Local Computation Algorithms:

- the model
- Maximal Independent Set

Local Computation Algorithms (LCAs)

when input <u>and</u> output large,

do we need to construct full output?

" " " " read " input?

LCAs provide fast query access to output

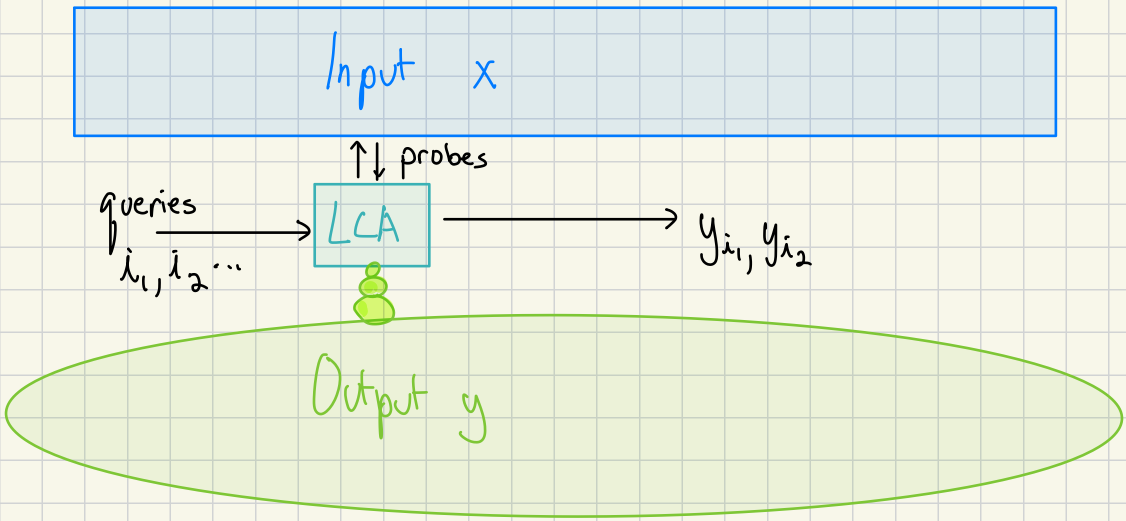

For example:

Maximal Independent Set: $X =$ description of graph

$Y = (y_1 \ldots y_n)$ s.t. $y_i = \begin{cases} 1 & \text{if node } i \text{ in MIS} \\ 0 & \text{o.w} \end{cases}$

def $U \subseteq V$ is a "maximal independent set" (MIS) if

(1) $\forall\, u_1, u_2 \in U \quad (u_1, u_2) \notin E$ "independent"

(2) $\nexists\, w \in V \setminus U$ s.t. $U \cup \{w\}$ is independent "maximal"

G has max degree Δ

Queries: is node u in MIS?

v ?

← do we need to compute whole MIS in order to answer?

Problem: lots of MIS's possible!

every u is in some MIS

Can we always answer "yes"?

How to define LCA?

more than one "legal" output?

depends on query order?

need to remember past answers?

would like answer to be "No" for above

def $\mathcal{A}$ is "Local computation algorithm" (LCA) for a problem Π if

- given:
  - probe access to input X
  - random bits r
  - local memory
- answers queries to bits/words of

$$\cancel{y = \Pi(x)}$$

$$y \in \Pi(x)$$

what if more than one soln?

maybe $y = \Pi_r(x)$? ← determined by random bits

- memory of LCA "wiped" clean between queries (Keep the random bits)
- $\forall$ queries i,

let $y_i \leftarrow$ output of LCA on i

note: can't depend on other queries since "wiped"

then y_i's must be consistent with a legal solution for input X. (e.g. $\in \Pi(x)$)

Can we provide "illusion" of fully constructed output?

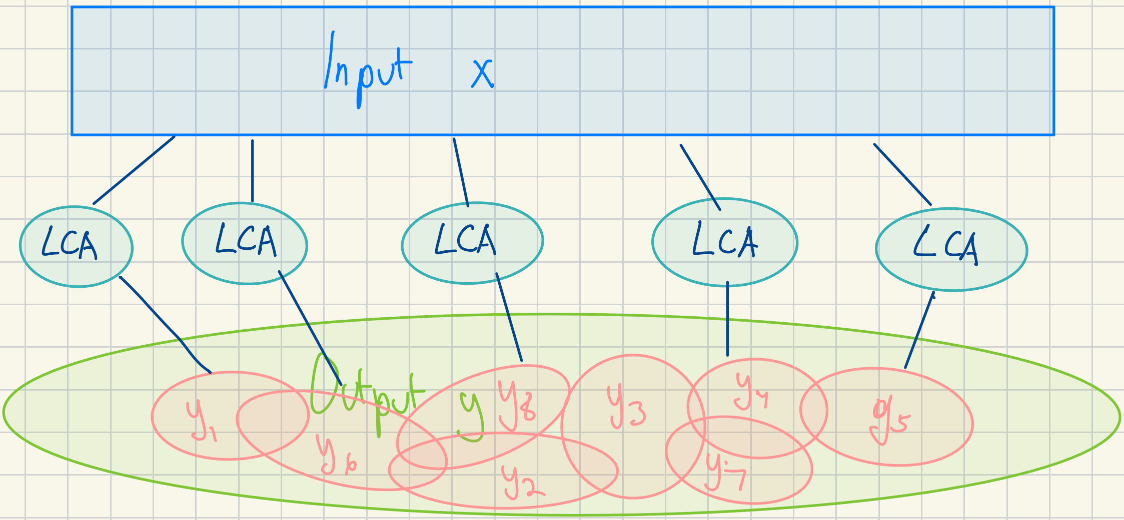

Equivalent model:

- many copies, each gets only <u>one</u> question
  - initially share (short?) random string
  - afterwards compute independently
- sublinear time, space, randomness

Maximal Independent Set:

$X =$ description of graph

$Y = (y_1 \dots y_n)$ s.t. $y_i = \begin{cases} 1 & \text{if node } i \text{ in MIS} \\ 0 & \text{o.w} \end{cases}$

<u>def</u> $U \subseteq V$ is a "maximal independent set" (MIS) if

(1) $\forall\, u_1, u_2 \in U \quad (u_1, u_2) \notin E$ "independent"

(2) $\nexists\, w \in V \setminus U$ s.t. $U \cup \{w\}$ is independent "maximal"

G has max degree Δ

<u>note</u> MIS can be solved via greedy

but maximum independent set is NP-complete

<u>Plan</u>
- show distributed alg for MIS in k rounds
- convert to $\Delta^{O(k)}$ sublinear query/time algorithm

Distributed Algorithm for MIS:

"Luby's Algorithm" (actually a variant...)

- $MIS \leftarrow \phi$
- all nodes set to "live"
- repeat k times in parallel:
  - $\forall$ nodes v, color self "red" with prob $\frac{1}{2\Delta}$ else "blue"
    send color to all nbrs
  - if v colors self red, & no other nbr of v colors self red, then
    - add v to MIS
    - remove v & all nbrs from graph (set to "inactive")

(For purposes of analysis, continue to select colors after removed but don't send to nbrs)

Thm $\Pr[\#\text{ phases til graph empty} \geq 8\Delta \log n] \leq \frac{1}{n}$

Corr $E[\#\text{ phases}]$ is $O(\Delta \log n)$ ← can improve!

Main Lemma: For live v, $\Pr[v \text{ added to MIS in a round}] \geq \frac{1}{4\Delta}$

Proof:

$\Pr[v \text{ colors self red}] = \frac{1}{2\Delta}$

$\Pr[\underline{\text{any}}\ w \in Nbr(v) \text{ colors self red}] \leq \sum_{w \in N(v)} \frac{1}{2\Delta}$ (union bnd)

$\leq \frac{1}{2}$ (degree bnd)

$\therefore \Pr[v \text{ colors self red} + \text{no other nbr colors self red}]$

$\leq \frac{1}{2\Delta} \cdot \left(1 - \frac{1}{2}\right) = \frac{1}{4\Delta}$ (independence)

$\blacksquare$

Lemma $\Rightarrow$ $\Pr[v \text{ live after } \underbrace{4 \cdot k' \cdot \Delta}_{=k} \text{ rounds}] \leq \left(1 - \frac{1}{4\Delta}\right)^{4\Delta \cdot k'}$

$\leq e^{-k'}$

Setting k:

if $k = \Theta(\Delta \log n)$, $\Pr[v \text{ live at end}] \leq e^{-\Theta(\log n)}$

$= \frac{1}{n^c}$

(can do better)

Problem: when sequentially simulate $\Theta(\Delta \log n)$ rounds,

$\approx n^{\Delta \log \Delta} \xrightarrow{\text{need}} \Delta^{\Theta(\Delta \log n)}$ complexity .... not sublinear in n ☹

Idea

- define $\boxed{\text{"Luby-status": run Luby for } K = \Theta(\Delta \log \Delta) \text{ rounds}}$

  at end, each $v \in V$ is one of

  $\begin{cases} \text{live} \\ \text{in MIS} & \leftarrow \text{set self to red \& no nbrs did} \\ \underline{\text{not}} \text{ in MIS} & \leftarrow \text{taken out by nbr in MIS} \end{cases}$

- use [Parnas-Ron] (distributed → sequential) reduction to simulate v's view of Luby-status in sequential manner: $\Delta^{K} = \Delta^{\Theta(\Delta \log \Delta)}$ queries

  ← rounds

  ↑ degree

- if v is in/not in then done

  else v is alive ← what do we do now?

Questions:

- what is probability that v still alive?

$$\Pr[v \text{ inactive}] \le \Pr[\text{survives } \Theta(\Delta \log \Delta) \text{ rounds}]$$
$$\le e^{-\Theta(\Delta \log \Delta)} \le \frac{1}{\Delta^c}$$

for any c & for proper choice of constant in # rounds

Very few stay active!

- how much work to "finish off" v?

Key insight distribution of live nodes!

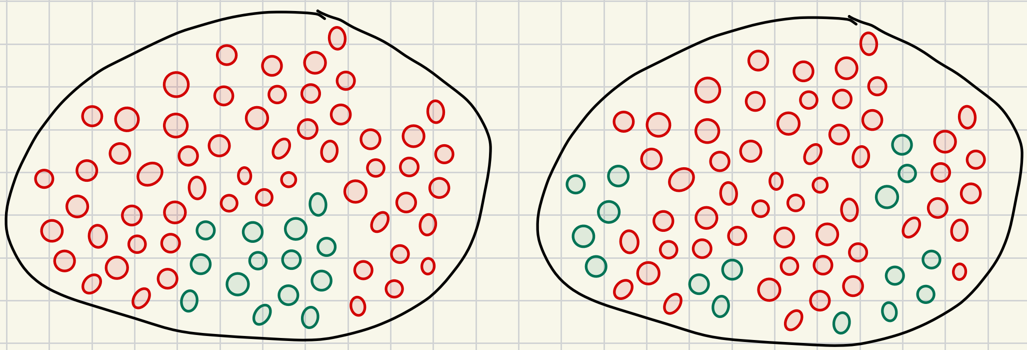

big clumps of surviving nodes?

NO!

surviving nodes in small connected components

relies on degree bound of graph

$\Rightarrow$ survival of components $\approx$ independent