

Think Global – Act Local



Roger Wattenhofer

Think global ...



act local.

THINK GLOBAL
ACT LOCAL



Town Planning *Patrick Geddes*






Architecture *Buckminster Fuller*


Computer Architecture *Caching*

space (addresses)




Robot Gathering

e.g., [Degener et al., 2011]



Natural Algorithms



[Bernard Chazelle, 2009]



game theory

Algorithmic Trading



Think Global – Act Local



...is there a theory?

Complexity Theory

Can a Computer Solve
Problem P in Time t ?

(Think Global - Act Local)

Distributed




Complexity Theory

Network

Can a ~~Computer~~ Solve
Problem P in Time t ?


Distributed (Message-Passing) Algorithms

- Nodes are agents with unique ID's that can communicate with neighbors by **sending messages**. In each **synchronous round**, every node can send a (different) message to each neighbor.



Distributed (Message-Passing) Algorithms

- Nodes are agents with unique ID's that can communicate with neighbors by **sending messages**. In each **synchronous round**, every node can send a (different) message to each neighbor.




each round:
every node:
1. send msgs
2. rcv msgs
3. compute

- Distributed (Time) Complexity**: How many rounds does problem take?

An Example


each round:
every node:
1. send msgs
2. rcv msgs
3. compute

How Many Nodes in Network?




each round:
every node:
1. send msgs
2. rcv msgs
3. compute

How Many Nodes in Network?




each round:
every node:
1. send msgs
2. rcv msgs
3. compute

How Many Nodes in Network?




each round:
every node:
1. send msgs
2. rcv msgs
3. compute


How Many Nodes in Network?




How Many Nodes in Network?



How Many Nodes in Network?




How Many Nodes in Network?




With a simple flooding/echo process, a network can find the number of nodes in **time** $O(D)$, where D is the diameter (size) of the network.

Diameter of Network?




- **Distance** between two nodes = Number of hops of shortest path

Diameter of Network?




- **Distance** between two nodes = Number of hops of shortest path

Diameter of Network?




- **Distance** between two nodes = Number of hops of shortest path
- **Diameter** of network = Maximum distance, between any two nodes


Diameter of Network?




Diameter of Network?




Diameter of Network?




Diameter of Network?




Diameter of Network?



Diameter of Network?




Diameter of Network?



Networks Cannot Compute Their Diameter in Sublinear Time!

(even if diameter is just a small constant)




Pair of rows connected neither left nor right? Communication complexity:
Transmit $\Theta(n^2)$ information over $O(n)$ edges $\rightarrow \Omega(n)$ time!

What about a “local” task?


Example: Minimum Vertex Cover (MVC)

- Given a **network** with n nodes, nodes have **unique IDs**.
- Find a **Minimum Vertex Cover (MVC)**
 - a minimum set of nodes such that all edges are adjacent to node in MVC




Example: Minimum Vertex Cover (MVC)

- Given a **network** with n nodes, nodes have **unique IDs**.
- Find a **Minimum Vertex Cover (MVC)**
 - a minimum set of nodes such that all edges are adjacent to node in MVC




Example: Minimum Vertex Cover (MVC)

- Given a **network** with n nodes, nodes have **unique IDs**.
- Find a **Minimum Vertex Cover (MVC)**
 - a minimum set of nodes such that all edges are adjacent to node in MVC



On MVC


- Find an MVC that is “close” to minimum (**approximation**)
- **Trade-off** between time complexity and approximation ratio



- MVC: Various simple (non-distributed) 2-approximations exist!
- What about **distributed algorithms**?!?


Finding the MVC (by Distributed Algorithm)

- Given the following bipartite graph with $|S_0| = \delta |S_1|$
- The MVC is just all the nodes in S_1
- Distributed Algorithm...




Finding the MVC (by Distributed Algorithm)

- Given the following bipartite graph with $|S_0| = \delta |S_1|$
- The MVC is just all the nodes in S_1
- Distributed Algorithm...




Finding the MVC (by Distributed Algorithm)

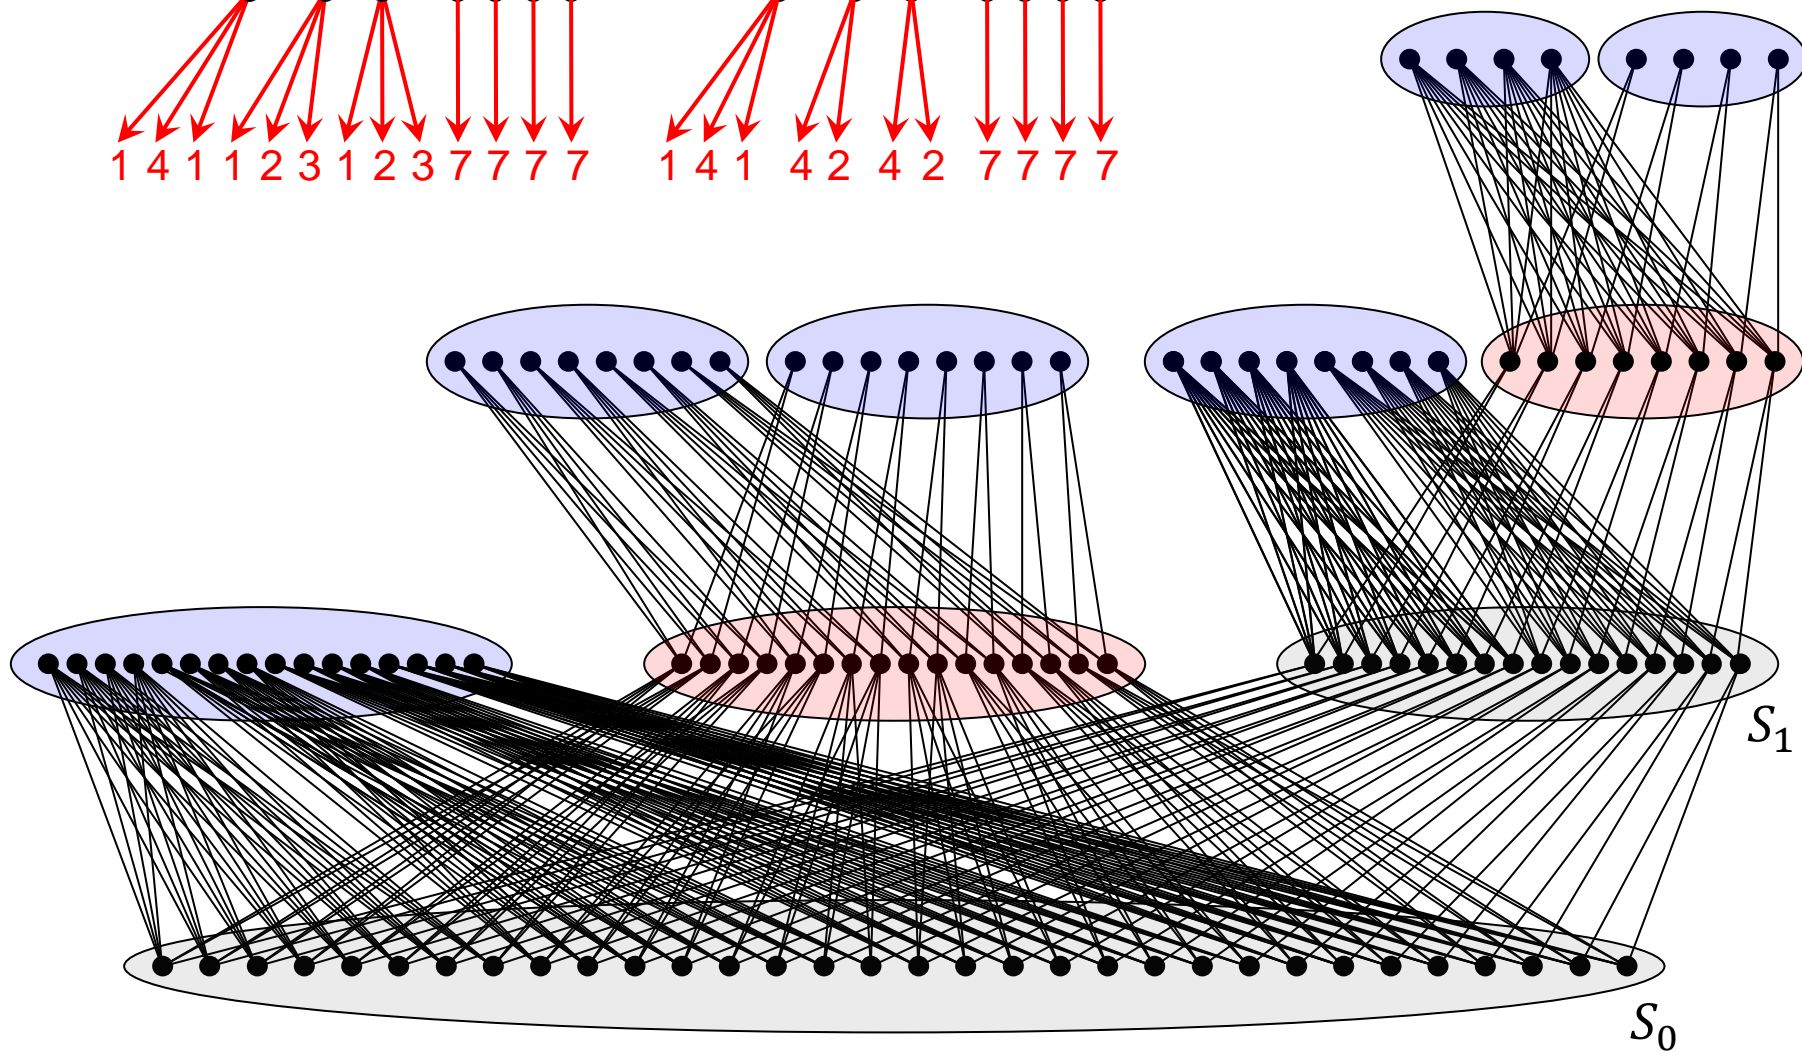

- Given the following bipartite graph with $|S_0| = \delta |S_1|$
- The MVC is just all the nodes in S_1
- Distributed Algorithm...




$N_2(\text{node in } S_0)$




$N_2(\text{node in } S_1)$




$N_2(\text{node in } S_0)$



$N_2(\text{node in } S_1)$



Graph is "symmetric", yet highly non-regular!



Lower Bound: Results

- We can show that for $\epsilon > 0$, in t time, the approximation ratio is at least

$$\Omega\left(n^{\frac{1/4-\epsilon}{t^2}}\right) \quad \text{and} \quad \Omega\left(\Delta^{\frac{1-\epsilon}{t+1}}\right)$$

- Constant approximation needs at least $\Omega(\log \Delta)$ and $\Omega(\sqrt{\log n})$ time.
- Polylog approximation $\Omega(\log \Delta / \log \log \Delta)$ and $\Omega(\sqrt{\log n / \log \log n})$.

Lower Bound: Results

- We can show that for $\epsilon > 0$, in t time, the approximation ratio is at least

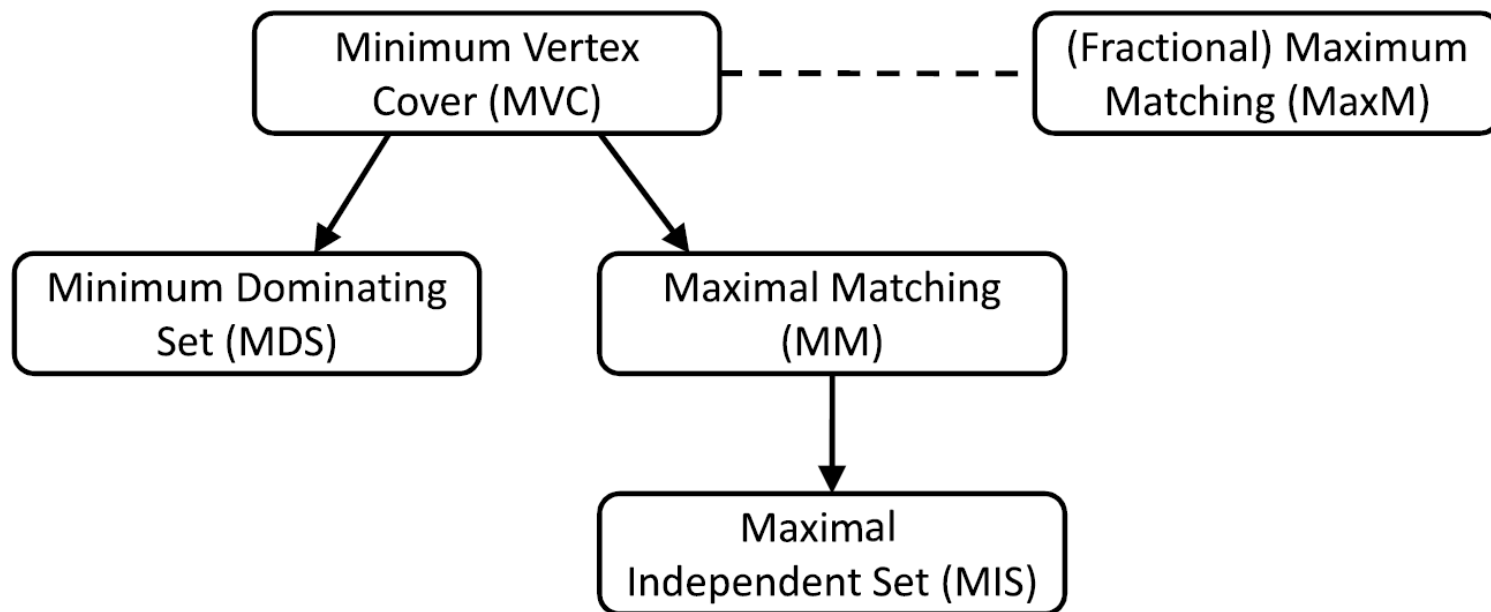
$$\Omega\left(n^{\frac{1/4-\epsilon}{t^2}}\right) \quad \text{and} \quad \Omega\left(\Delta^{\frac{1-\epsilon}{t+1}}\right)$$

tight for MVC

- Constant approximation needs at least $\Omega(\log \Delta)$ and $\Omega(\sqrt{\log n})$ time.
- Polylog approximation $\Omega(\log \Delta / \log \log \Delta)$ and $\Omega(\sqrt{\log n / \log \log n})$.

Lower Bound: Reductions

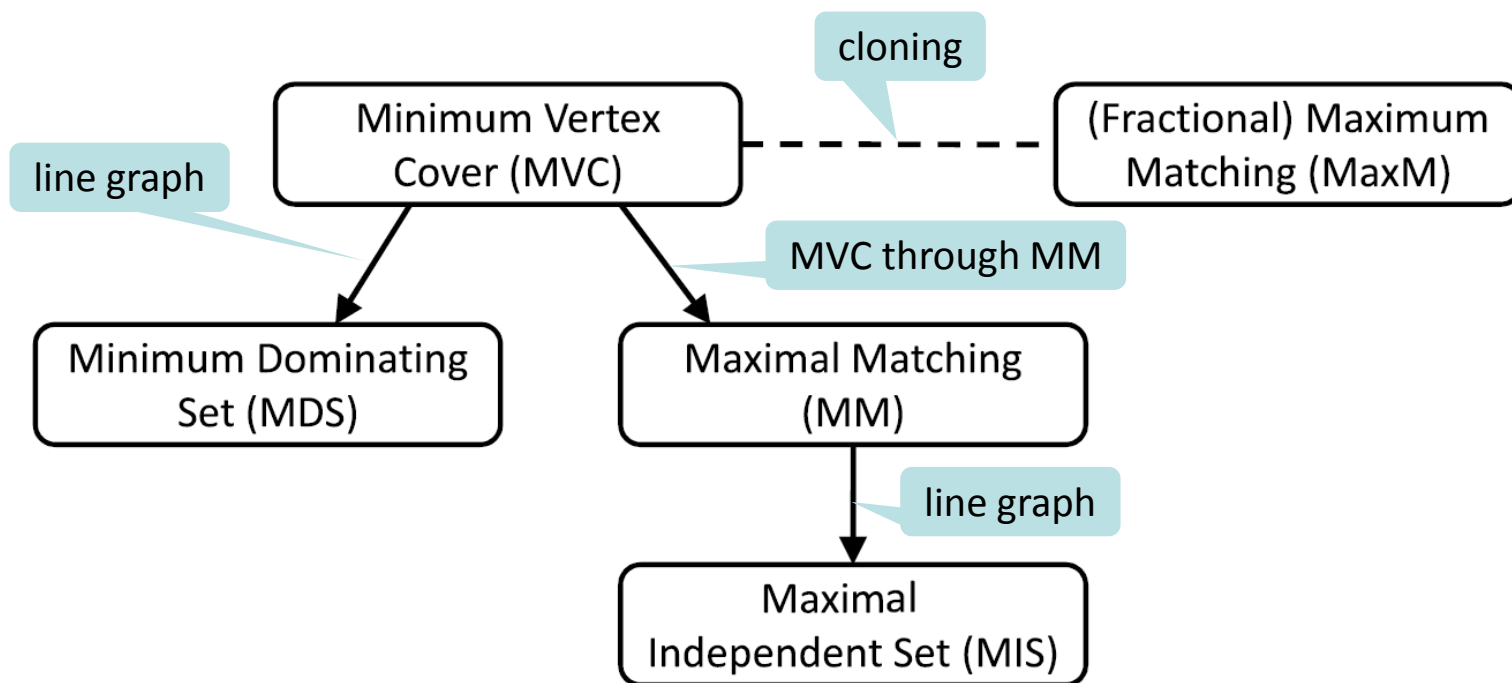
- Many “local looking” problems need non-trivial t , in other words, the bounds $\Omega(\log \Delta)$ and $\Omega(\sqrt{\log n})$ hold for a variety of classic problems.



[Kuhn, Moscibroda, W, journal version in submission]


Lower Bound: Reductions

- Many “local looking” problems need non-trivial t , in other words, the bounds $\Omega(\log \Delta)$ and $\Omega(\sqrt{\log n})$ hold for a variety of classic problems.




[Kuhn, Moscibroda, W, journal version in submission]

Olympics!



Distributed Complexity Classification



e.g., dominating set approximation in planar graphs

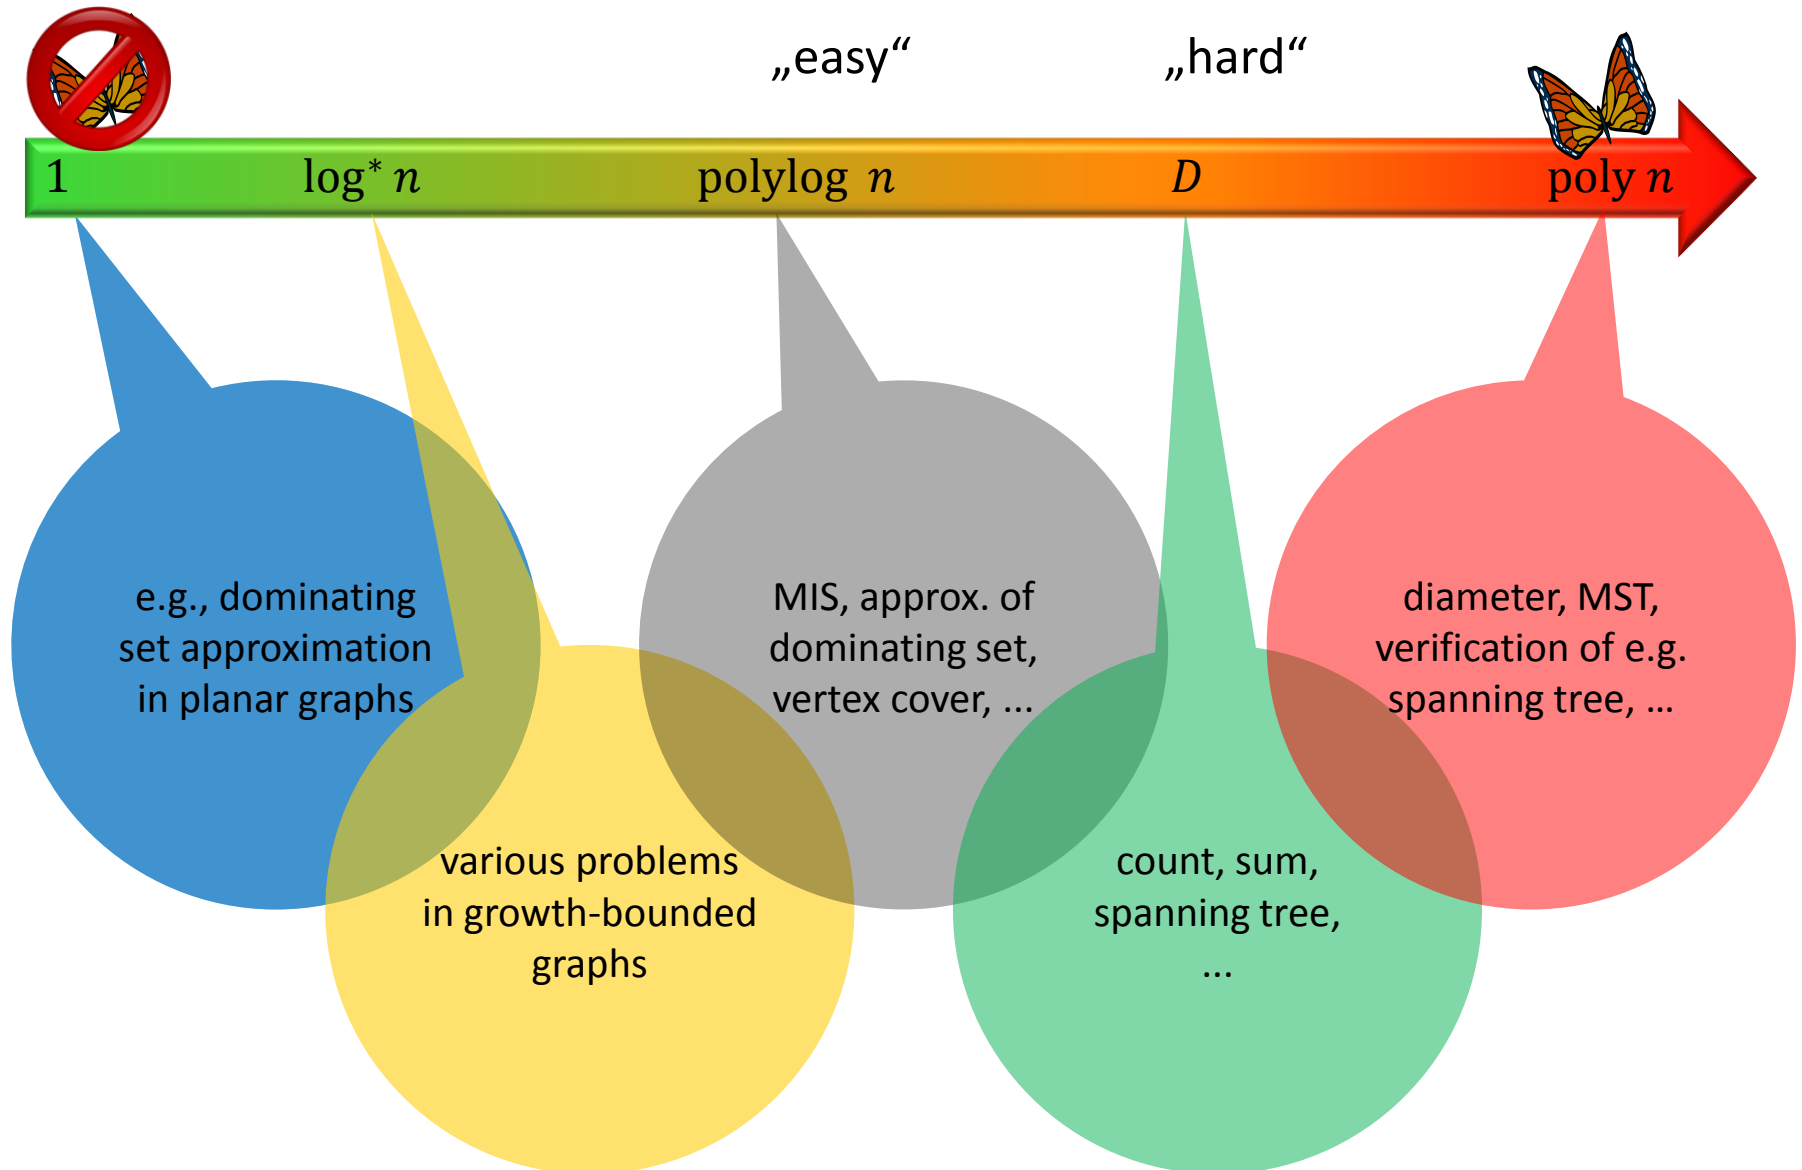
various problems in growth-bounded graphs

MIS, approx. of dominating set, vertex cover, ...

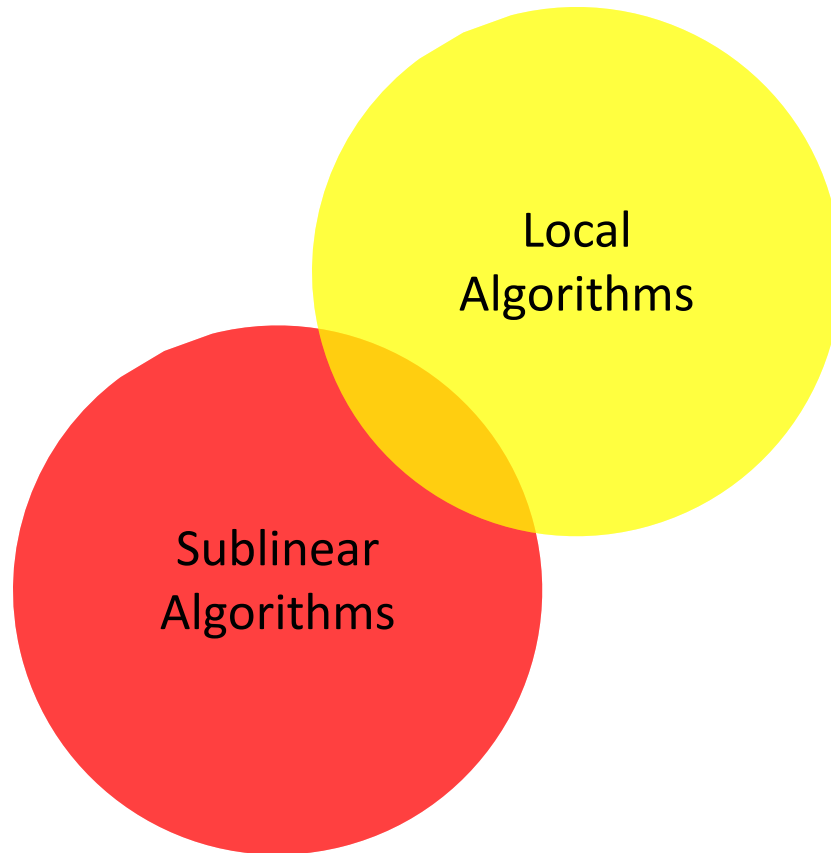
count, sum, spanning tree, ...

diameter, MST, verification of e.g. spanning tree, ...


Distributed Complexity Classification



Locality



Locality is Everywhere!



Self-
Assembly

Applications
e.g. Multi-Core

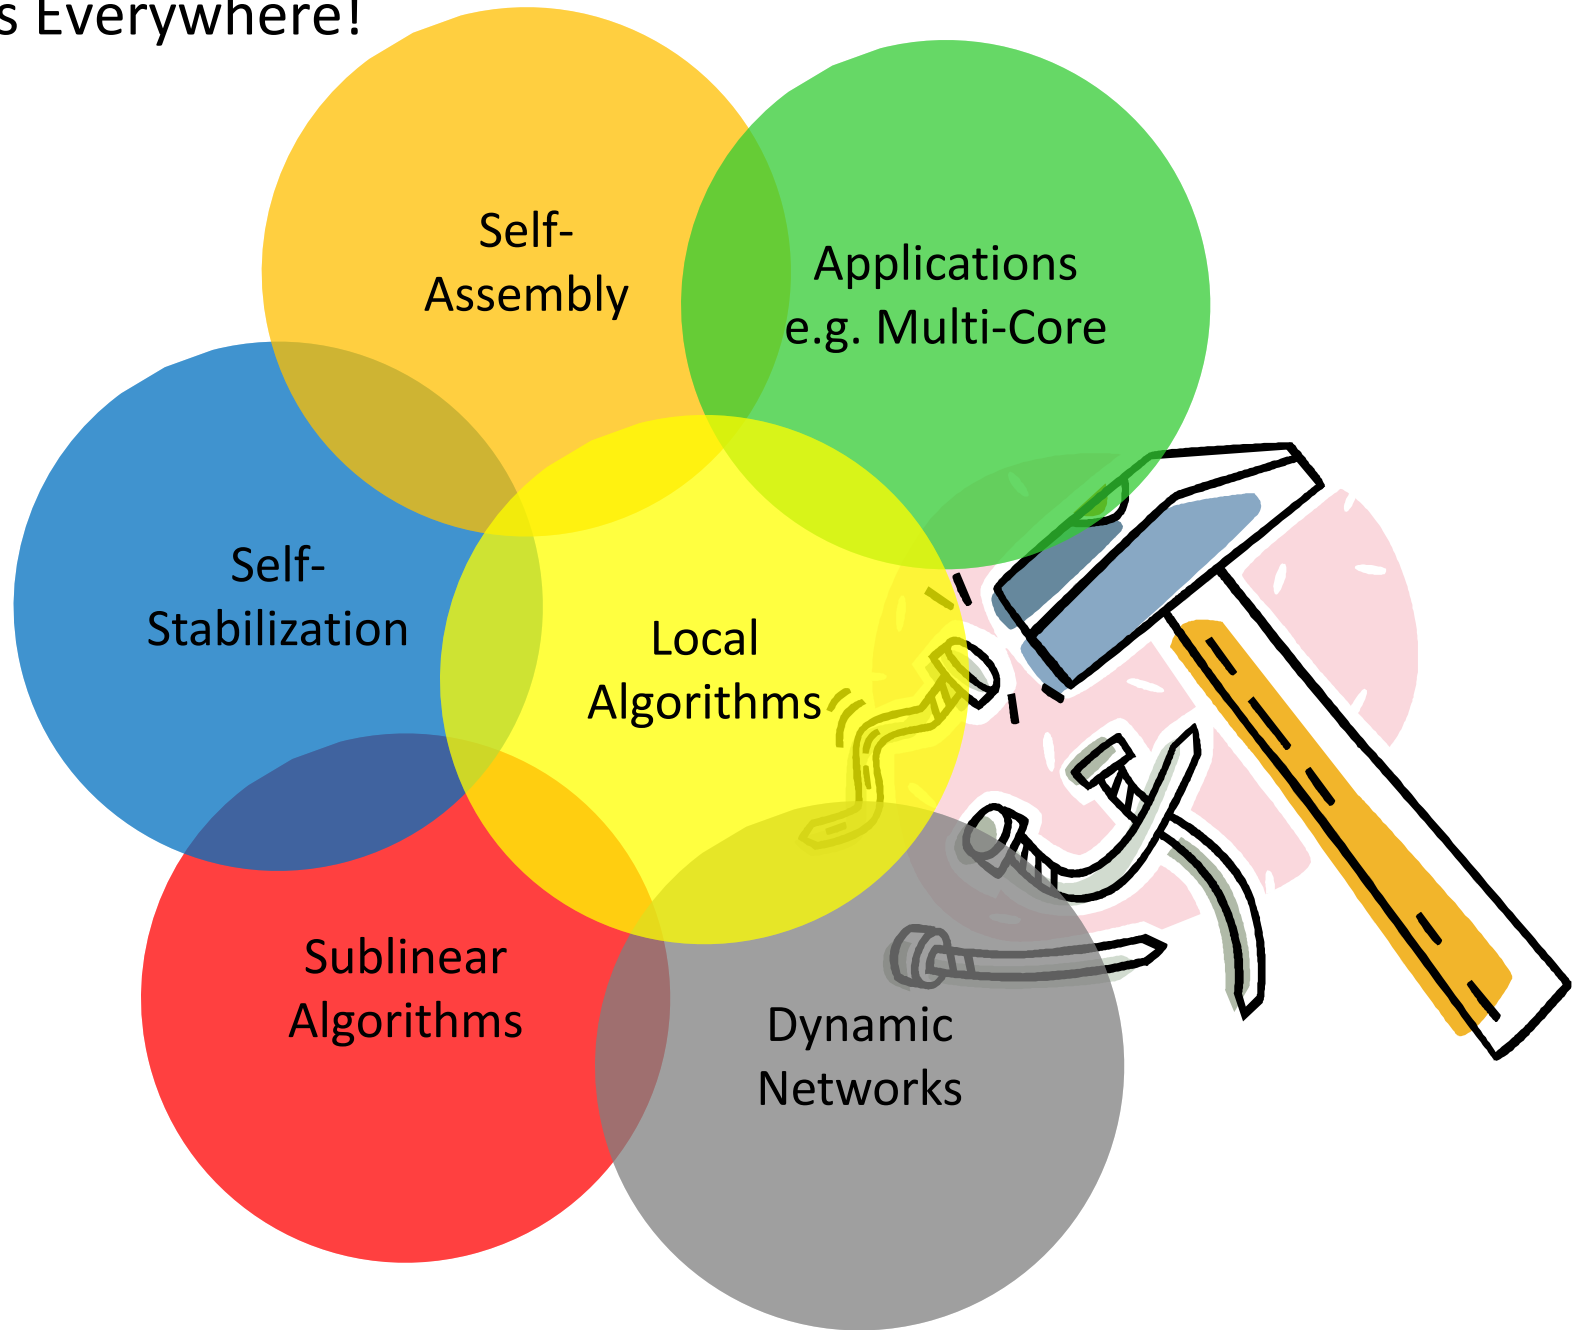
Self-
Stabilization


Local
Algorithms


Sublinear
Algorithms

Dynamic
Networks


Locality is Everywhere!








each round:
every node:
1. send msgs
2. rcv msgs
3. compute

The image features a network of nodes and edges. The nodes are represented by semi-transparent, glowing cyan shapes that overlap and connect to form a complex web. Each node contains a small, dark, irregularly shaped object, possibly representing a seed or a particle. The edges are thin, glowing cyan lines that connect the nodes. The background is a solid black color. In the lower right quadrant, there is a yellow rectangular box containing a list of operations.

each round:
every node:
1. send msgs
2. rcv msgs
3. compute


Maximal Independent Set (MIS)

- Given a **network** with n nodes, nodes have **unique IDs**.
- Find a **Maximal Independent Set (MIS)**
 - a non-extendable set of pair-wise non-adjacent nodes




Maximal Independent Set (MIS)

- Given a **network** with n nodes, nodes have **unique IDs**.
- Find a **Maximal Independent Set (MIS)**
 - a non-extendable set of pair-wise non-adjacent nodes



Maximal Independent Set (MIS)

- Given a **network** with n nodes, nodes have **unique IDs**.
- Find a **Maximal Independent Set (MIS)**
 - a non-extendable set of pair-wise non-adjacent nodes




```
given: id, degree
synchronized while (true) {
  p = 1 / (2 * degree);
  if (random value between 0 and 1 < p) {
    transmit "(degree, id)";
    ...
  }
}
```




```
given: id, degree
synchronized while (true) {
  p = 1 / (2 * degree);
  if (random value between 0 and 1 < p) {
    transmit "(degree, id)";
    ...
  }
}
```

?!








Distributed Computing
Without Computing!



each round:
every node:
1. send msgs
2. rcv msgs
3. compute



Nano-Magnetic Computing





Stone Age

Distributed Computing

nFSM: networked Finite State Machine

- Every node is the **same finite state machine**, e.g. no IDs
- Apart from their state, nodes **cannot store** anything
- Nodes **know nothing about the network**, including e.g. their degree
- Nodes **cannot explicitly send messages** to selected neighbors, i.e. nodes can only implicitly communicate by changing their state
- Operation is **asynchronous**
- **Randomized** next state okay, as long as constant number
- Nodes **cannot compute**, e.g. cannot count





One, Two, Many Principle

Piraha




Walpiri


One, Two, Many Principle


- Not okay
 - while ($k < \log n$) {
 - At least half of neighbors in state s ?
 - More neighbors in state s than in state t ?
- Okay
 - No neighbor in state s ?
 - Some neighbor in state s ?
 - At most two neighbors in state s ?

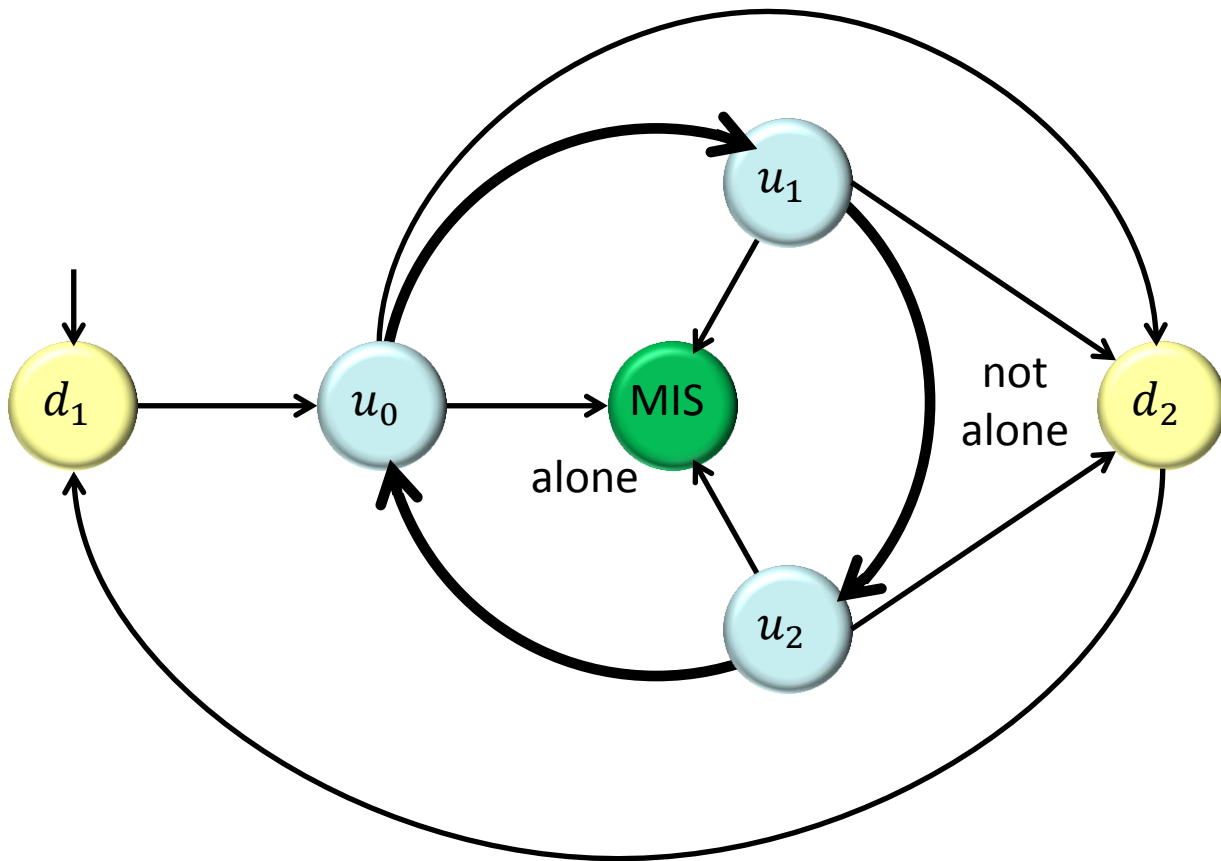



PRIMITIVE CULTURES DEVELOP
SESAME STREET.










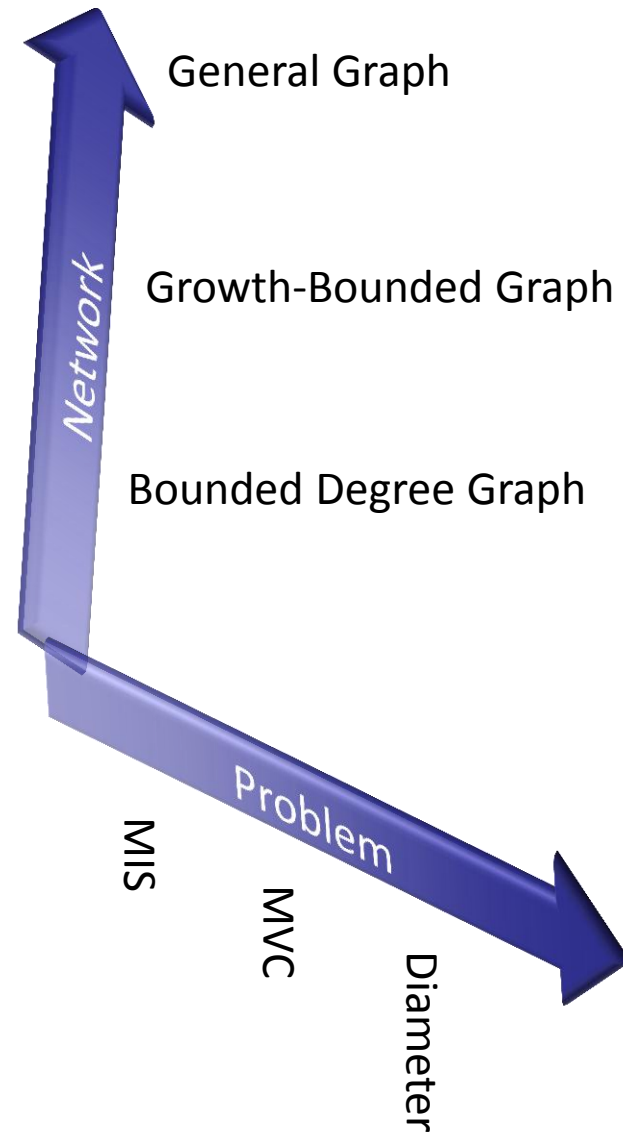


nFSM solves MIS whp in time $O(\log^2 n)$



[Emek, Smula, W, in submission]

Overview



$O(1)$ -APX,
 $O(1)$ -time

$w(n)$ -APX
 \log^* -time

Series-parallel
→ planar
↑
trees

planar
proj.
plane

planar
2-fold
cover

(bounded tree-w.)

triangle-free

some forbidden ind. subgr.

no $K_{3,3}$

no $K_{3,5}$

some
forbidden
minor

no K_5

sparse

sparse,
 d_1, d_2, d_3

trees

bounded arb.

bound
indep.
dom. p.

claw-free
line graph
 $f(n)$ -reg.

d-regular

sparse,
 d_1, d_2

bounded
degree

growth-
bounded

$O(1)$ -APX
 \log^* -time


cliques

bounded
diam.




gb +
sparse



Overview



Summary



Thank You!

Questions & Comments?



Thanks to my co-authors

Yuval Emek

Silvio Frischknecht

Stephan Holzer

Fabian Kuhn

Thomas Moscibroda

Jasmin Smula

www.disco.ethz.ch