Local Computation Algorithms

Ronitt Rubinfeld

MIT

Large Inputs

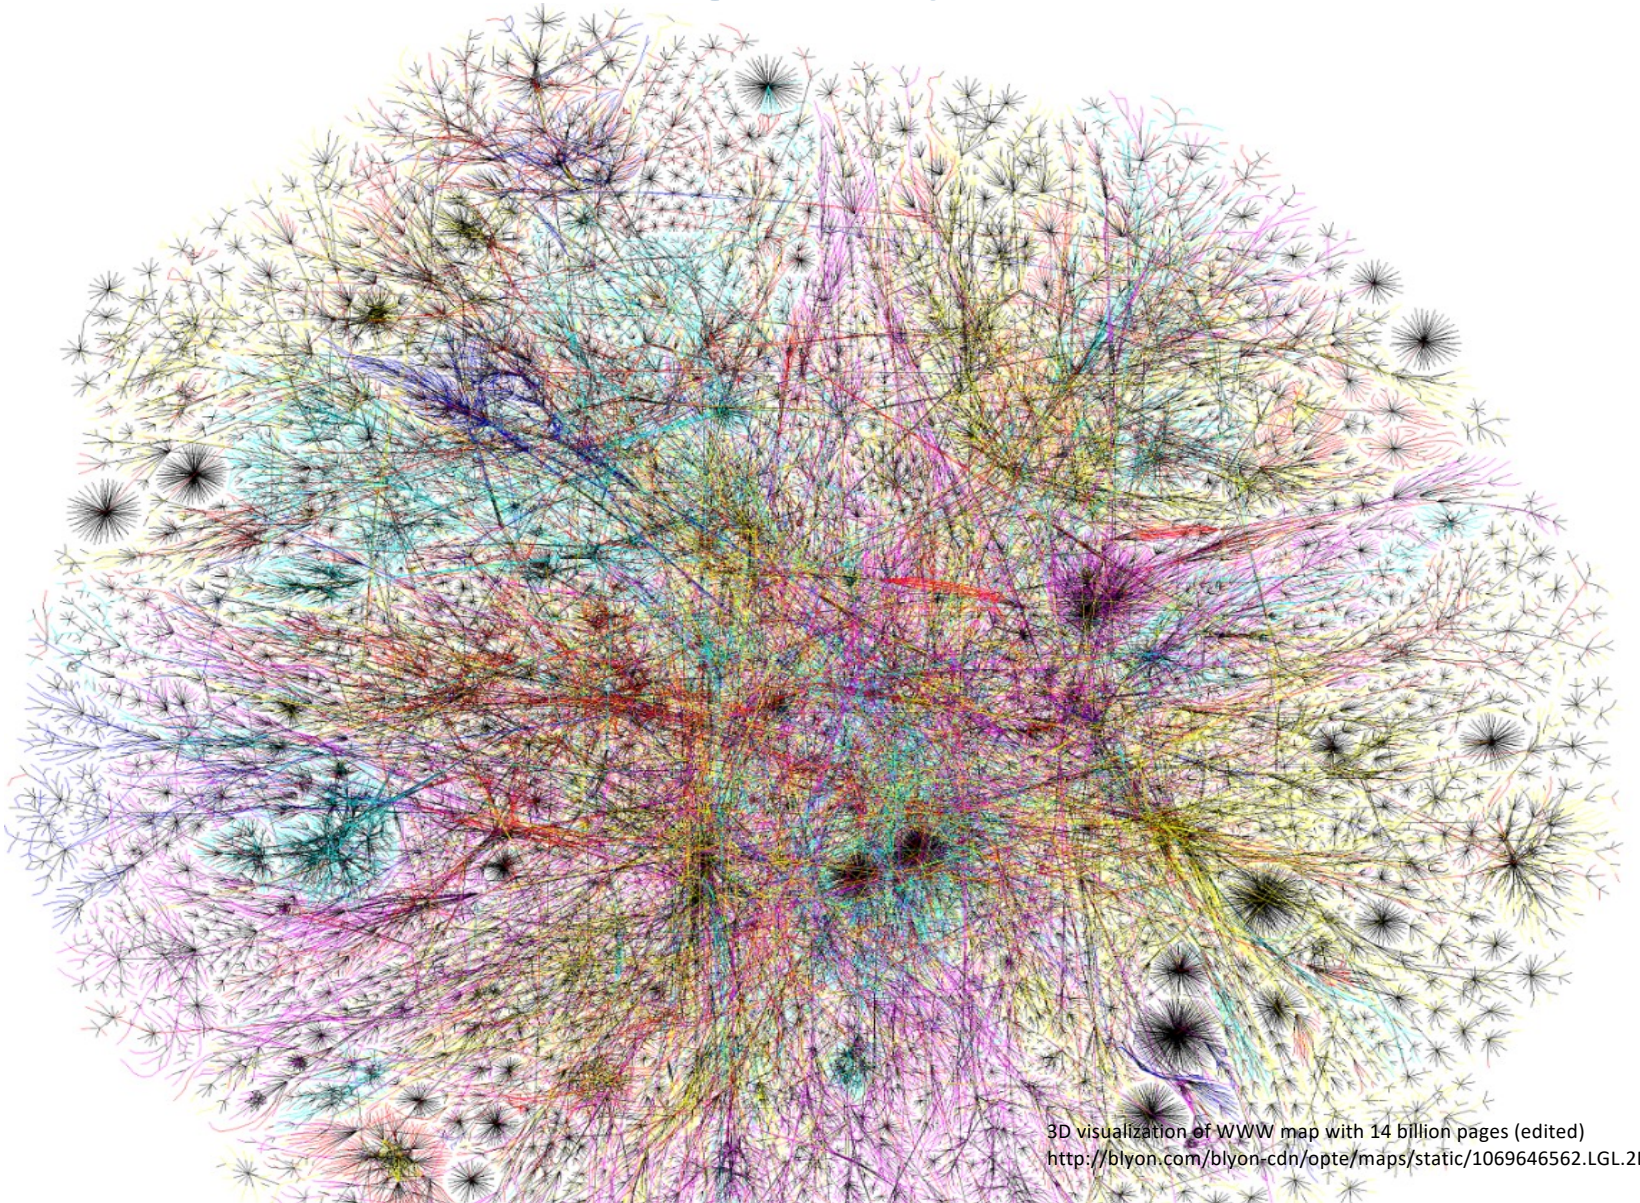

3D visualization of WWW map with 14 billion pages (edited)
http://blyon.com/blyon-cdn/opte/maps/static/1069646562.LGL.2D.4096x4096.png

Large inputs

Large outputs

When we don’t need to see the whole output…

do we need to compute the whole output?

do we need to see all the input?

Some examples

Locally decodable codes

Local property reconstruction

Local decompression

Estimating graph parameters: page rank, communities, dominating set, …

Local Computation Algorithms: a model

[R Tamir Vardi Xie ‘11] [Alon R Vardi Xie ‘12]

Also known as: Centralized Local Model

Local Computation Algorithms

[R Tamir Vardi Xie] [Alon R Vardi Xie]

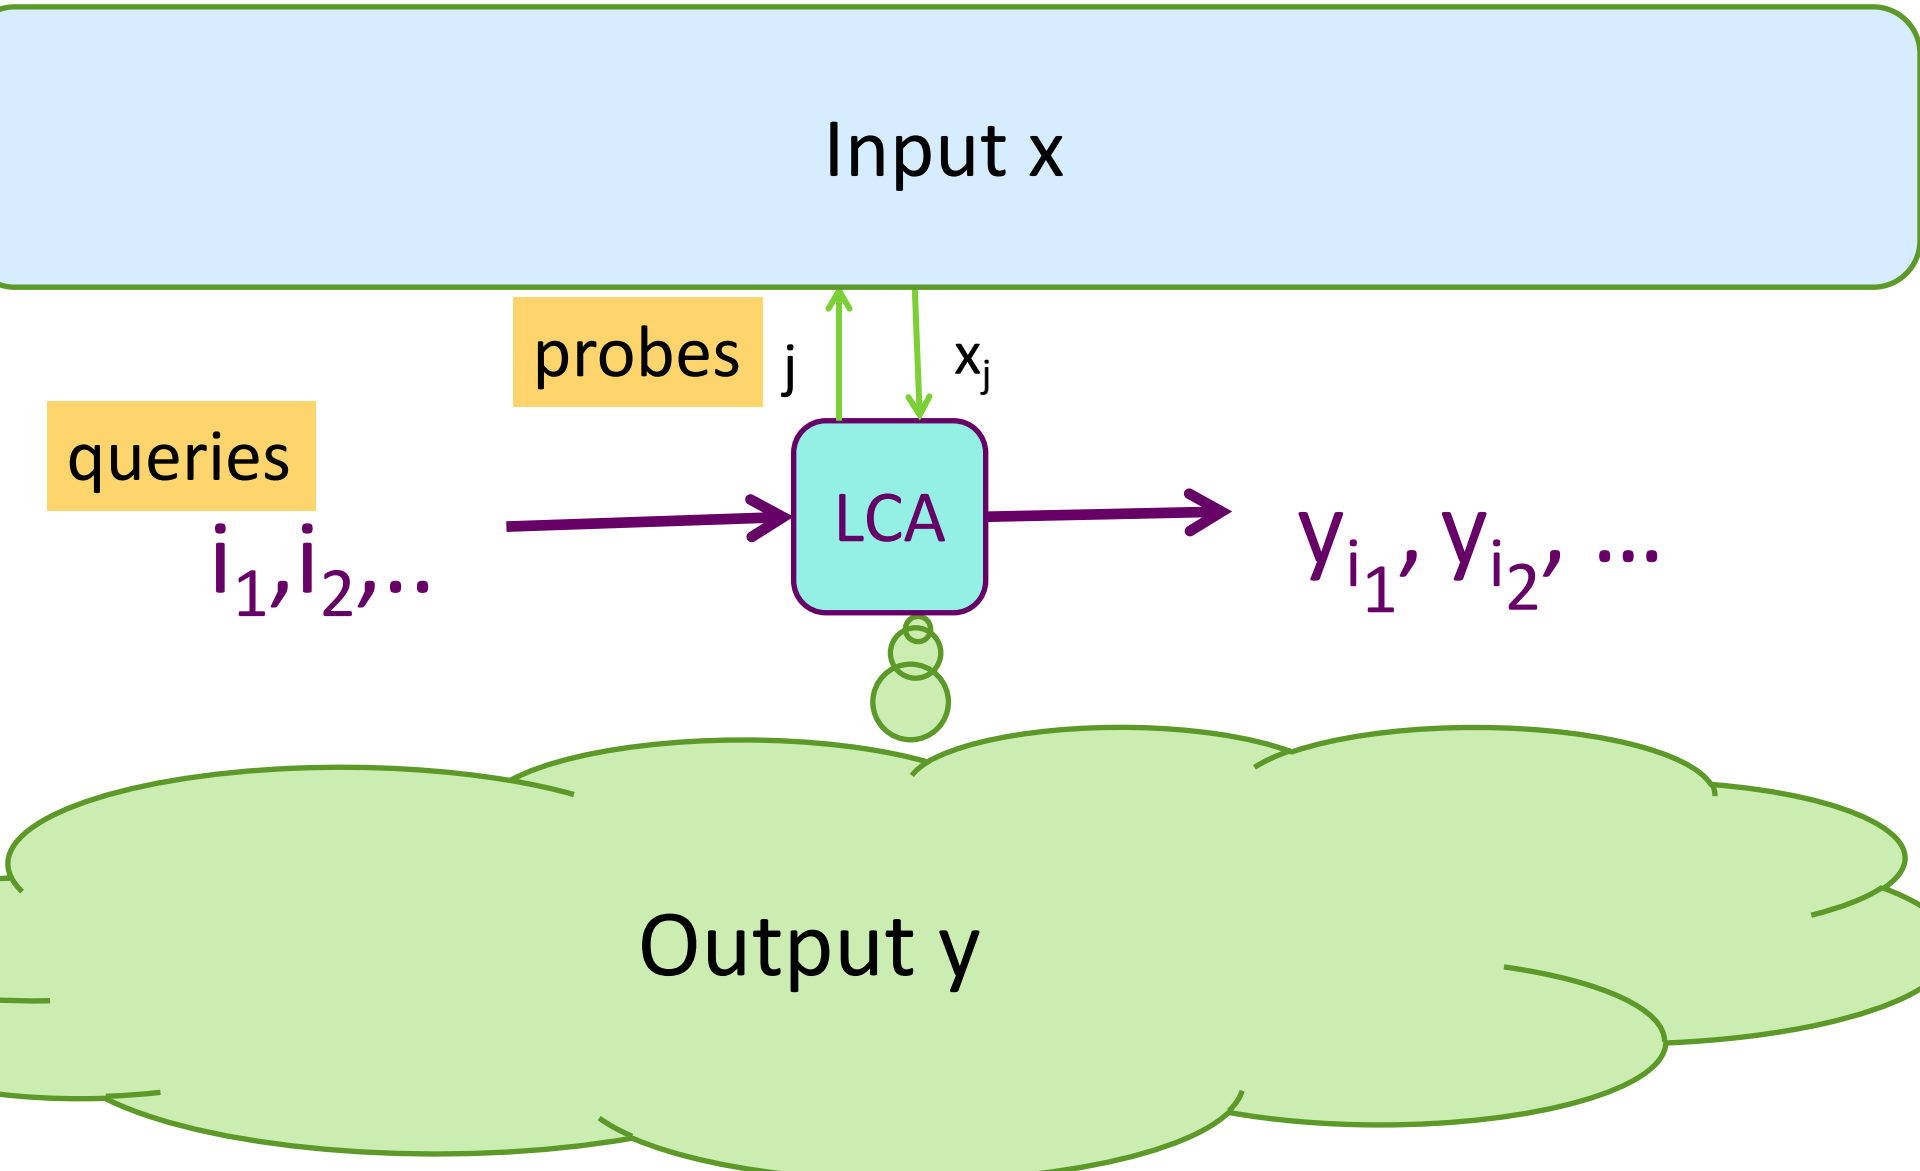

Some problems considered in this model

- Maximal Independent set [Barenboim Elkin] [R Tamir Vardi Xie] [Alon R Vardi Xie] [Barenboim Elkin Pettie Schneider] [Even Medina Ron][Reingold Vardi][Chung Pettie Su] [Levi R Yodpinyanee] [Ghaffari][Ghaffari Uitto]…
- Approximate maximum matching, bipartite weighted vertex cover [Mansour Vardi] [Even Medina Ron] [Fei… Mansour Schapire]
- Radio net…
- Graph, Hyp… maj Mansour Var…
- k-CNF [RTVX]
- Local computatio… r Vardi]
- Online algorithms [Mansour Rub… Xie]
  - load balancing balls and bins

Polyn…

Used…

Polylog probes (sequential)

Difficulty:

There may be more than one “legal” output

So what?

“Swarms” of LCAs

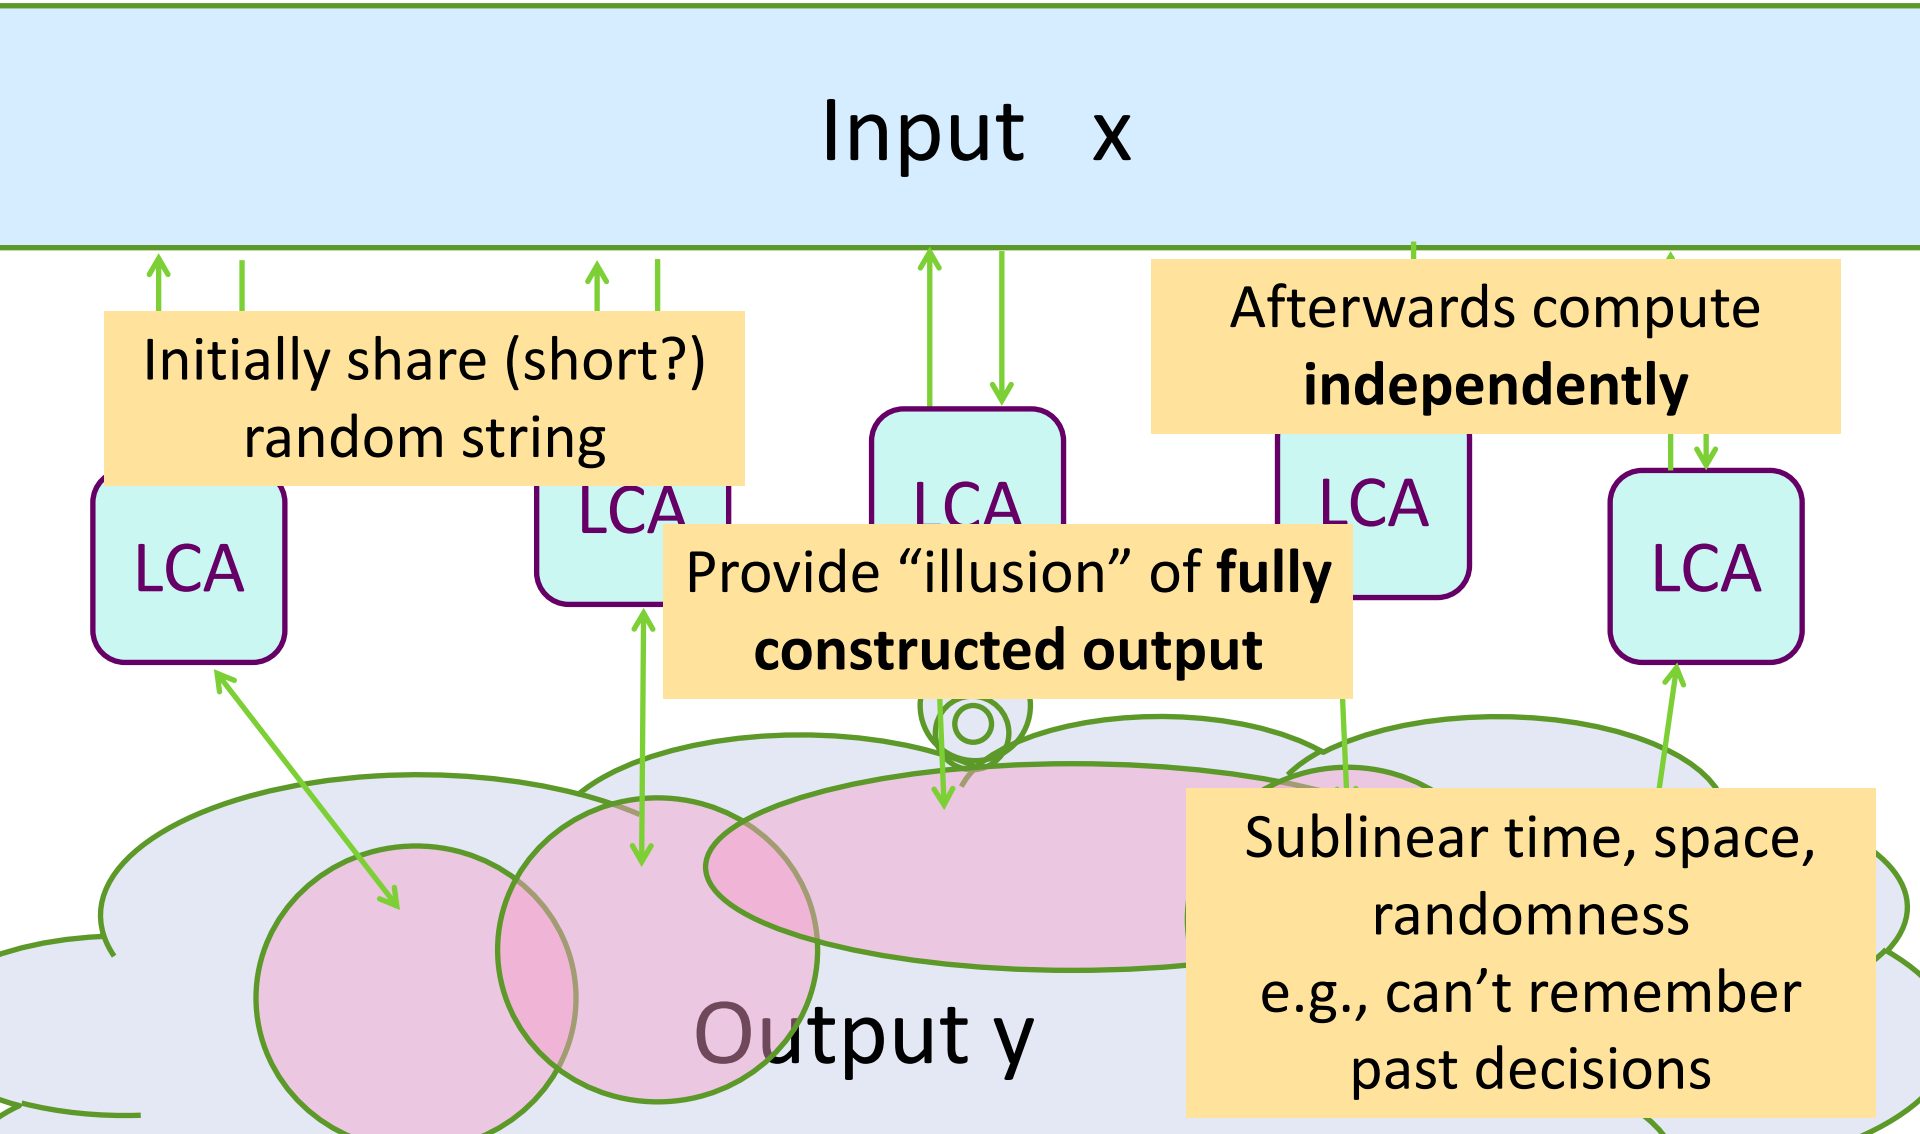

A challenge:

Consistency!

Query order oblivious?

How do we design good LCAs?

Maximal independent set

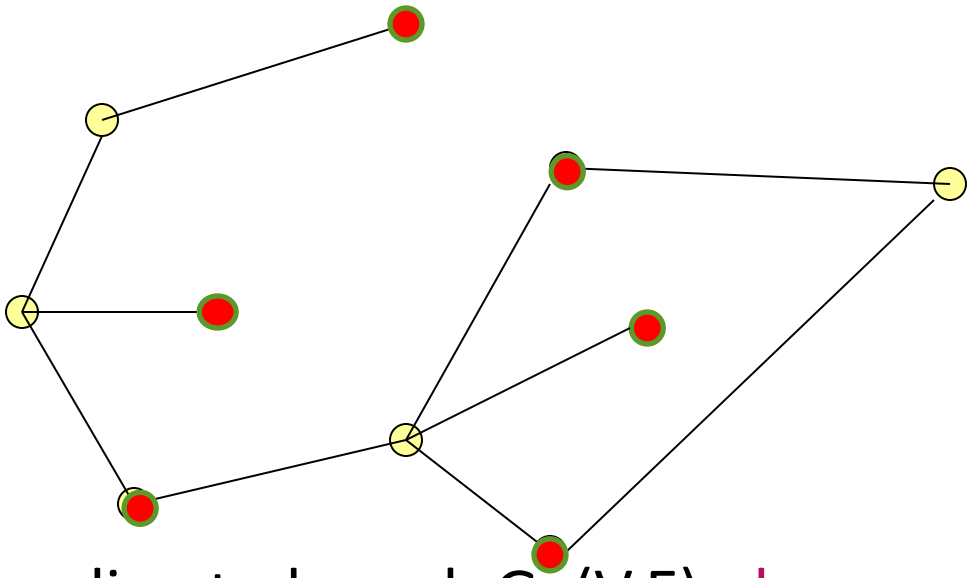

- Sparse undirected graph G=(V,E), degree at most d (constant)
- Independent set: subset V' of V such that no two vertices connected by an edge
- Maximal independent set: can't add any nodes to it

Idea 1:
Distributed Algorithms to the rescue!

Distributed algorithms give LCAs

[Parnas Ron]

- If there is a k round distributed algorithm for MIS, then:
  - v's output depends only on inputs and computations of k-radius ball around v
  - Can read/simulate in d^k probes!

- But how big is k?

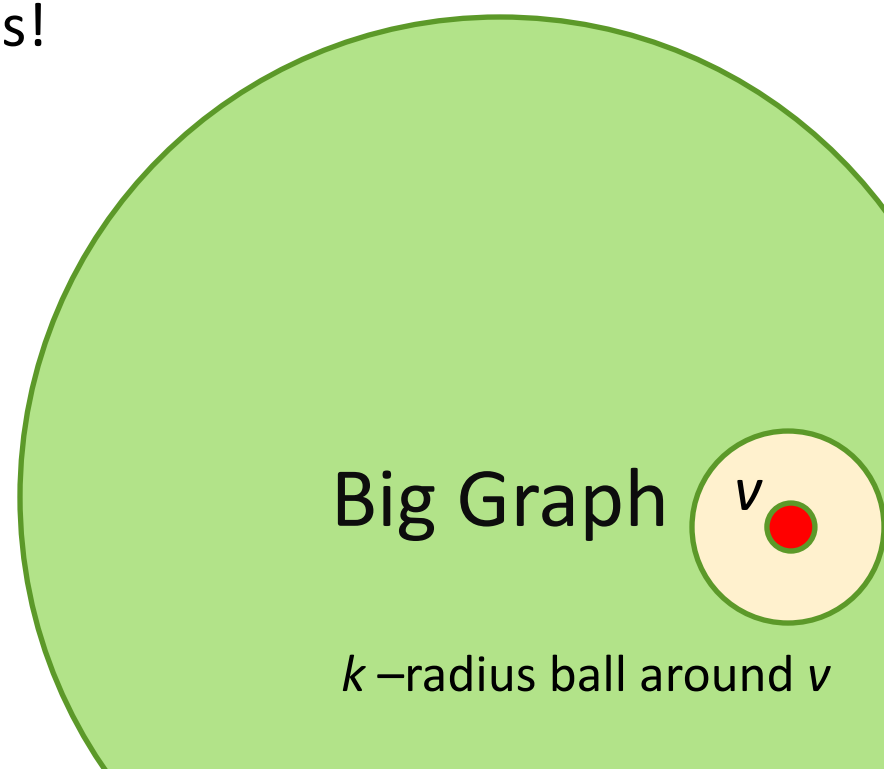

Local *distributed* algorithms

In this context:

Local = Constant rounds

fantastic progress in local distributed algorithms!!!

How fast can MIS be computed in a distributed setting?

- Lexicographically-first-MIS is P-complete [Cook]

- Randomized O(log n) rounds [Luby][Alon Babai Itai]
  - Yields $d^{\text{clog n}} = 2^{O(\log d \log n)}$ time LCA

Idea 1’:
Use distributed algorithms to “shatter” the graph

[R Tamir Vardi Xie]

Broad outline of many distributed MIS algorithms:

- All nodes start out “live”
- In each round:
  - Live nodes toss coins
  - Live nodes use coins and interaction with neighbors to decide whether to join MIS
  - If node or its neighbor joined MIS, node “dies”

Usually:
repeat until all nodes dead
$O(\log n)$ rounds

Here:
repeat until constant fraction dead
$O(d)$ rounds

After running constant number rounds of distributed algorithm …

[R Tamir Vardi Xie]

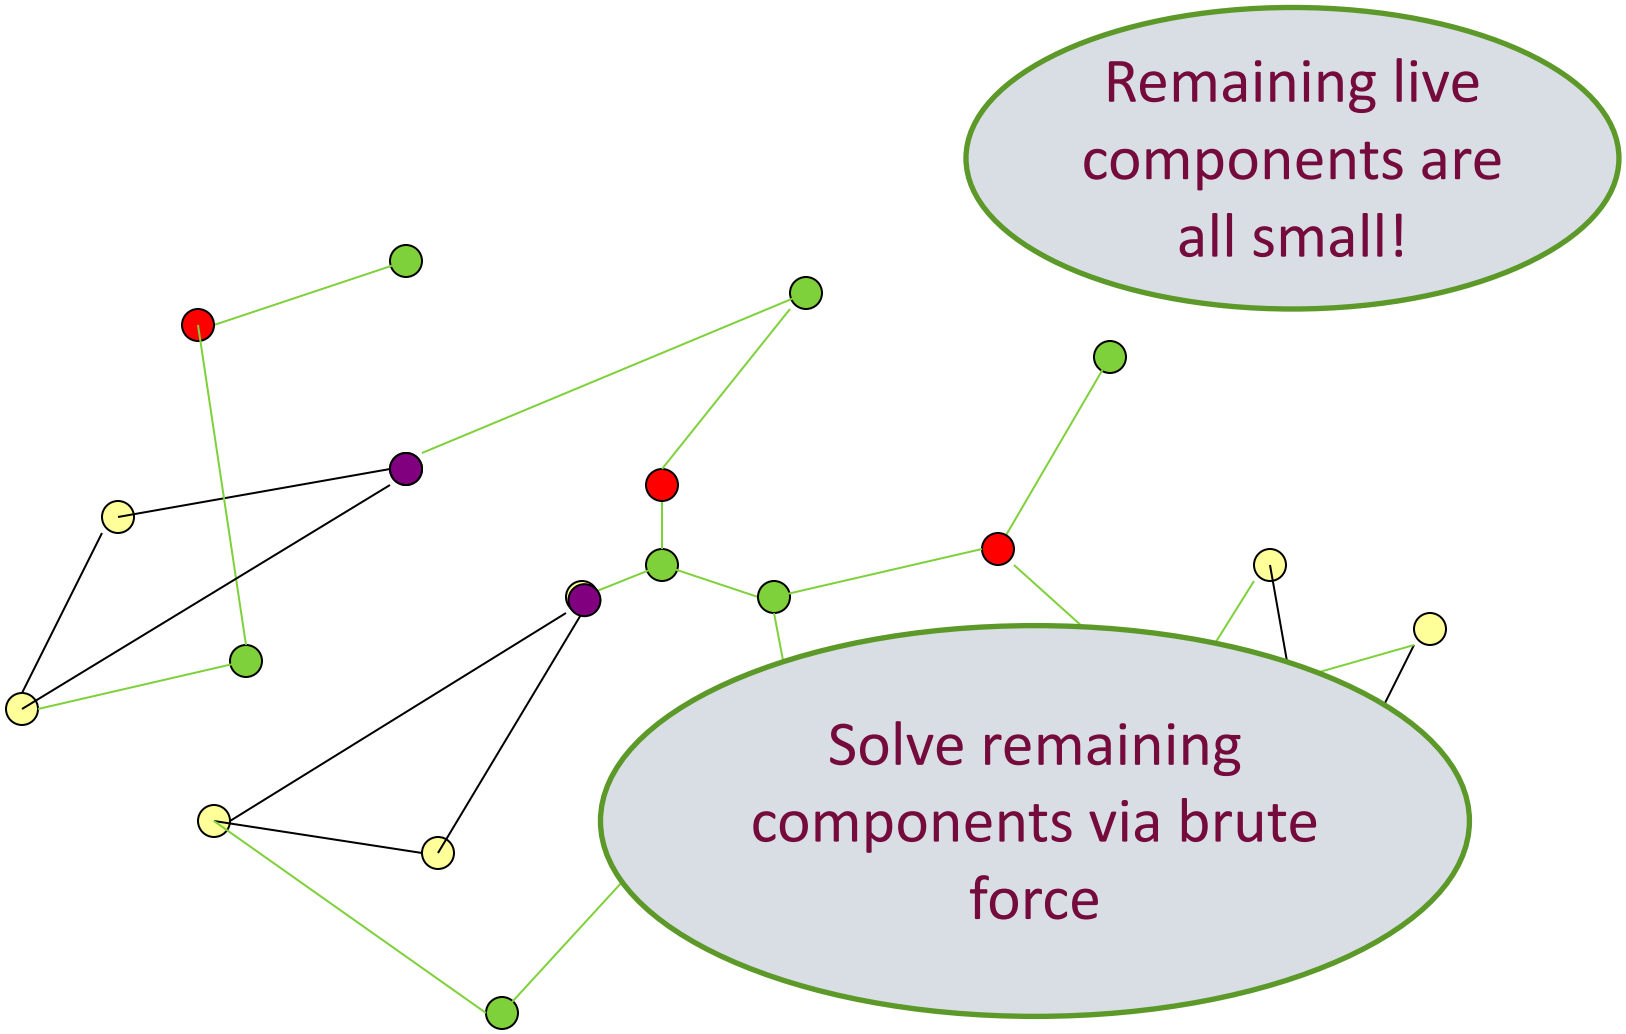

An example for MIS [RTVX,ARVX]

O(d) round distributed algorithm:

- Run *O(d)* rounds of variant of Luby's algorithm
- Prove remaining "live" connected components are small (logarithmic) via Beck-like analysis for algorithmic Lovasz Local Lemma

LCA(v)

ALIVE(v): sequential Parnas-Ron simulation of v's local view of distributed algorithm in $d^{O(d)}$ queries

- If v's output determined by ALIVE(v), output it
- Else:
  - Find v's small "live" connected component via BFS, calling ALIVE(w)
  - Output "yes" if v is in lexicographically first MIS in live component

Idea later used/improved in

[Barenboim Elkin Pettie Schneider] [Ghaffari] …