UNI FREIBURG

On the Complexity of Local Distributed Graph Problems

Fabian Kuhn
University of Freiburg
Germany

joint work with
Mohsen Ghaffari (ETH Zurich) & Yannic Maus (U. Freiburg)

The LOCAL Model

- Synchronous message passing on a graph $G = (V, E)$

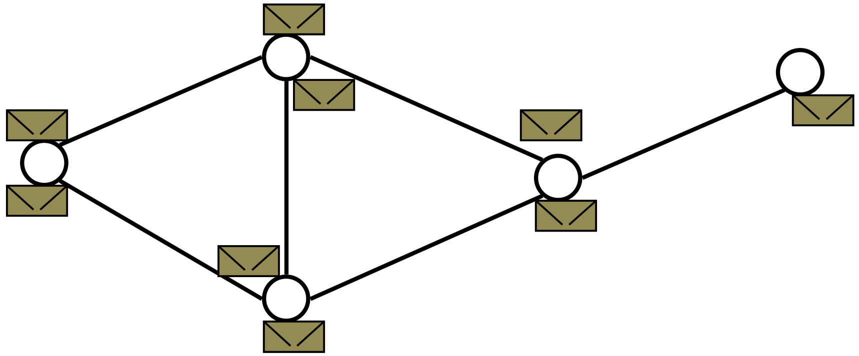

  - message size / internal computations per round are unbounded
  - each node has a unique $O(\log n)$-bit ID
  - time complexity = number of rounds

- Model was first studied by Nati Linial [FOCS ’87; SICOMP ‘92]
  - Upper and lower bounds for distributed graph coloring

Distributed Graph Problems

Distributed Graph Problem on $G = (V, E)$

- each node $v \in V$ has an input x_v
- each node $v \in V$ needs to compute an output y_v
- problem defined by pairs of valid input / output vectors

Classic Distributed Graph Problems

- Distributed Graph Coloring

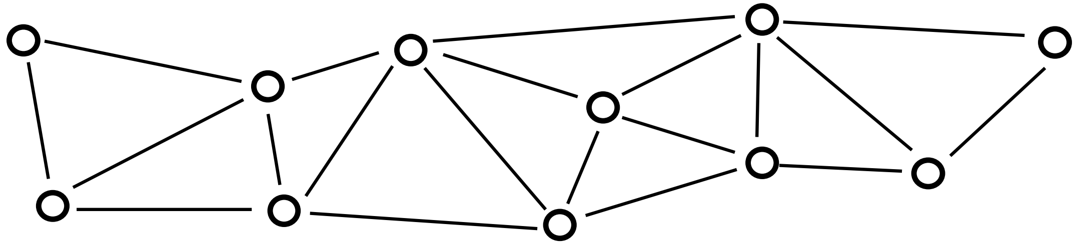

- typical goal: $\Delta + 1$ colors (Δ: max. degree of G)
- sequential greedy algorithm colors with $\leq \Delta + 1$ colors

Distributed Graph Problems

Classic Distributed Graph Problems

- Maximal Independent Set (MIS)

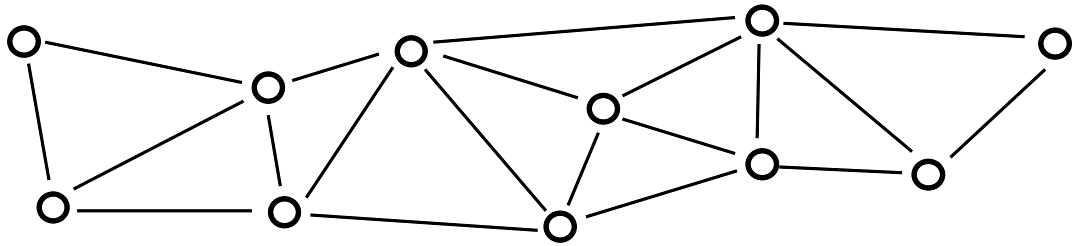

- Maximal Matching

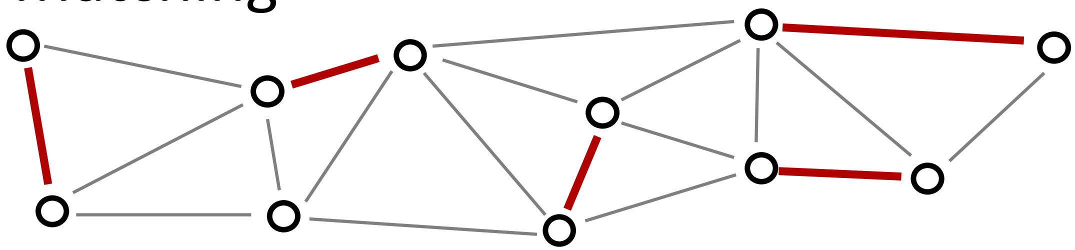

- Minimum dominating set / vertex cover approximation
- Matching / independent set approximation
- Many graph coloring variants
- …

$(\Delta + 1)$-Coloring: Current State

Lower Bound

- $\Omega(\log^* n)$ rounds needed even on the ring [Linial '87]

Efficient Randomized Algorithms

- Simple randomized $O(\log n)$-time algorithms

[Luby '86; Alon, Isreali, Itai '86; Linial '87]

- Best current upper bound: $O\left(\sqrt{\log \Delta}\right) + 2^{O\left(\sqrt{\log \log n}\right)}$

[Harris, Schneider, Su '16]

Best Deterministic Algorithm

- Based on network decomposition: $2^{O\left(\sqrt{\log n}\right)}$

[Panconesi, Srinivasan '92]

Maximal Independent Set: Current State

Lower Bound

- $\Omega\left(\sqrt{\log n / \log\log n}\right)$ rounds needed

[Kuhn, Moscibroda, Wattenhofer '04]

Efficient Randomized Algorithms

- Simple randomized $O(\log n)$-time algorithms

[Luby '86; Alon, Isreali, Itai '86]

- Best current upper bound: $O(\log \Delta) + 2^{O\left(\sqrt{\log\log n}\right)}$

[Ghaffari '16]

Best Deterministic Algorithm

- Based on network decomposition: $2^{O\left(\sqrt{\log n}\right)}$

[Panconesi, Srinivasan '92]

Exponential Gap

Current State for $(\Delta + 1)$-coloring & MIS

- $O(\log n)$-time **randomized** algorithms
- Best **deterministic** algorithm: $2^{O(\sqrt{\log n})}$

There is an exponential separation between the best randomized and deterministic algorithms

- The same is true for many other distributed graph problems
  - Dominating set / independent set approximation
  - $(2\Delta - 1)$-edge coloring
  - Network decomposition
  - …
- A major open problem already mentioned in [Linial '87]

Challenges in the LOCAL Model

Locality

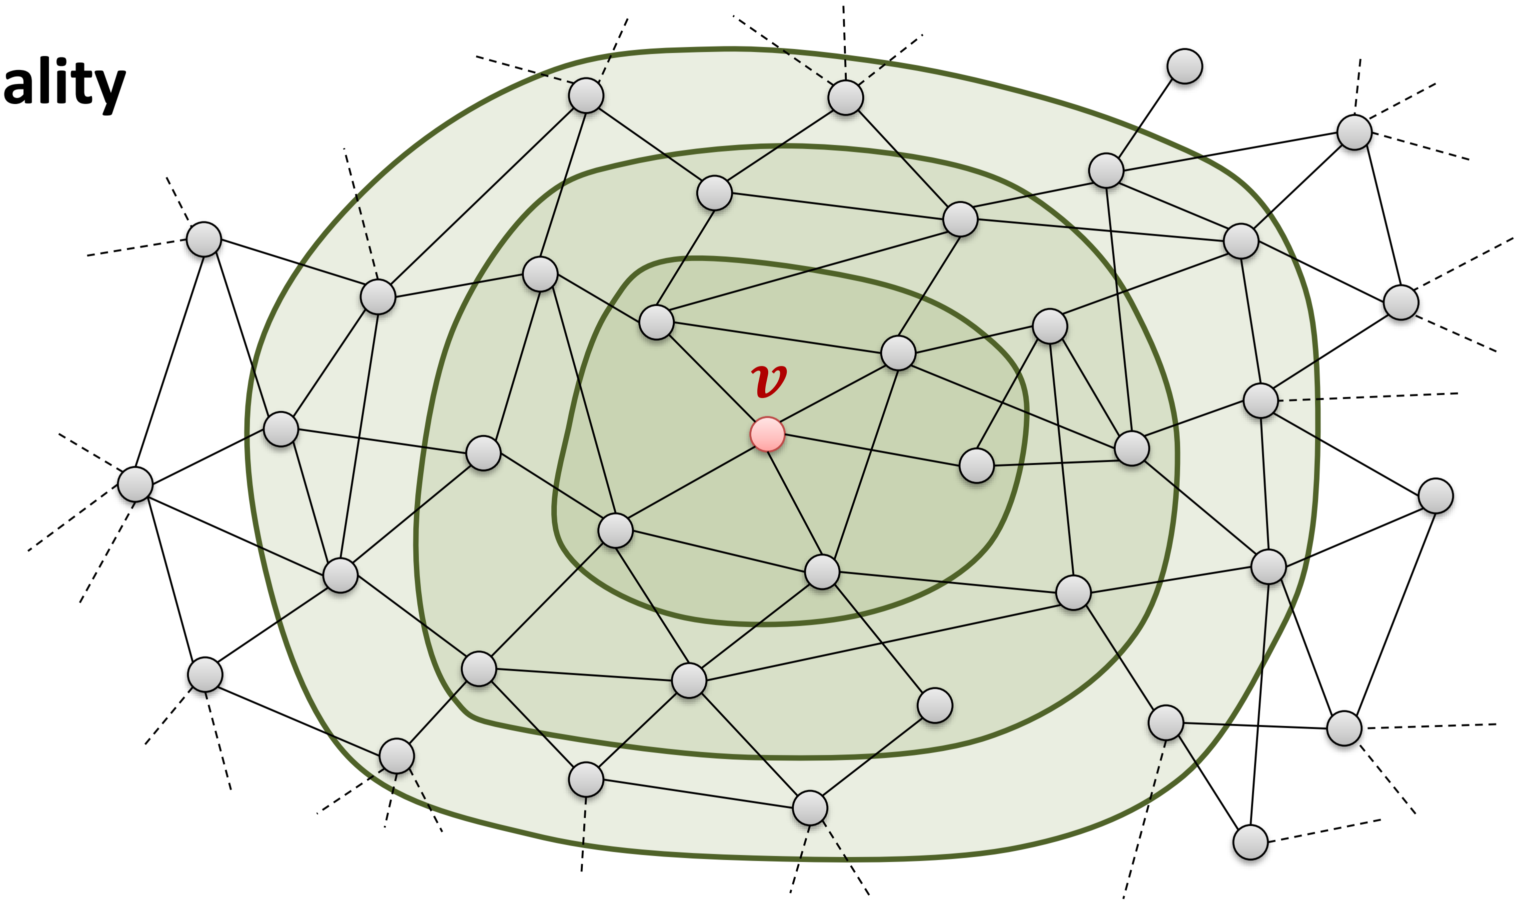

r-Round Algorithm:

Each node computes its output as a function of the initial state of its r-neighborhood

Challenges in the LOCAL Model

Local Coordination / Symmetry Breaking

- Nearby (symmetric) nodes need to output different values
  - Neighboring nodes need different colors
  - No adjacent nodes in MIS, each node not in MIS has neighbor in MIS
  - …

- Nodes decide in parallel based on their r-neighborhoods

Main Challenge:
Locally coordinate among nearby nodes

- Randomization naturally helps
  - E.g., choose random color, keep if no conflict with neighbors

Sequential LOCAL Model

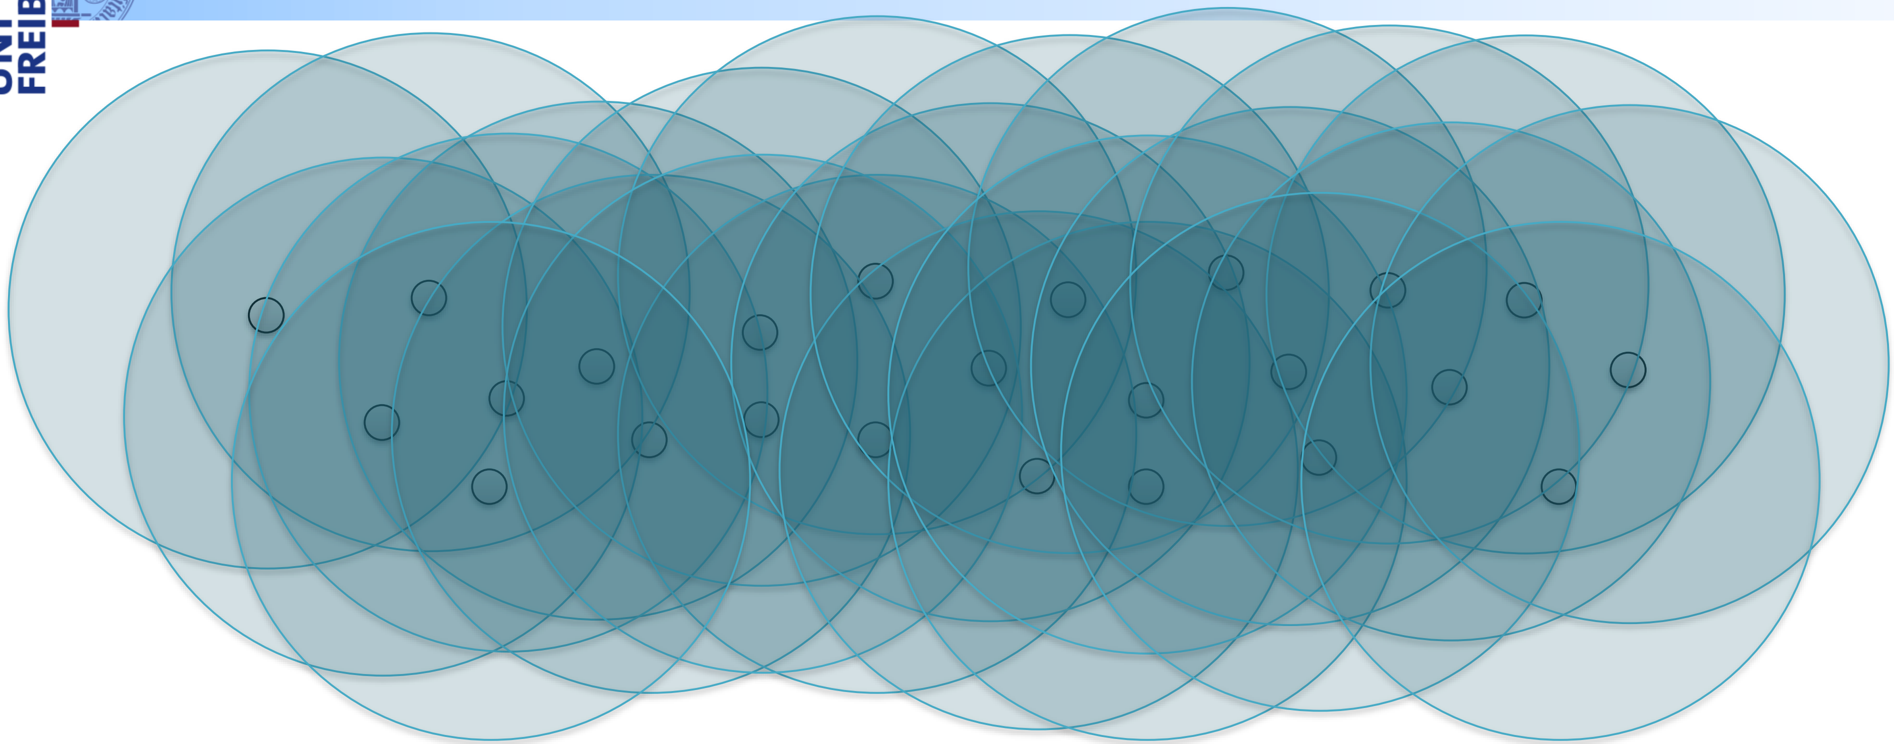

SLOCAL Model

- locality parameter $r(n)$
- sequentially go over all nodes
- compute output of a node based on the current state of its $r(n)$-neighborhood

SLOCAL Model

SLOCAL model is much more powerful than LOCAL model

- $(\Delta + 1)$-coloring and MIS can easily be solved with locality 1
  - The sequential greedy algorithm is an SLOCAL-algorithm
  - The output a node v only depends on the outputs of neighbors which were processed before v
- SLOCAL is a generalization of sequential greedy algorithms
  - if for each node, one only looks at previous nearby nodes

Network Decomposition

Introduced in [Awerbuch, Goldberg, Luby, Plotkin '89]

Definition: $(d(n), c(n))$-decomposition of $G = (V, E)$
- Partition of V into clusters of diameter $\leq d(n)$
- Coloring of cluster graph with $c(n)$ colors

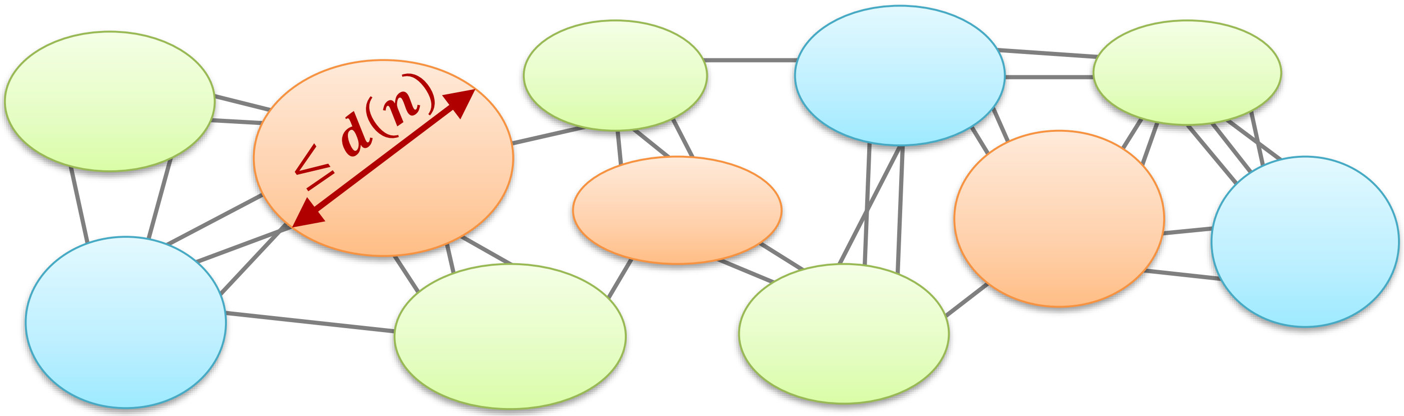

[AGLP '89]: Can be computed deterministically in $2^{O\left(\sqrt{\log n \log\log n}\right)}$ rounds for $d(n) = c(n) = 2^{O\left(\sqrt{\log n \log\log n}\right)}$

Distributed Comp. & Network Decomp.

Claim: Given a $(d(n), c(n))$-decomp. of $G = (V, E)$, one can compute a $(\Delta + 1)$-coloring or MIS in $c(n) \cdot d(n)$ rounds.

- Iterate through $c(n)$ colors
- For each cluster, compute solution in $d(n)$ rounds

- This works for all SLOCAL algorithms with locality 1

Lemma:
Given a $(d(n), c(n))$-decomp. of $G^{r(n)}$, one can run any SLOCAL alg. with locality $\leq r(n)$ in $r(n) \cdot c(n) \cdot d(n)$ rounds.

– $G^{r(n)}$: edge between any two nodes at dist. $\leq r(n)$ in G

Network Decomposition Algorithms

Existential Result [Awerbuch, Peleg '90], [Linial, Saks '93]

- Every graph G has an $\big(O(\log n), O(\log n)\big)$-decomposition
  - clusters of diameter $O(\log n)$, clusters colored with $O(\log n)$ colors

Complexity of computing $\big(O(\log n), O(\log n)\big)$-decomposition

- **Deterministic SLOCAL Model:** locality $O(\log^2 n)$
  - simple adaptation of alg. by [Awerbuch, Peleg '90], [Linial, Saks '93]
- **Deterministic LOCAL Model:** $2^{O(\sqrt{\log n})}$ rounds
  - combination of algorithms by [Panconesi, Srinivasan '92] and [Awerbuch, Berger, Cowen, Peleg '96]
- **Randomized LOCAL Model:** $O(\log^2 n)$ rounds
  - randomized distributed algorithm by [Linial, Saks '93]

Complexity Classes

LOCAL$(t(n))$

- graph problems that can be solved deterministically in $t(n)$ rounds in the LOCAL model

SLOCAL$(t(n))$

- graph problems that can be solved deterministically with locality $t(n)$ in the SLOCAL model
  - MIS, $(\Delta + 1)$-coloring $\in$ SLOCAL(1)

P-LOCAL $:=$ **LOCAL$(\mathbf{poly} \log n)$**

P-SLOCAL $:=$ **SLOCAL$(\mathbf{poly} \log n)$**

Randomized classes: RLOCAL, RSLOCAL, P-RLOCAL, P-RSLOCAL

Exponential Separation Revisited

Basic Facts

- LOCAL$\big(t(n)\big) \subseteq$ SLOCAL$\big(t(n)\big)$
- P-LOCAL $\subseteq$ P-SLOCAL

$\big(O(\log n), O(\log n)\big)$-decomposition of $G^{\text{poly}\log(n)}$

$\Longrightarrow$ deterministic poly $\log n$-round algorithm for any problem in P-SLOCAL in the LOCAL model.

- **P-SLOCAL $\subseteq$ LOCAL$\left(2^{O\left(\sqrt{\log n}\right)}\right)$**
  - deterministic $2^{O\left(\sqrt{\log n}\right)}$-round distr. alg. for all problems in P-SLOCAL
- **P-SLOCAL $\subseteq$ P-RLOCAL**
  - randomized poly $\log n$-round distr. alg. for all problems in P-SLOCAL

Exponential Separation Revisited

All P-SLOCAL problems have deterministic $2^{O\left(\sqrt{\log n}\right)}$-round and randomized poly $\log n$-round algorithms in the LOCAL model.

Open Problem

- Is there an asymptotic separation for deterministic algorithms between P-LOCAL and P-SLOCAL in the LOCAL model?

$$\textbf{P-LOCAL} \stackrel{?}{=} \textbf{P-SLOCAL}$$

- There is no separation for rand. alg.: P-RLOCAL = P-RSLOCAL

Are there any complete problems in P-SLOCAL?

– If $\big(O(\log n), O(\log n)\big)$-decomposition is in P-LOCAL, all problems in P-SLOCAL are in P-LOCAL.

P-SLOCAL Completeness

Local Reduction

- We say that a distr. graph problem P_1 is polylog-reducible to P_2 if a deterministic $\operatorname{poly}\log n$-round distr. algorithm for P_2 implies a deterministic $\operatorname{poly}\log n$-round distr. algorithm for P_1.

P-SLOCAL Completeness

- A problem P in P-SLOCAL is called P-SLOCAL-complete if every problem P' in P-SLOCAL is polylog-reducible to P

$\big(O(\log n), O(\log n)\big)$-decomposition is P-SLOCAL-complete

- $\big(O(\log n), O(\log n)\big)$-decomp. is in SLOCAL$\left(O\left(\log^2 n\right)\right)$
- polylog-round decomp. alg. $\Longrightarrow$ polylog-round P-SLOCAL alg.

Local Splitting Problem

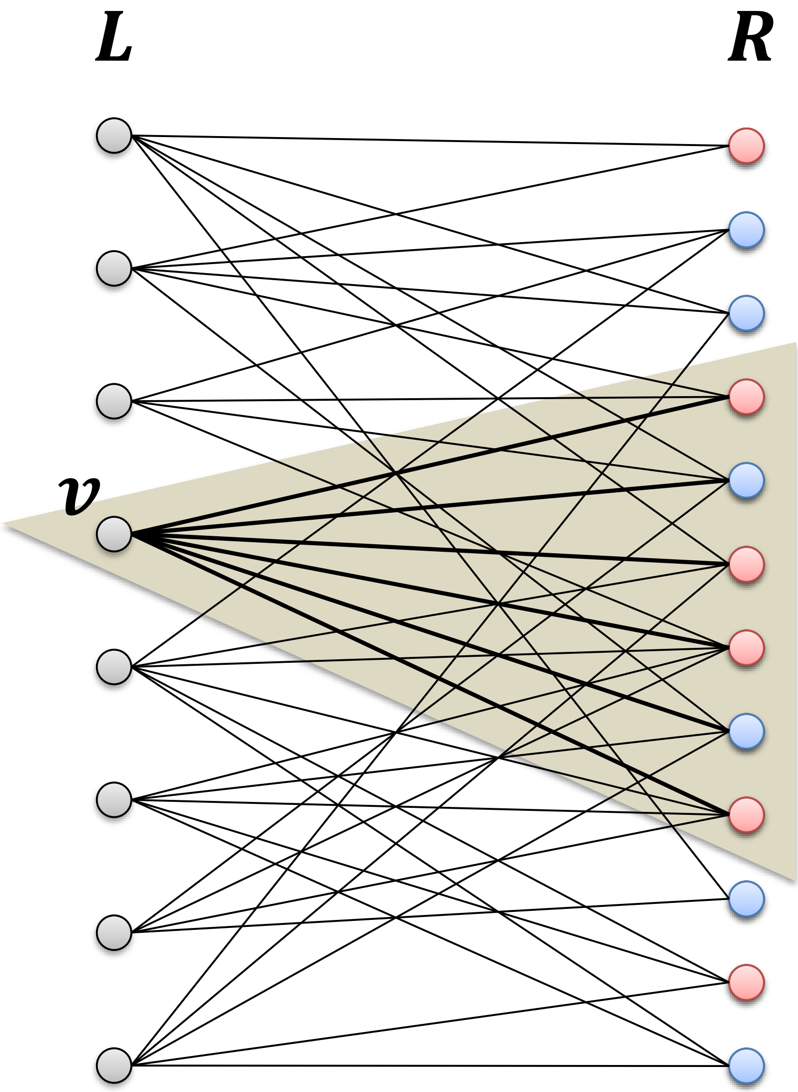

λ-Local Splitting $(\lambda \in [0, 1/2])$:

Every $v \in L$ has $\geq \lfloor \lambda \deg(v) \rfloor$ neighbors of each color.

Weak Local Splitting:

Every $v \in L$ has at least one neighbor of each color

Trivial Randomized Solution:

Independently color red/blue with probability $^1/_2$

- works w.h.p. if all degrees in L are $\Omega(\log n)$ and if λ is not too close to $1/2$

Local Splitting is P-SLOCAL-Complete

Theorem: If all nodes in L have degree $\Omega(\log^2 n)$,

a) weak local splitting is P-SLOCAL-complete and

b) λ-local splitting is P-SLOCAL-complete for any $\lambda = \frac{1}{\text{poly} \log n}$.

network decomposition

↓ polylog-reducible

conflict-free coloring

↓ polylog-reducible

λ-local splitting

↓ polylog-reducible

weak local splitting

Local Splitting is P-SLOCAL-Complete

Theorem: If all nodes in L have degree $\Omega(\log^2 n)$,

a) weak local splitting is P-SLOCAL-complete and

b) λ-local splitting is P-SLOCAL-complete for any $\lambda = \frac{1}{\text{poly}\log n}$.

- There is a 0-round randomized algorithm for both problems
- Can be seen as a rounding fractional values to integer values
  - initially, each node in R is red and blue with value ½ each

Coarsely rounding fractional values to integer values is the only obstacle to obtaining efficient (polylog-time) deterministic algorithms in the LOCAL model.

Related Work

Complexity of Local Decision Problems

[Fraigniaud, Korman, Peleg ’11]

- Studies complexity classes for distributed decision problems
- Our basic complexity classes can be seen as a generalization:

$$\mathrm{LD}\big(t(n)\big) \subset \mathrm{LOCAL}\big(t(n)\big)$$

Exponential Separation [Chang, Kopelowitz, Pettie ‘16]

- If we do not ignore log-factors, there is an exponential separation between rand. and det. alg. In the LOCAL model
- There are problems with $O(\log\log n)$-round randomized algorithms and an $\Omega(\log n)$ deterministic lower bound

Open Problems

- Is $(\Delta + 1)$-coloring / MIS P-SLOCAL-complete?

- Local splitting seems an important problem:
  Can we solve it efficiently for special cases
  - existing polylog-time maximal matching and edge coloring algorithms use variants of local splitting (split edges of each node)
  - generalizing this to 3-uniform hypergraphs would already lead to interesting results

- Simple complete problems might help to develop lower bounds or find better deterministic distributed algorithms

- Use SLOCAL model to develop new rand. distributed alg.
  - allowed us to obtain polylog-round approximation schemes for minimum dominating set / maximum independent set

Thanks!

Questions ?
Comments?