

Distributed Algorithms vs. sublinear time algorithms on SPARSE graphs

↑
max deg $\leq d$

Again, Sparse graphs: max degree: d
adj list representation

A problem to solve:

Vertex Cover

$V' \subseteq V$ is "Vertex Cover" (VC) if $\forall (u, v) \in E$

either $u \in V'$ or $v \in V'$

VC Question: What is min size of VC?

Note: in $\text{deg} \leq d$ graph, $|VC| \geq \frac{m}{d}$ since each node can cover $\leq d$ edges

(VC is NP-complete, but there is a polytime 2-multiplicative approximation)

Can you approximate V.C. in sublinear time?

multiplicative? no! graph with n edges $|VC|=0$ \rightarrow can't distinguish these cases in sublinear time but must answer 0 in first case & >0 in second.

additive? hard! need some mult error

computationally hard to approx to

better than 1.36 factor (maybe even 2)

Combination?

def. \hat{y} is (α, ϵ) -estimate of soln value y for minimization problem if

$$y \leq \hat{y} \leq \alpha y + \epsilon$$

allows mult + additive error

(analogous defn for maximization problems)

Some Background on distributed Algorithms


- Network
 - processors \rightarrow max degree d known to all
 - links
- Communication round
 - nodes send messages to neighbors

def. Vertex Cover problem for distributed networks:

- Network graph = input graph (i.e. network computes on itself) ↙ not some other graph
- at end, each node knows if in or out of VC (doesn't know about others necessarily)

Main insight on why fast distributed \Leftrightarrow sublinear time:

in k round algorithm, output of node v only depends on nodes at distance at most k from v . At most d^k of these!



⇒ Can ^{sequentially} simulate V 's view of distributed computation in $\leq d^k$ time
↓ figure out if v in or out of VC

Comment: if algorithm is randomized, v needs to know random bits (or be able to construct) of all d^k nbors. \leftarrow must be consistent

∴ fast distributed alg ⇒ "oracle" which tells you if v is in VC


But are there fast VC distributed algorithms?

YES, will see some soon

↑ often called "local distributed algorithms"

How do you use this to approximate VC in sublinear time?

Parnas-Ron framework:



Sample nodes of graph v_1, \dots, v_r

for each v_i , simulated distributed algorithm to see if $v_i \in VC$

Output $\frac{\#v_i\text{'s in VC}}{r}$ on

Runtime $O(r \cdot d^{k/r}) \approx O(\frac{1}{\epsilon^2} d^{k/r})$

Proof of correctness Chernoff/Hoeffding bnds

fast distributed algorithm for VC:

$i \leftarrow 1$

While edges remain:

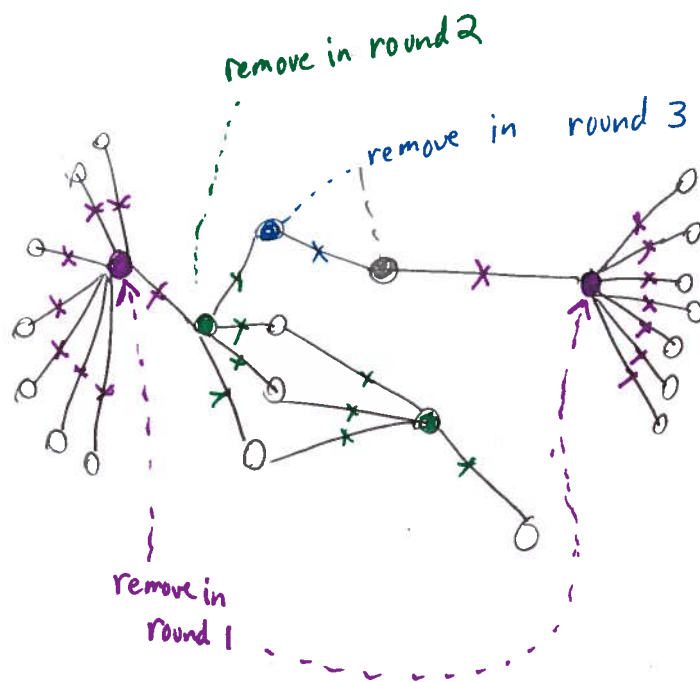
- remove vertices of degree $\geq d/2^i$ + adjacent edges
- update degrees of remaining nodes
- increment i

put these in vertex cover

Output all removed nodes as VC

#rounds: $\log d$

example:



is it a VC?

no edges remain at end

all removed along with some adjacent vertices

Is it a good approximation?


Thm let $VC(G)$ = size of min VC of G

Then, $VC(G) \leq \text{output} \leq (2 \log d + 1) VC(G)$
 ↑ since output is VC ↑ to prove

Proof.

Claim: in each iteration, add $\leq 2 VC_G$ new vertices
 ← any thing bigger removed earlier

why: all nodes removed have deg bet $\frac{d}{2^{i+1}} + \frac{d}{2^i}$
 $\Theta \leftarrow$ any min VC (so any edge has to have ≥ 1 endpt in Θ)



so $|X| \cdot \frac{d}{2^i} \leq |\Theta| \cdot \frac{d}{2^{i-1}}$

so $|X| \leq 2|\Theta|$