Distributed Testing of Conductance

Hendrik Fichtenberger, Yadu Vasudev

August 31, 2018

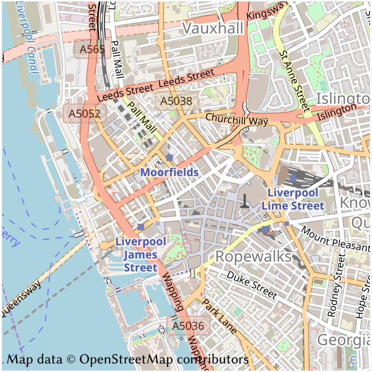

classic / global algorithm

see everything

complexity $\Omega(n)$

output solution

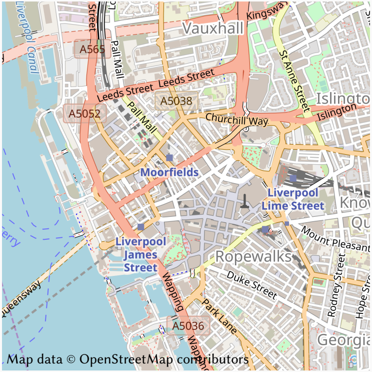

classic / global algorithm

see everything

complexity $\Omega(n)$

output solution

sublinear algorithm

see only small parts

complexity $o(n)$

estimate solution's value

Property Testing

Given a graph $G = (V, E)$, decide with prob. $\geq 2/3$

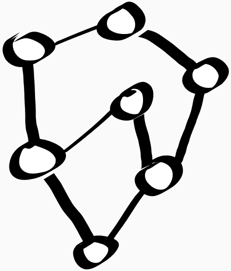

C_3-free
accept

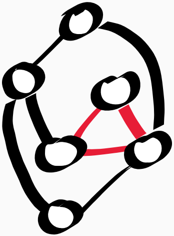

ϵ-close to C_3-free
don't care

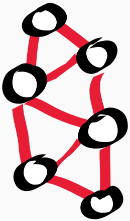

ϵ-far from C_3-free
reject

distance "ϵ-far from" = need to modify more than $\epsilon|E|$ edges

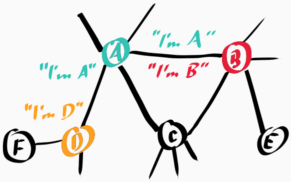

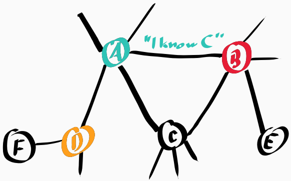

- input graph $G = (V, E)$
- each vertex has id $\in \mathsf{poly}(n)$
- processor on each vertex $v \in V$

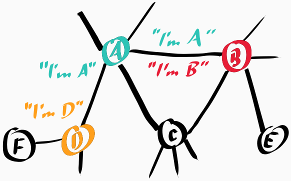

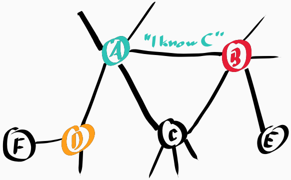

- input graph $G = (V, E)$
- each vertex has id $\in \text{poly}(n)$
- processor on each vertex $v \in V$
- synchronized rounds
  one round for vertex $v \in V$:
  1. unlimited local computation
  2. $\forall u \in \Gamma(v)$: send $O(\log n)$ bits to u
  3. $\forall u \in \Gamma(v)$: receive message from u

Distributed Property Testing in the CONGEST model

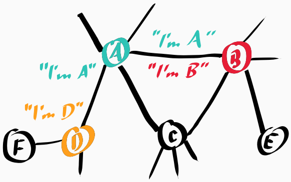

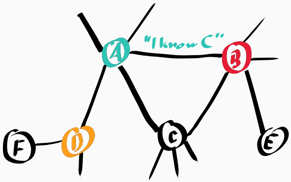

- input graph $G = (V, E)$
- each vertex has id $\in \text{poly}(n)$
- processor on each vertex $v \in V$
- synchronized rounds
  one round for vertex $v \in V$:
  1. unlimited local computation
  2. $\forall u \in \Gamma(v)$: send $O(\log n)$ bits to u
  3. $\forall u \in \Gamma(v)$: receive message from u
- after last round
  - every vertex votes *accept* or *reject*
  - tester rejects iff at least one vertex votes *reject*

Distributed Property Testing in the CONGEST model

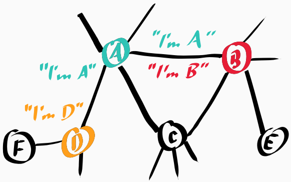

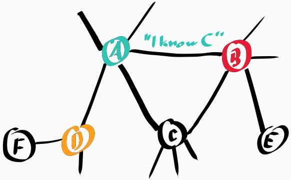

- input graph $G = (V, E)$
- each vertex has id $\in \text{poly}(n)$
- processor on each vertex $v \in V$
- synchronized rounds
  one round for vertex $v \in V$:
  1. unlimited local computation
  2. $\forall u \in \Gamma(v)$: send $O(\log n)$ bits to u
  3. $\forall u \in \Gamma(v)$: receive message from u
- after last round
  - every vertex votes *accept* or *reject*
  - tester rejects iff at least one vertex votes *reject*
- complexity measure: #rounds

Conductance In Pictures

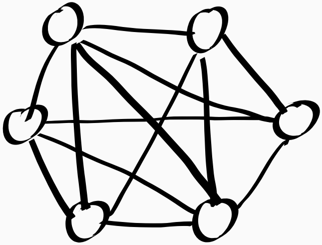

Conductance In Pictures

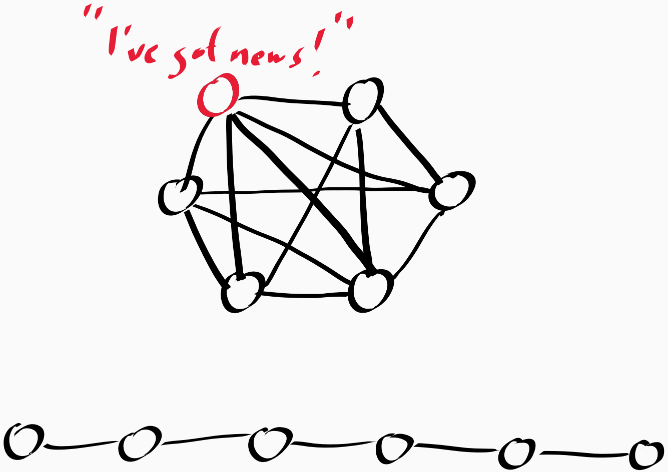

Conductance In Pictures

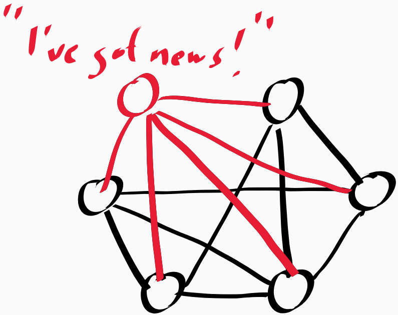

Conductance In Pictures

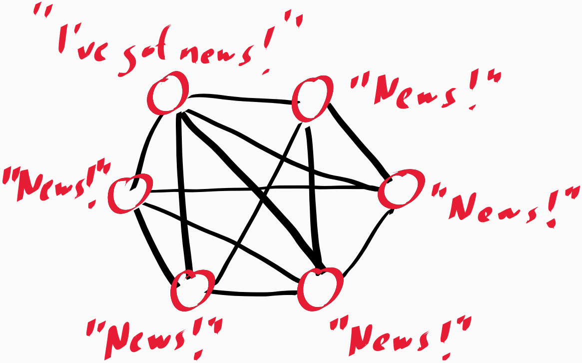

Conductance In Pictures

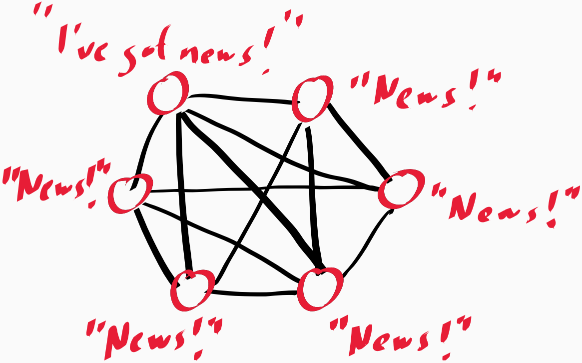

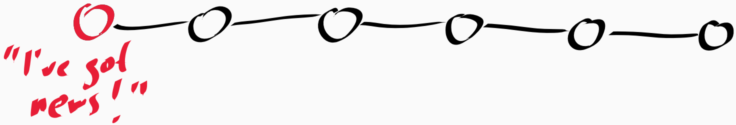

Conductance In Pictures

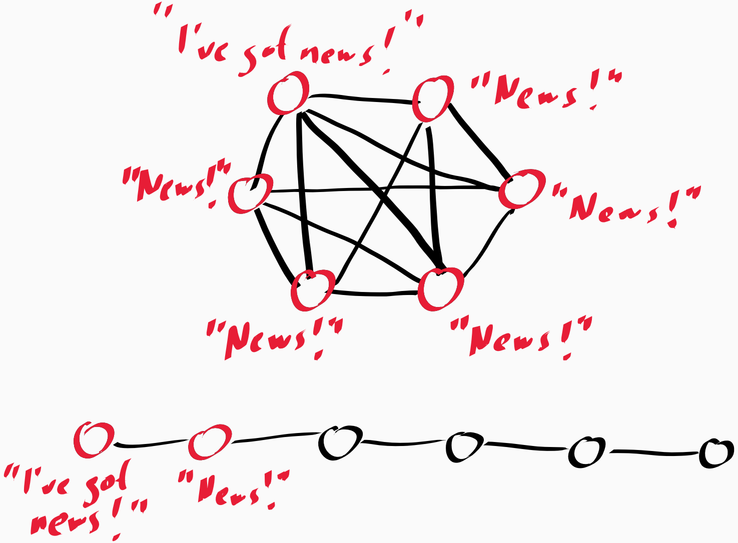

Conductance In Pictures

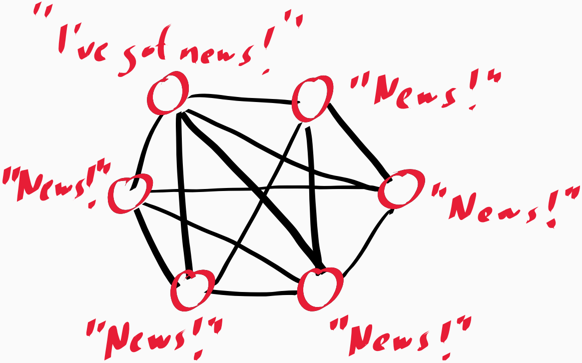

Conductance In Pictures

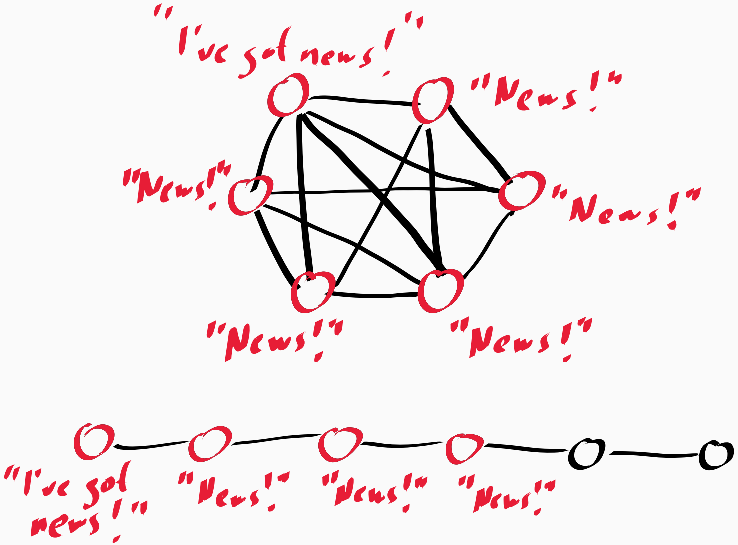

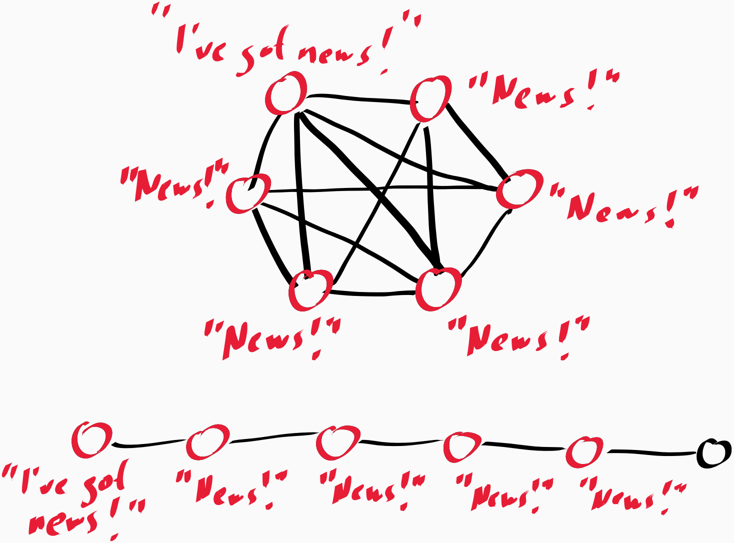
"I've got news!"
"News!"
"News!"
"News!"
"News!"
"News!"
"I've got news!"
"News!"
"News!"
"News!"
"News!"

Conductance In Pictures

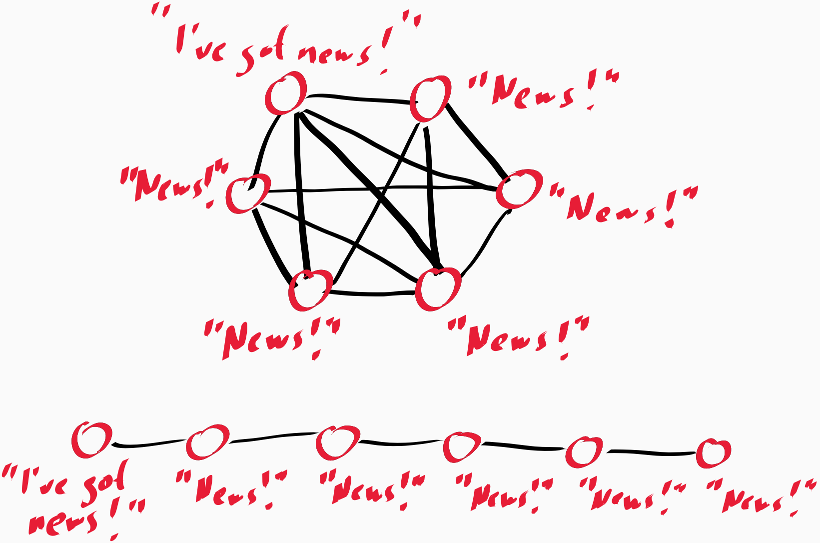

Conductance More Formally

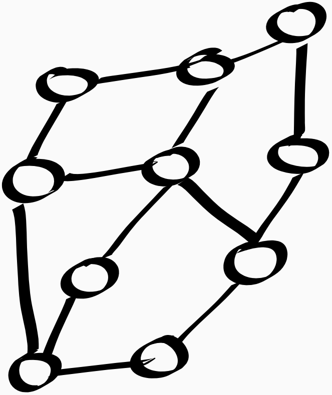

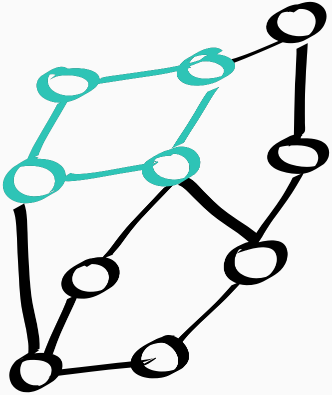

For $S \subseteq V$ ●,

Conductance More Formally

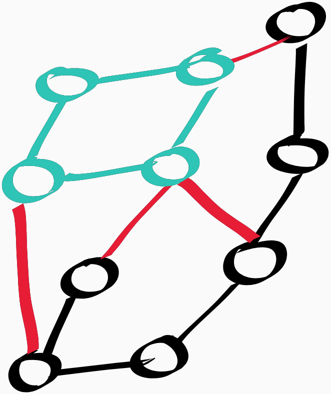

For $S \subseteq V$ ●, $\Phi(S) = \frac{|E(S, V \setminus S)|}{|(S \times V) \cap E|}$ ● ●●

Conductance More Formally

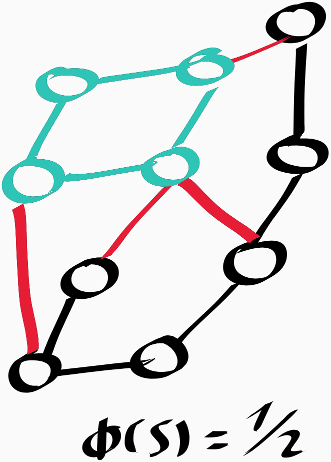

For $S \subseteq V$ ●, $\Phi(S) = \dfrac{|E(S, V \setminus S)| \; ●}{|(S \times V) \cap E| \; ●●}$

Conductance More Formally

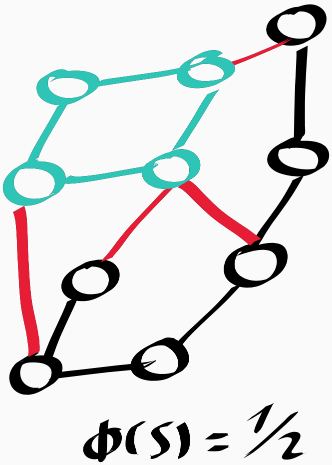

For $S \subseteq V$ ●, $\Phi(S) = \dfrac{|E(S, V \setminus S)| \text{ ●}}{|(S \times V) \cap E| \text{ ●●}}$

$$\Phi(G) = \min_{\substack{S \subset V \\ |E(S,S)| \leq |E(\bar{S},\bar{S})|}} \Phi(S)$$

Testing of Conductance

Theorem

There is a tester for conductance Φ in the CONGEST model with round complexity $O(\frac{\log n}{\epsilon\Phi^2})$, and a lower bound of $\Omega(\log n)$.

Testing of Conductance

Theorem

There is a tester for conductance Φ in the CONGEST model with round complexity $O(\frac{\log n}{\epsilon\Phi^2})$, and a lower bound of $\Omega(\log n)$.

- tester works also for connected graphs of unknown size
- votes can be made all accept / all reject

Lazy Random Walks

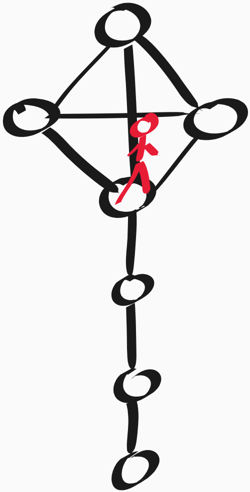

- random walker starts on $s \in V$

Lazy Random Walks

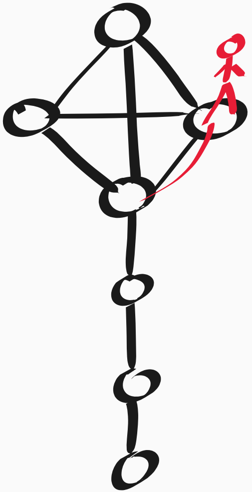

- random walker starts on $s \in V$
- goes $u \to v$, $v \in \Gamma(u)$ with probability

$$p(v,u) = \begin{cases} \frac{1}{2d(u)} & \text{if } u \neq v \\ \frac{1}{2} & \text{if } u = v \end{cases}$$

Lazy Random Walks

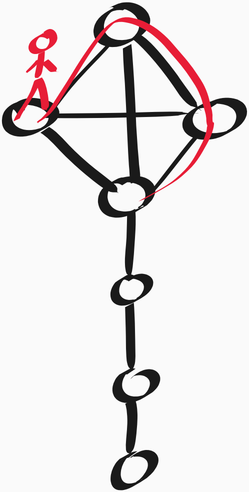

- random walker starts on $s \in V$
- goes $u \to v$, $v \in \Gamma(u)$ with probability

$$p(v,u) = \begin{cases} \frac{1}{2d(u)} & \text{if } u \neq v \\ \frac{1}{2} & \text{if } u = v \end{cases}$$

Lazy Random Walks

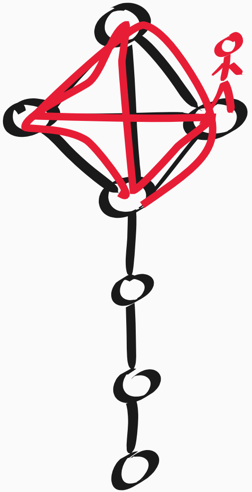

- random walker starts on $s \in V$
- goes $u \to v$, $v \in \Gamma(u)$ with probability

$$p(v,u) = \begin{cases} \frac{1}{2d(u)} & \text{if } u \neq v \\ \frac{1}{2} & \text{if } u = v \end{cases}$$

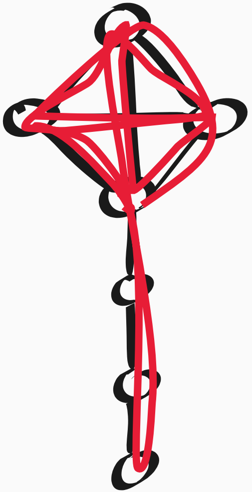

- random walker starts on $s \in V$
- goes $u \to v$, $v \in \Gamma(u)$ with probability

$$p(v,u) = \begin{cases} \frac{1}{2d(u)} & \text{if } u \neq v \\ \frac{1}{2} & \text{if } u = v \end{cases}$$

Lazy Random Walks

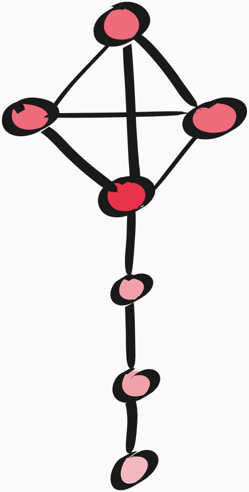

- random walker starts on $s \in V$
- goes $u \to v$, $v \in \Gamma(u)$ with probability

$$p(v, u) = \begin{cases} \frac{1}{2d(u)} & \text{if } u \neq v \\ \frac{1}{2} & \text{if } u = v \end{cases}$$

- stationary distribution

$$\vec{\pi}_v = d(v)/(2m)$$

Lazy Random Walks

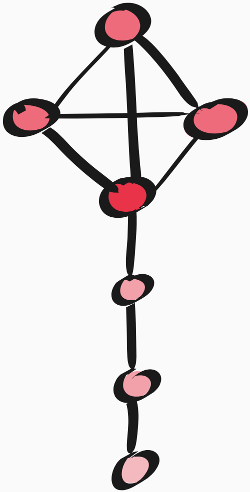

- random walker starts on $s \in V$
- goes $u \to v$, $v \in \Gamma(u)$ with probability

$$p(v,u) = \begin{cases} \frac{1}{2d(u)} & \text{if } u \neq v \\ \frac{1}{2} & \text{if } u = v \end{cases}$$

- stationary distribution

$$\vec{\pi}_v = d(v)/(2m)$$

- walk mixes, that is, converges to $\vec{\pi}$

$$\lim_{t\to\infty} \|P^t \vec{\mathbb{1}}_s - \vec{\pi}\| = 0$$

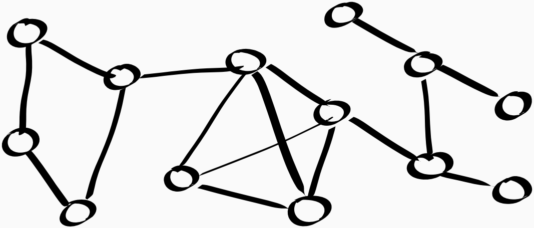

idea test for vertices with large mixing time

Idea of the Algorithm

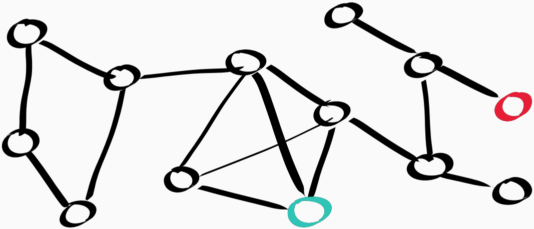

idea test for vertices with large mixing time

algorithm 1. sample $\Theta(1/\epsilon)$ vertices S ● ●

Idea of the Algorithm

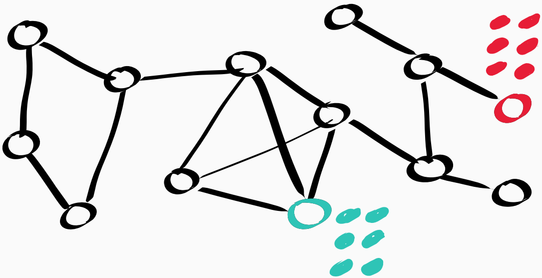

idea test for vertices with large mixing time

algorithm
1. sample $\Theta(1/\epsilon)$ vertices S ● ●
2. perform poly(n) random walks from S

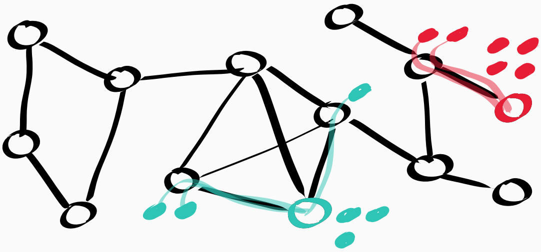

idea test for vertices with large mixing time

algorithm
1. sample $\Theta(1/\epsilon)$ vertices S ● ●
2. perform $\mathsf{poly}(n)$ random walks from S

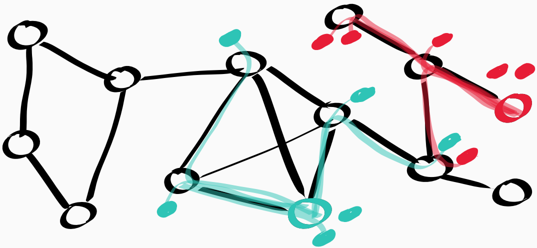

idea test for vertices with large mixing time

algorithm
1. sample $\Theta(1/\epsilon)$ vertices S ● ●
2. perform $\mathsf{poly}(n)$ random walks from S

Idea of the Algorithm

idea test for vertices with large mixing time

algorithm
1. sample $\Theta(1/\epsilon)$ vertices S ● ●
2. perform $\mathrm{poly}(n)$ random walks from S
3. check if walks for some $v \in S$ mixed poorly after $\Theta(\log n)$ steps

Idea of the Algorithm

idea test for vertices with large mixing time

algorithm

1. sample $\Theta(1/\epsilon)$ vertices S ● ●
2. perform $\mathsf{poly}(n)$ random walks from S
3. check if walks for some $v \in S$ mixed poorly after $\Theta(\log n)$ steps

...but keeping all traces is costly: $> \mathsf{poly}(n)$ bits

Reducing Congestion

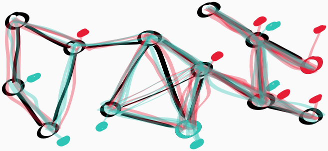

1. attempt: transmit full traces

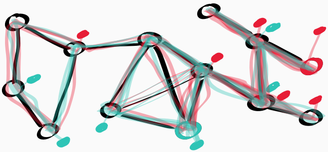

1. attempt: transmit full traces

can approximate $\|P^t \vec{\mathbb{1}}_s - \pi\| \ \forall s \in S$

Reducing Congestion

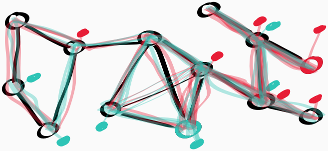

1. attempt: transmit full traces
can approximate $\|P^t \vec{\mathbb{1}}_s - \pi\| \; \forall s \in S$
but $\log n \cdot \text{poly}(n) = \Omega(n)$ rounds

Reducing Congestion

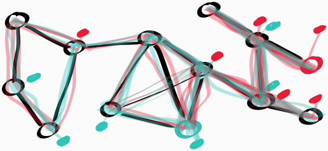

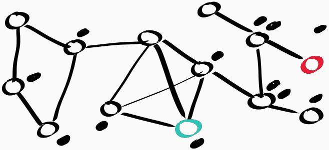

1. attempt: transmit full traces
   can approximate $\|P^t \vec{\mathbb{1}}_s - \pi\| \; \forall s \in S$
   but $\log n \cdot \mathrm{poly}(n) = \Omega(n)$ rounds

2. attempt: transmit only tokens

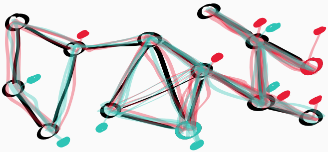

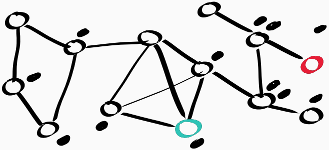

1. attempt: transmit full traces
can approximate $\|P^t \vec{\mathbb{1}}_s - \pi\| \; \forall s \in S$
but $\log n \cdot \mathrm{poly}(n) = \Omega(n)$ rounds

2. attempt: transmit only tokens
requires only $\mathcal{O}(\log n)$ rounds

Reducing Congestion

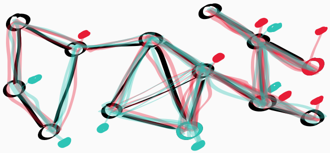

1. **attempt: transmit full traces**
can approximate $\|P^t \vec{\mathbb{1}}_s - \pi\| \; \forall s \in S$
but $\log n \cdot \text{poly}(n) = \Omega(n)$ rounds

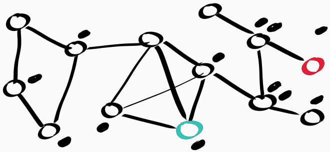

2. **attempt: transmit only tokens**
requires only $\mathcal{O}(\log n)$ rounds
but approx. $\|P^t \frac{1}{|S|} \sum_{s \in S} \vec{\mathbb{1}}_s - \pi\|$ only

Reducing Congestion

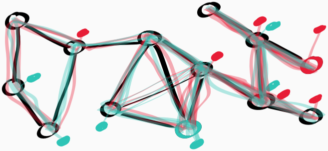

1. **attempt: transmit full traces**
   can approximate $\|P^t\vec{\mathbb{1}}_s - \pi\| \ \forall s \in S$
   but $\log n \cdot \mathrm{poly}(n) = \Omega(n)$ rounds

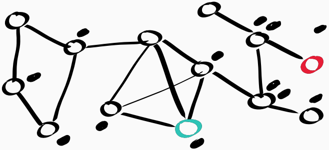

2. **attempt: transmit only tokens**
   requires only $\mathcal{O}(\log n)$ rounds
   but approx. $\|P^t \frac{1}{|S|} \sum_{s \in S} \vec{\mathbb{1}}_s - \pi\|$ only

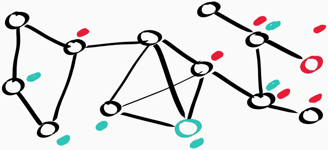

3. **attempt: transmit start vertices**

Reducing Congestion

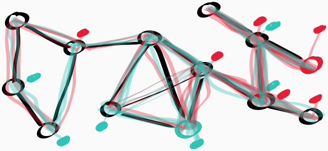

1. **attempt: transmit full traces**
   can approximate $\|P^t \vec{\mathbb{1}}_s - \pi\| \; \forall s \in S$
   but $\log n \cdot \mathrm{poly}(n) = \Omega(n)$ rounds

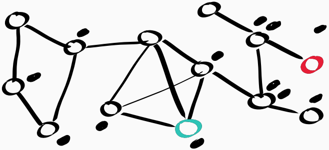

2. **attempt: transmit only tokens**
   requires only $\mathcal{O}(\log n)$ rounds
   but approx. $\|P^t \frac{1}{|S|} \sum_{s \in S} \vec{\mathbb{1}}_s - \pi\|$ only

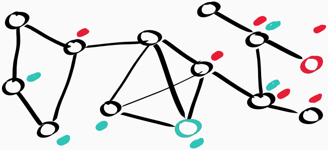

3. **attempt: transmit start vertices**
   can approximate $\|P^t \vec{\mathbb{1}}_s - \pi\| \; \forall s \in S$

Reducing Congestion

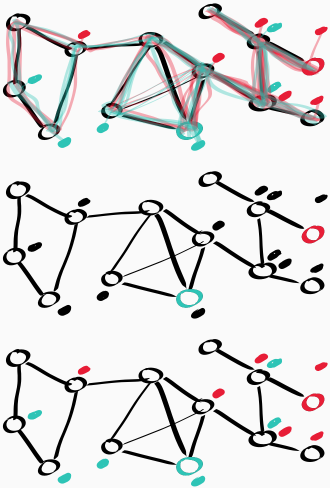

1. **attempt: transmit full traces**
can approximate $\|P^t \vec{\mathbb{1}}_s - \pi\| \ \forall s \in S$
but $\log n \cdot \mathrm{poly}(n) = \Omega(n)$ rounds

2. **attempt: transmit only tokens**
requires only $\mathcal{O}(\log n)$ rounds
but approx. $\|P^t \frac{1}{|S|} \sum_{s \in S} \vec{\mathbb{1}}_s - \pi\|$ only

3. **attempt: transmit start vertices**
can approximate $\|P^t \vec{\mathbb{1}}_s - \pi\| \ \forall s \in S$
requires only $\mathcal{O}(\log n)$ rounds

Summary

Theorem

There is a tester for conductance Φ in the CONGEST model with round complexity $O(\frac{\log n}{\epsilon \Phi^2})$, and a lower bound of $\Omega(\log n)$.

- lower bound is based on high girth expander graphs

Summary

Theorem

There is a tester for conductance Φ in the CONGEST model with round complexity $O(\frac{\log n}{\epsilon\Phi^2})$, and a lower bound of $\Omega(\log n)$.

- lower bound is based on high girth expander graphs
- first two-sided error distributed tester
    - voting rule taken from one-sided error testing
    - power of other rules?

Summary

Theorem

There is a tester for conductance Φ in the CONGEST model with round complexity $O(\frac{\log n}{\epsilon\Phi^2})$, and a lower bound of $\Omega(\log n)$.

- lower bound is based on high girth expander graphs
- first two-sided error distributed tester
    - voting rule taken from one-sided error testing
    - power of other rules?
- lower bound for one-sided error tester of conductance?