $p_{12} = p_{21} = 1$

    Then $p_{11}(n) = p_{22}(n) = \begin{cases} 0 & n \text{ odd} \\ 1 & n \text{ even} \end{cases}$

    This does not converge!

- Theorem: for an irreducible aperiodic chain (not necessarily finite),

$$p_{ij}(n) \to \frac{1}{\mu_j} \qquad \text{for all } i \text{ and } j$$

- Interpretation: the long-term probability of visiting state j “from anywhere” converges to the rate of revisiting that state
- Can we compute that rate? Yes!

Stationary distribution π

- For an irreducible aperiodic finite Markov chain, the transition matrix $P \in \mathbb{R}^{d \times d}$:
  - Has a single dominant eigenvalue 1; its associated eigenvector is (proportional to) the stationary distribution π
  - All other eigenvalues $\lambda_2, \lambda_3, \dots, \lambda_d$ are $|.| < 1$

- Computing the unique stationary distribution: solve $\pi = \pi P$ (with $\pi 1^T = 1$)

- For large systems, can be computed by the power method:
  - Take a vector $e = [1\ 1\ 1 \dots 1]/d$ (for example), repeatedly multiply with P: $eP^n \to \pi$
  - Recall simple random walk:
    $P_{10} = \begin{bmatrix} 0.600 & 0.400 \\ 0.600 & 0.400 \end{bmatrix}$ → $\pi = [3/5, 2/5]$

Example: random walks on general graphs

- Undirected graph $G(V, E)$
  - Assume connected (otherwise assume G is the GC for the actual network)
  - Random Walk = Markov chain whose state space is $\mathcal{S} = V$:
    - Node at time n: $X_n \in V$
    - At each time step, go to a neighbor of X_n uniformly at random → X_{n+1}
- Random walk on the graph: random memoryless exploration of the graph

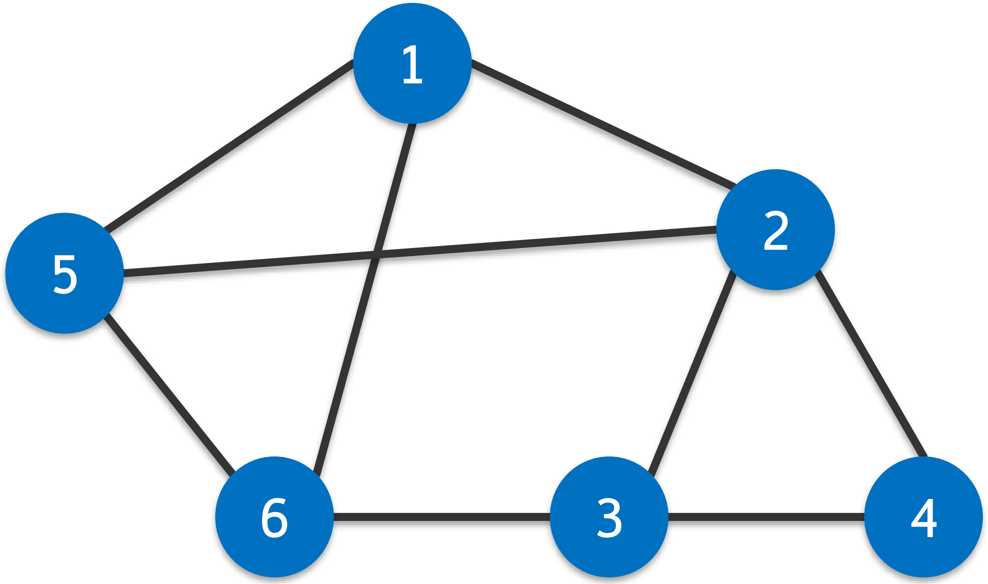

Random walk as Markov chain

- Transition matrix P:
  - $P = \begin{cases} p_{ij} = 1/d_i & (i,j) \in E \\ 0 & \text{otherwise} \end{cases}$
- If $G(V,E)$ is undirected, connected and non-bipartite, then $\{X_n\}$ is an ergodic (irreducible, aperiodic) Markov chain
- Ergodicity:
  - Stationary distribution π
  - $p_{ij}(n) = \mathbb{P}(\text{at vertex } j \text{ after } n \text{ steps} | \text{starting at vertex } i)$
    - $\boldsymbol{p_{ij}(t) \to \pi_j}$ **for all** $\boldsymbol{i,j \in V}$
  - RW "forgets" starting point i

Stationary distribution π

- Lemma:
  - $\pi \propto [d_1, d_2, \dots, d_n]$
- Proof:
  - Def of stationary distribution: $\pi = \pi P$
  - $[d_1, d_2, d_3, \dots, d_n]P = x$
  - $x_j = \sum_i d_i\, p_{ij} = \sum_i 1_{\{(i,j)\in E\}} = d_j$
  - $[d_1, \dots, d_n]$ is a left-eigenvector with eigenvalue = 1 → $\pi \propto [d_1, \dots, d_n]$

- Intuition:
  - Random walk “sees” uniformly random edges
  - The nodes visited by RW therefore biased by # of edges = degree
- Note: this is specific to undirected graphs, no such compact characterization of π exists for directed

Summary

- Markov chains: a very versatile and rich model, only one small step beyond independence:
  - Conditional on the current value, the past and the future are independent
- Classification of states:

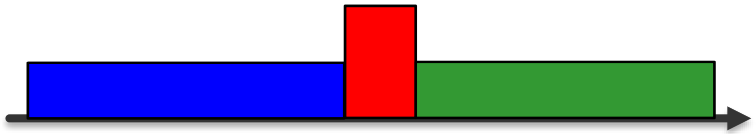

  - Transient: eventually abandoned
  - Recurrent: always coming back
    - Positive-recurrent: always coming back in finite expected time
    - Null-recurrent: always coming back, but in infinite expected time
  - Null-recurrent is “on the edge”, and only happens in infinite state spaces
- Ergodicity: the chain forgets where it came from, and visits all states at some long-term average rate given by the stationary distribution
  - Time averages converge to ensemble averages
- It is hard to overstate the range of applications of Markov chains in machine learning, communications, queueing theory, supply chains, finance,…
- We have only seen discrete-space, discrete-time: much more to learn! ;)