ergodic
MC defined by transition matrix P on Ω with stationary dist. π.

<u>Conductance</u> measures edge expansion of (Ω, P)
(called bottleneck ratio in [LPW])

For a subset $S \subset \Omega$,

$$\Phi(S) = \Pr(X_{t+1} \notin S \mid X_t \in S, X_t \sim \pi)$$

Prob. escape S in one step.

in S at time t with dist. π

$$\Phi(S) = \frac{\sum_{y \in S, z \notin S} \pi(y) P(y,z)}{\pi(S)}$$

Let $\Phi_* = \min_{\substack{S \subset \Omega: \\ \pi(S) \le 1/2}} \Phi(S)$

<u>Theorem:</u> $\frac{1}{4\Phi_*} \le T_{mix} \le O\left(\frac{1}{\Phi_*^2} \log\left(\frac{1}{\pi_{min}}\right)\right)$

$\pi_{min} = \min_{\sigma \in \Omega} \pi(\sigma)$.

Example:

Consider random walk (RW) on a d-regular connected graph?

Let $G=(V,E)$, & $\pi = \text{uniform}(V)$.

Note, for $(v,w)\in E$, $P(v,w)=\frac{1}{d}$ since d-regular.

let's say $P(v,w)=\frac{1}{2d}$ & $P(v,v)=1/2$

So it's clearly ergodic.

Consider $S\subset V$, where $|S|\le\frac{|V|}{2}$,

$$\Phi(S)=\frac{\sum_{y\in S, z\notin S}\pi(y)P(y,z)}{\pi(S)}=\frac{\sum_{y,z}\frac{1}{\cancel{n}}\frac{1}{d}}{\frac{|S|}{\cancel{n}}}$$

$$=\frac{1}{d}\frac{|E(S,\bar{S})|}{|S|}=\frac{\#\text{edges } S\to\bar{S}}{d|S|}$$

So if good edge expansion for every cut then ~~rapid~~ fast mixing

& if there's a bottleneck then slow mixing.

Proof of lower bound that $T_{mix} \geq \frac{1}{4\Phi^*}$

See also [LPW, Theorem 7.4]

Let $X_0 \sim \pi$. Hence, for all $t \geq 0$, $X_t \sim \pi$,
since π is stationary.

for any $S \subset \Omega$, any $t \geq 1$,

$$\Pr_\pi(X_0 \in S, X_t \notin S)$$

$$\leq \sum_{r=1}^{t} \Pr_\pi(X_{r-1} \in S, X_r \notin S)$$

$$= t \Pr_\pi(X_0 \in S, X_1 \notin S)$$

$$= t \sum_{y \in S, z \notin S} \pi(y) P(y,z)$$

Note,

$$\Pr_\pi(X_t \notin S \mid X_0 \in S) = \frac{\Pr_\pi(X_0 \in S, X_t \notin S)}{\Pr_\pi(X_0 \in S)}$$

$$\leq \frac{t \sum_{y \in S, z \notin S} \pi(y) P(y,z)}{\sum_{y \in S} \pi(y)} = t\Phi(S).$$

We have: $\Pr_{\pi}(X_t \notin S \mid X_0 \in S) \leq t\Phi(S)$

So if $X_0 \in S$ & $X_0 \sim \pi$ then $\Pr(X_t \notin S) \leq t\Phi(S)$,
and hence there must be at least one $\sigma \in S$
where: if $Y_0 = \sigma$, then $\Pr(Y_t \notin S) \leq t\Phi(S)$
(because this is averaging over σ wrt π)
So at least one σ achieves it.

We have for $Y_0 = \sigma$, $\Pr(Y_t \notin S) \leq t\Phi(S)$.

in other words, $P^t(\sigma, \bar{S}) \leq t\Phi(S)$

and hence, $P^t(\sigma, S) \geq 1 - t\Phi(S)$.

Look at variation distance of Y_t & π (for $Y_0 = \sigma$):

$$d_{TV}(Y_t, \pi) = \frac{1}{2} \sum_{\tau \in \Omega} \left| \Pr(Y_t = \tau) - \pi(\tau) \right|$$

$$= \max_{A \subset \Omega} \Pr(Y_t \in A) - \pi(A)$$

$$\geq \Pr(Y_t \in S) - \pi(S)$$

$$\geq 1 - t\Phi(S) - \pi(S).$$

If $\pi(S) \leq \frac{1}{2}$ & $t < \frac{1}{4\Phi(S)}$

then: $d_{TV}(Y_t, \pi) > \frac{1}{4}$

So $T_{mix} \geq \max\limits_{S} \frac{1}{4\Phi(S)} = \frac{1}{4} \frac{1}{\min \Phi(S)} = \frac{1}{4\Phi_*}$ $\blacksquare$

Consider random walk (RW) on the following graph known as dumbbell graph:

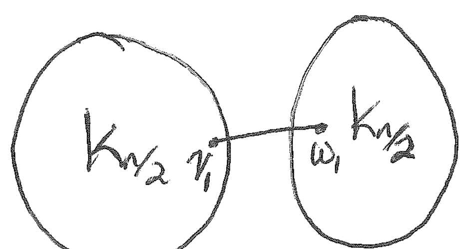

S T

$$\Phi(S) = \frac{\pi(v_1)P(v_1,w_1)}{\pi(S)} = \frac{\frac{1}{n} \cdot \frac{1}{n/2}}{\frac{1}{2}} = \frac{4}{n^2}$$

Thus, $T_{mix} \geq \frac{n^2}{16} = \Omega(n^2)$

Can show coupling argument to prove:

$T_{mix} = O(n^2)$

therefore, $T_{mix} = \Theta(n^2)$. $\square$

Why only S where $\pi(S) \le \frac{1}{2}$?

If S is huge (almost all of Ω) then few transitions out since $|\bar{S}| \ll |S|$ but this is not a real bottleneck.

Symmetrized conductance:

$$\hat{\Phi}(S) = \frac{\sum_{y \in S, z \in \bar{S}} \pi(y) P(y,z)}{\pi(S)\,\pi(\bar{S})}$$

$$\& \quad \hat{\Phi}_* = \min_{S \subset \Omega} \hat{\Phi}(S)$$

all S regardless of $\pi(S)$.

For reversible chains, $\hat{\Phi}(S) = \hat{\Phi}(\bar{S})$.

if $\pi(S) \le \pi(\bar{S})$ then $\pi(S) \le \frac{1}{2}$, $\pi(\bar{S}) \ge \frac{1}{2}$

thus

$$\Phi(S) \le \hat{\Phi}(S) \le 2\Phi(S)$$

So

$$\Phi_* \le \hat{\Phi}_* \le 2\Phi_*.$$

Upper bound: $T_{mix} = O(\Phi_*^{-2} \log(1/\pi_{min}))$

we'll use spectral techniques to prove later.

Let's see how to bound Φ_*

<u>idea:</u> min cut size = max-flow value

though we're looking at normalized cut size

we'll define a "good" multicommodity flow.

<u>Canonical paths technique:</u>

For every ordered pair $x, y \in \Omega$,

define a path γ_{xy} along transitions, i.e.,

$$\gamma_{xy} = (z_0, z_1, \ldots, z_\ell)$$

where $z_0 = x$, $z_\ell = y$

and forall $0 \le i < \ell$, $P(z_i, z_{i+1}) > 0$.

for transition $\vec{uv}$ (i.e., $u, v \in \Omega$, $P(u,v) > 0$)

let $\mathcal{P}_{uv} = \{(x,y) \in \Omega : \gamma_{xy} \ni \vec{uv}\}$ = set of paths that go through $\vec{uv}$

Congestion:

$$\rho(\vec{uv}) = \frac{1}{\underbrace{\pi(u)P(u,v)}_{\text{capacity of } \vec{uv}}} \sum_{(x,y)\in\mathcal{P}_{uv}} \overbrace{\pi(x)\pi(y)}^{\pi(x)\pi(y) \text{ flow } x \rightsquigarrow y}$$

Suppose $\pi = \text{uniform}(\Omega)$ & $P(u,v) = \frac{1}{2d}$

then $\rho(\vec{uv}) \leq |\Omega| \times (2d) \times |\mathcal{P}_{uv}| \times \frac{1}{|\Omega|^2} = O\left(\frac{d|\mathcal{P}_{uv}|}{|\Omega|}\right)$

Let $\rho = \max\limits_{\vec{uv}:P(u,v)>0} \rho(\vec{uv})$

then $$T_{mix} = O\left(\rho^2 \log\left(\frac{1}{\pi_{min}}\right)\right)$$

Can improve to: $$T_{mix} = O\left(\rho\, \ell_{max} \log\left(\frac{1}{\pi_{min}}\right)\right)$$

where $\ell_{max} = \max\limits_{x,y\in\Omega} |\gamma_{xy}|$ = max path length

(usually only gives poly(n) bounds, not optimal)

Proof: Consider $S \subset \Omega$ where $\pi(S) \le 1/2$.

Suppose $\pi = \text{uniform}(\Omega)$, & $P(u,v) = \frac{1}{2d}$

Then $$\Phi(S) = \frac{|E(S,\bar{S})|}{d|S|}$$

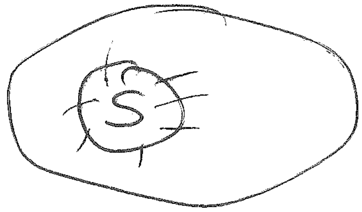

There are $|S||\bar{S}|$ paths from $S \rightsquigarrow \bar{S}$,

each path crosses $S \longrightarrow \bar{S}$ ≥ 1 time,

& each $e \in E(S,\bar{S})$ is crossed

$\le |\mathcal{P}(e)|$ times.

Note, $$|\mathcal{P}(e)| \le \frac{\rho|\Omega|}{d}$$

since $|S| \le \frac{|\Omega|}{2}$ & $|\bar{S}| \ge \frac{|\Omega|}{2}$

Hence, $$|E(S,\bar{S})| \ge \frac{|S||\bar{S}|}{\rho|\Omega|/d} \ge \frac{|S|d}{2\rho}$$

& thus, $$\Phi(S) \ge \frac{1}{2\rho} \Rightarrow T_{mix} = O\left(\rho^2 \log\left(\frac{1}{\pi_{min}}\right)\right)$$

since $T_{mix} = O(\Phi_*^{-2} \log(1/\pi_{min}))$. ∎

Whats an example of canonical paths?

Consider RW on n-dimensional hypercube $\{0,1\}^n$.

For $I, F \in \Omega$, we'll define $\gamma_{IF} = (z_0, z_1, \ldots, z_l)$
where $l \leq n$.

Let $z_0 = I$
For $i = 1 \to n$:
for all $j \neq i$, $z_i(j) = z_{i-1}(j)$
& $z_i(i) = F(i)$
(so set i^{th} bit to F & rest same)

Note, $z_n = F$, & z_i, z_{i-1} differ in ≤ 1 bit
so $P(z_{i-1}, z_i) = \frac{1}{2n}$
(& skip if $z_{i-1} = z_i$)

How to bound $\rho(\overrightarrow{uv})$?

— Suppose $\overrightarrow{uv}$ flips bit i.

We'll define a map $\eta_{\overrightarrow{uv}} : \mathcal{P}_{uv} \to \Omega$

which will be injective $\left(\begin{array}{l} \text{for } \tau \in \Omega, \\ |\eta^{-1}(\tau)| \le 1 \end{array} \right)$

& hence $|\mathcal{P}_{uv}| \le |\Omega|$

& thus $\rho(\overrightarrow{uv}) = O(n)$

think of η as "encoding": from η & $\overrightarrow{uv}$

we can uniquely determine $(I,F) \in \mathcal{P}_{uv}$.

Let $E = \eta_{\overrightarrow{uv}}(I,F)$

j^{th} bit of $\eta_{uv}(I,F)$

& $E(j) = \begin{cases} F(j) & \text{for } j > i \\ I(j) & \text{for } j \le i \end{cases}$ $\left(\begin{array}{c} \text{bit } i \\ \text{is flipped} \\ \text{by } \overrightarrow{uv} \end{array} \right)$

Given $E = \eta_{uv}(I,F)$ & $\overrightarrow{uv}$

then we can determine I, F:

I agrees with $\begin{cases} E & \text{on bits} \le i \\ u & \text{on bits} > i \end{cases}$

F agrees with $\begin{cases} v & \text{on bits} \le i \\ E & \text{on bits} > i \end{cases}$

Therefore, $|\eta^{-1}(E)| \le 1$ so η is injective

& hence $\rho = O(n)$.

Since $\pi_{min} = 2^{-n}$ so $\log\left(\frac{1}{\pi_{min}}\right) = O(n)$

then we get $T_{mix} = O(n^3)$.

(Note, we already know $T_{mix} = O(n \log n)$ from a coupling argument.)

Next class: Use this approach for a random matching (of any G).