ADFOCS Lectures

- ☐ Asynchronous Crash-Prone Distributed Computing
- ☑ **Locality in Distributed Network Computing**
- ☐ Congestion-Prone Distributed Network Computing
- ☐ Other Aspects of Distributed Computing

Various Models

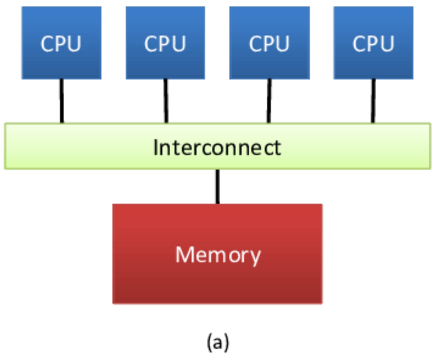

Shared Memory

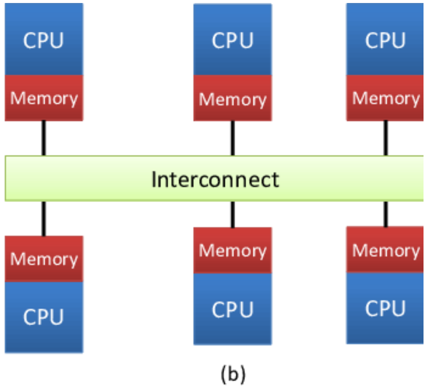

Message Passing

Synchronous

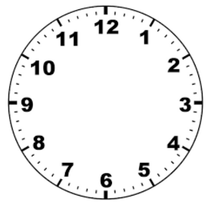

Asynchronous

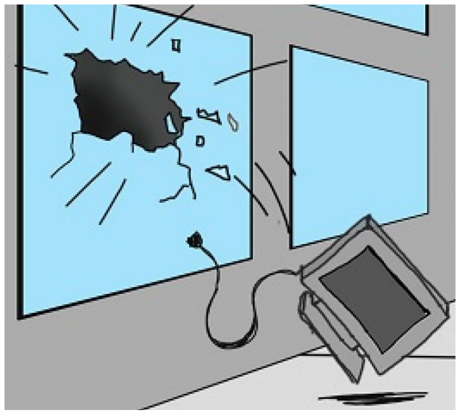

Failures: crash, transient, Byzantine, etc.

Networks

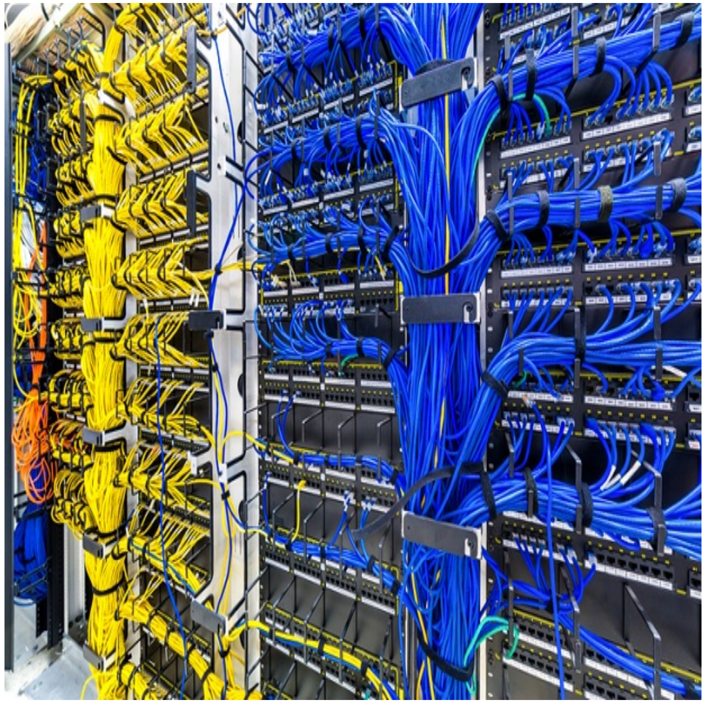

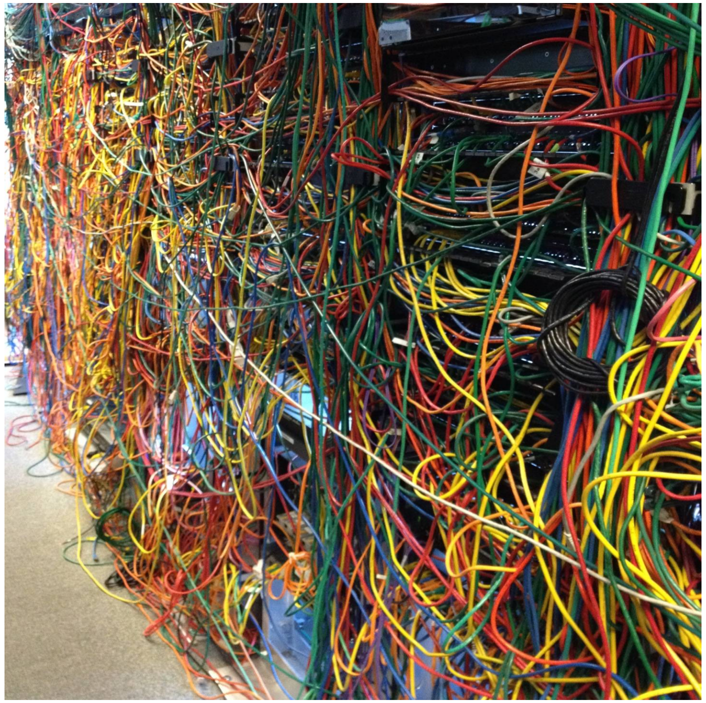

Two major technological constraints:

- Latency / Locality
- Bandwidth / Information

LOCAL Model

- Each process is located at a node of a network modeled as an n-node graph (n = #processes)
- Each process has a unique ID in $\{1,\ldots,n\}$
- Computation proceeds in synchronous rounds during which every process:

    1. Sends a message to each neighbor
    2. Receives a message from each neighbor
    3. Performs individual computation (same algorithm for all nodes)

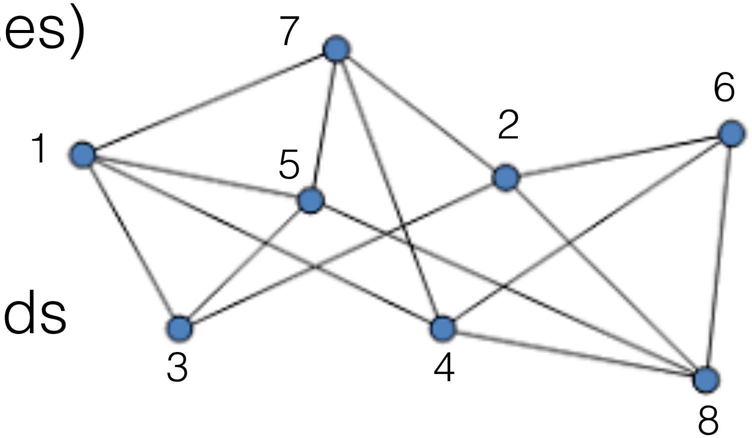

Complexity = #rounds

Lemma If a problem P can be solved in t rounds in the LOCAL model by an algorithm A, then there is a t-round algorithm B solving P in which every node proceeds in two phases:

Phase 1. Gather all data in the t-ball around it

Phase 2. Compute the solution

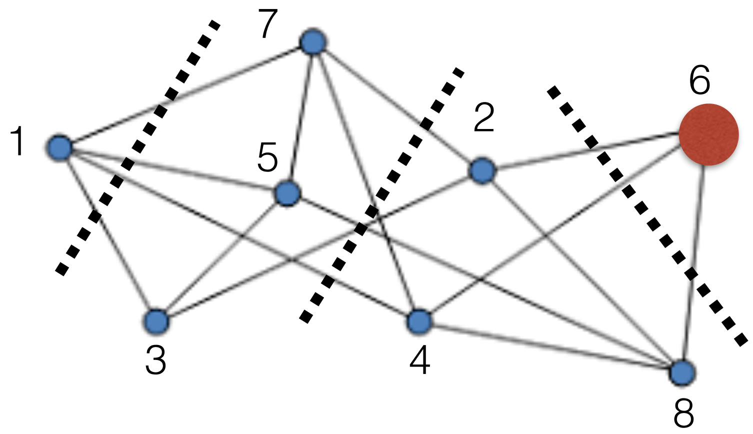

Graph problems

Vertex coloring

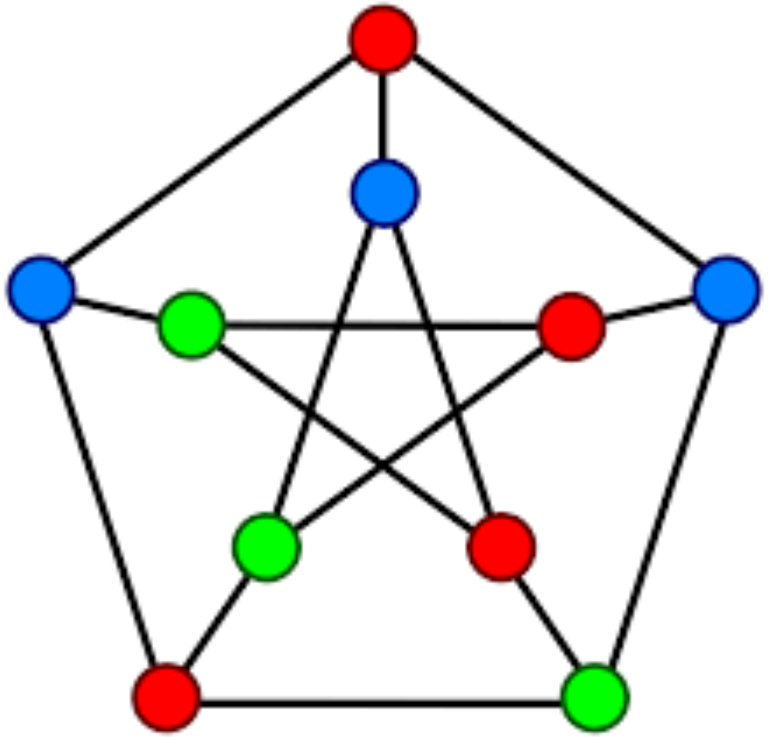

Independent set

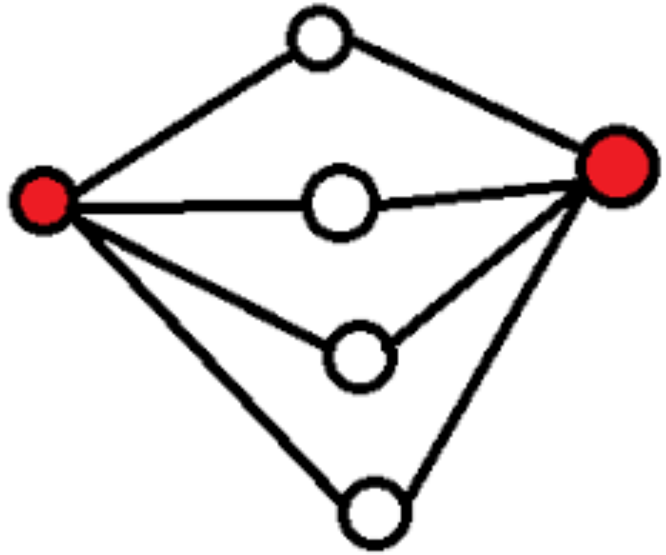

(Δ+1)-coloring

Δ = maximum node degree of the graph

$(\Delta+1)$-coloring = assign colors to nodes such that every pair of adjacent nodes are assigned different colors.

Lemma Every graph is $(\Delta+1)$-colorable

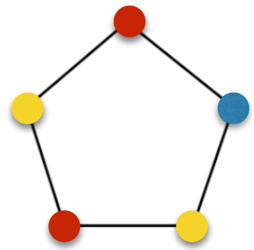

Theorem (Brooks, 1941)
Every graph G is Δ-colorable, unless G is a complete graph, or an odd cycle.

Maximal Indepent Set (MIS)

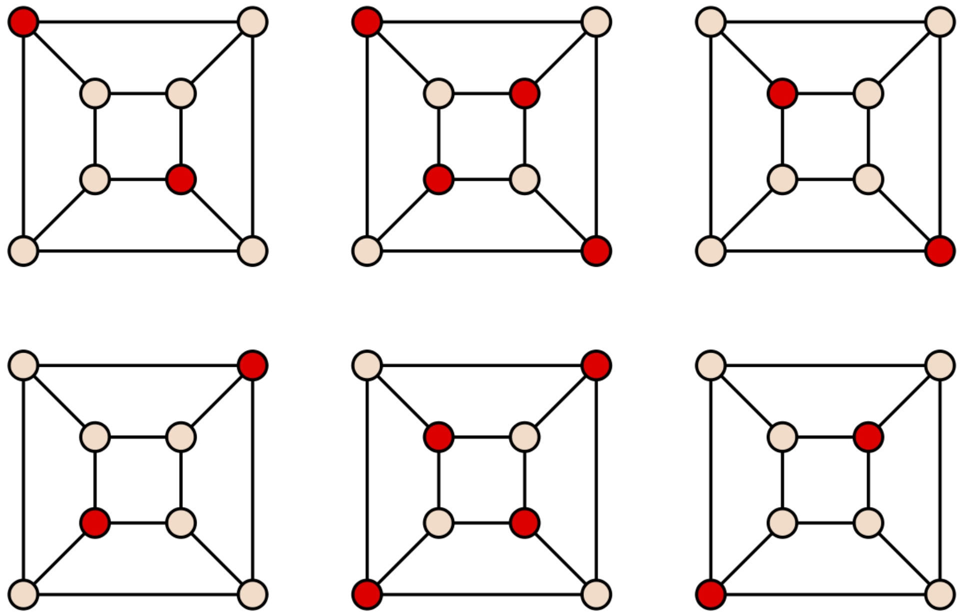

- Maximal, not maximum!

Roadmap

1. Deterministic algorithms

2. Randomized algorithms

3. Strong links between deterministic and randomized algorithms

Deterministic Algorithms

3-coloring the n-node cycle C_n

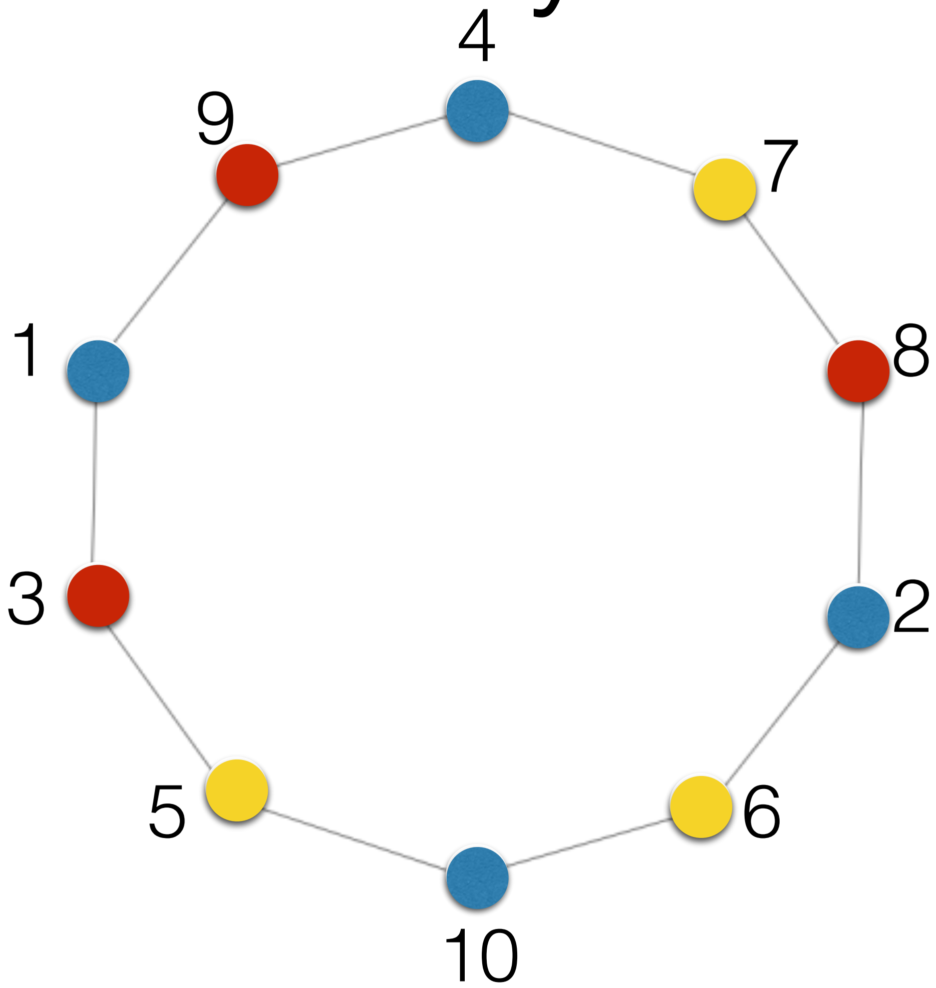

How many rounds for 3-coloring the n-node cycle?

Round complexity of 3-coloring C_n

Theorem (Cole and Vishkin, 1986) There exists an algorithm for 3-coloring C_n performing in $O(\log^* n)$ rounds.

Iterated logarithms:

- $\log^{(0)} x = \log x \qquad \log^{(k+1)} x = \log \log^{(k)} x$
- $\log^* x$ = smallest k such that $\log^{(k)} x < 1$
- $\log^* 10^{100} = 5$

Theorem (Linial, 1992) Any 3-coloring algorithm for C_n performs in $\Omega(\log^* n)$ rounds.

Dijkstra Prize 2013

Cole-Vishkin Algorithm

Initial color = ID
Express colors in binary

Assume: n is known, and
consistent sens of direction

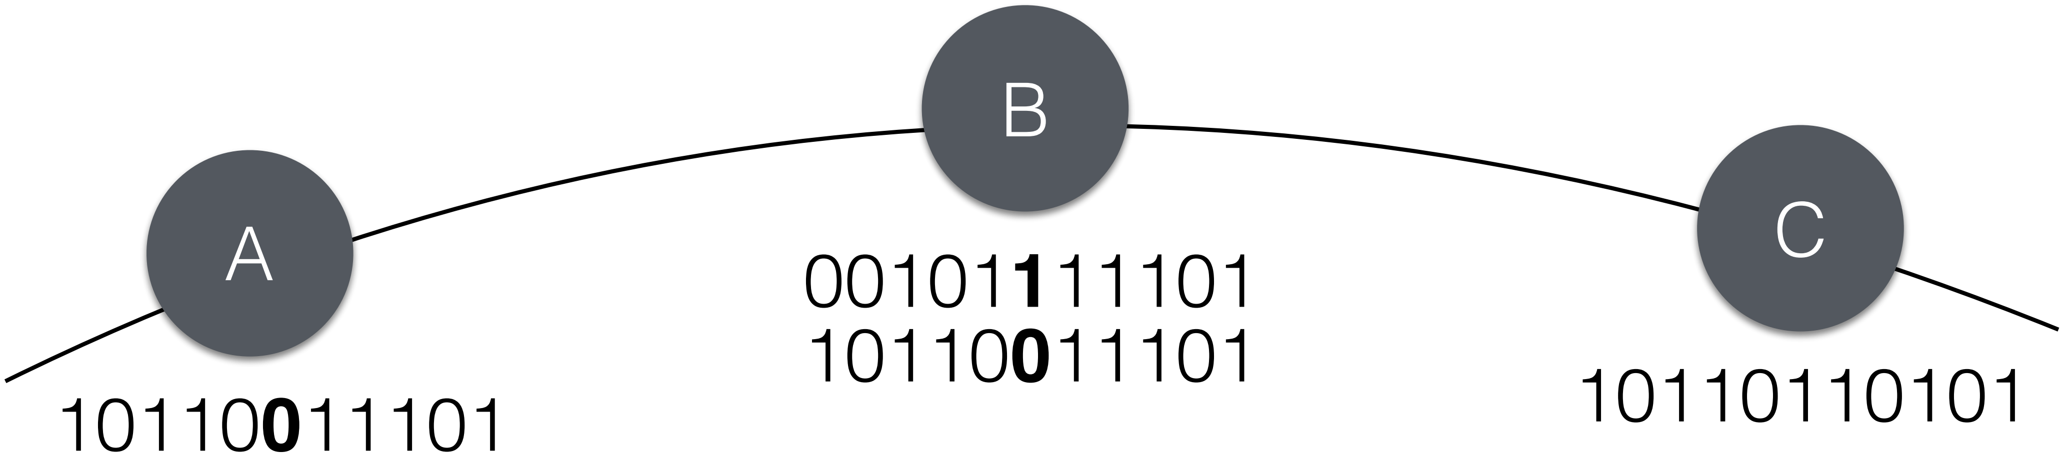

new = (position,bit) = (5,1) = 1011

(p',b') (p,b)

p ≠ p' ⇒ proper coloring

p = p' ⇒ b ≠ b' ⇒ proper coloring

Number of iterations

- k-bit colors ⇒ new colors on $\lceil \log_2 k \rceil + 1$ bits
- $\log^* n + O(1)$ rounds to reach colors on 3 bits
- 8 colors down to 3 colors in 5 rounds
- Total number of rounds = **$\log^* n + O(1)$**

❏

Speeding up the Algorithm

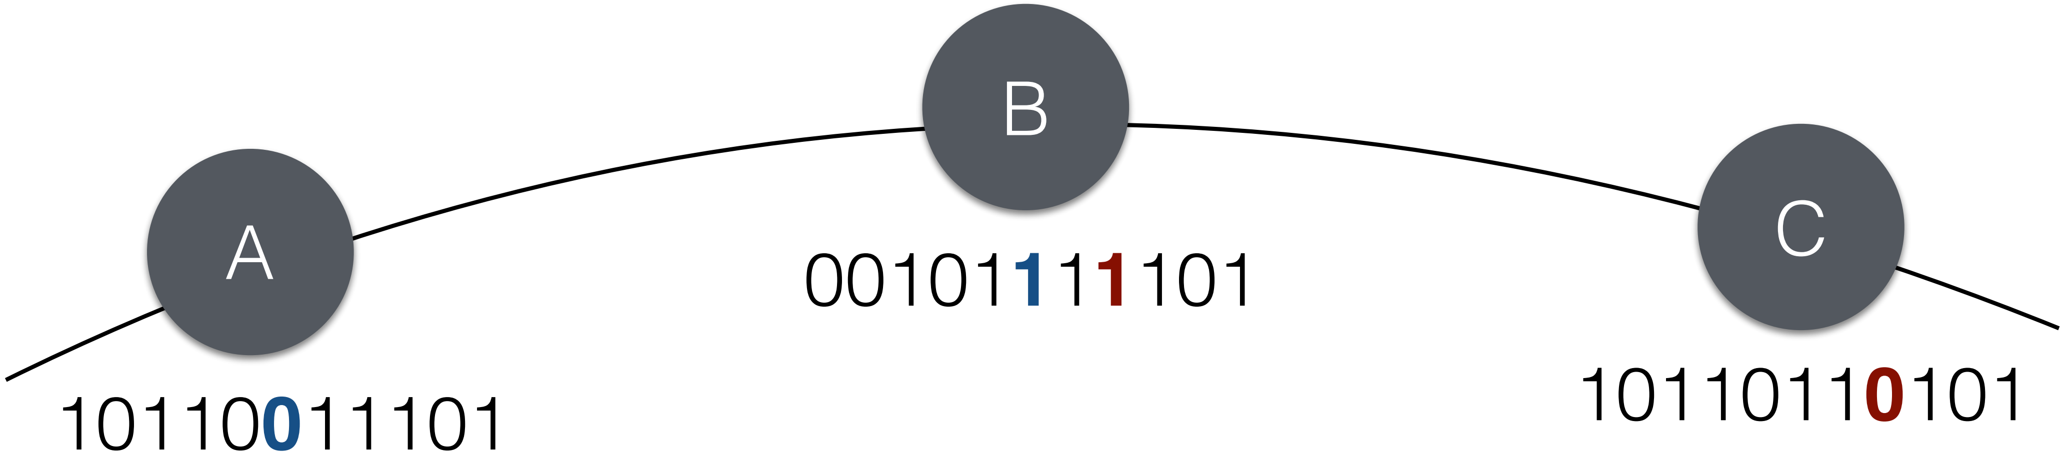

- Every node can simulate 2 rounds in just 1 round
- left round + right round ➡ implemented in 1 round
- Total number of rounds = **½ log*n + O(1)** ❑

Linial's Lower Bound

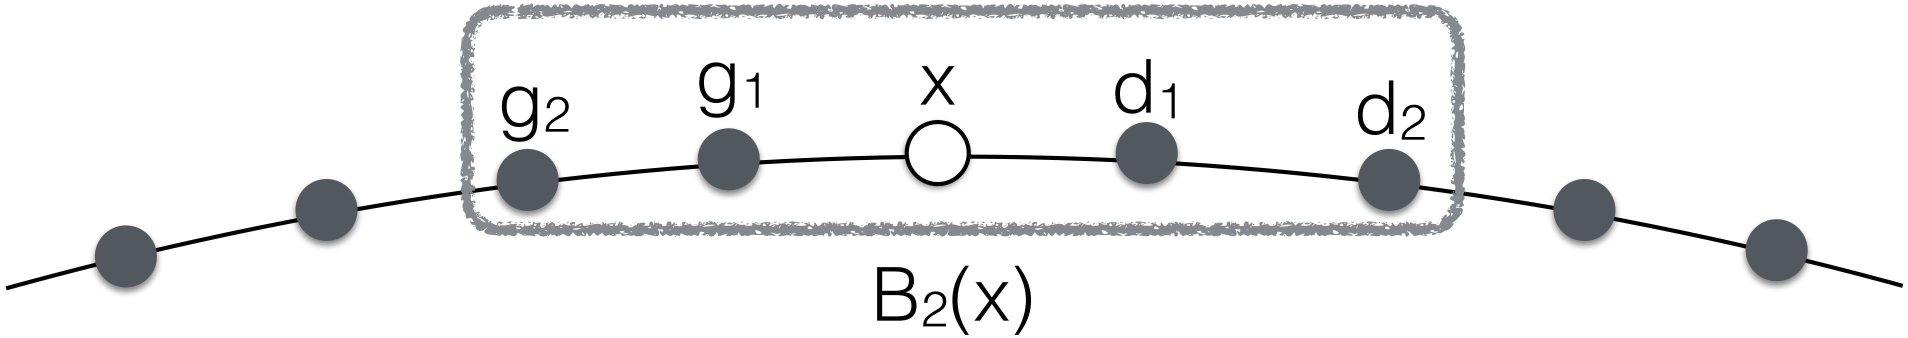

t-round algorithm ⟷ every node x decides as a function $\mathcal{A}$ applied to $B_t(x)$ where $B_t(x)=(g_t,g_{t-1},\ldots,g_1,x,d_1,\ldots,d_{t-1},d_t)$

Configuration Graph $G_{t,n}$

vertices = $\{ (g_t,\ldots,g_1,x,d_1,\ldots,d_t) \in \{1,\ldots,n\}^{2t+1} \}$

edges = $\left\{ (g_t,\ldots,g_1,x,d_1,\ldots,d_t) \quad (g_{t-1},\ldots,g_1,x,d_1,\ldots,d_t,d_{t+1}) \right\}$

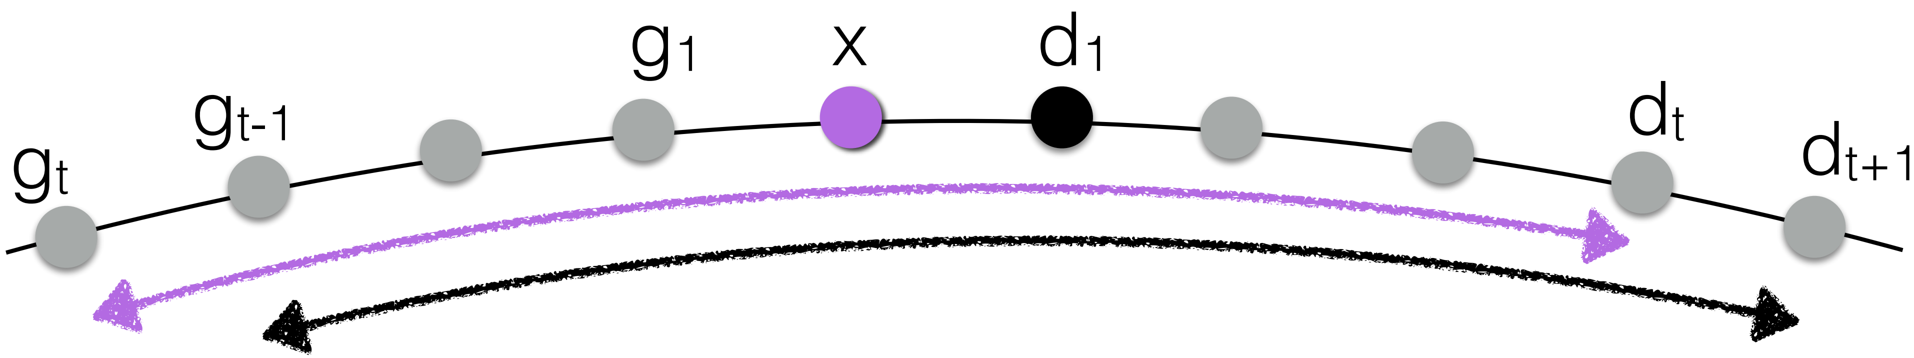

1. t-round 3-coloring algorithm for $C_n \Rightarrow \chi(G_{t,n}) \leq 3$
2. $t < ½ \log^* n - O(1) \Rightarrow \chi(G_{t,n}) > 3$

Step 1

Lemma t-round c-coloring algo for $C_n \Rightarrow \chi(G_{t,n}) \leq c$

Proof Algo $\mathcal{A}$ ➡ vertex $(g_t,\ldots,g_1,x,d_1,\ldots,d_t)$ colored

$$\mathcal{A}(g_t,\ldots,g_1,x,d_1,\ldots,d_t)$$

Coloring is proper as

$$(g_t,\ldots,g_1,x,d_1,\ldots,d_t) \text{ and } (g_{t-1},\ldots,g_1,x,d_1,\ldots,d_{t+1})$$

can appear as view of x and d_1 in some instances of ID assignment to the nodes of the ring. ❏

Application

Corollary (Linial, 1992) For n even, 2-coloring C_n requires $\Omega(n)$ rounds.

Proof Assume t rounds, with $t \leq n/2 - 2 \Rightarrow 2t+1 \leq n-3$.

1. $(x_1,x_2,\ldots,x_{2t+1})$
2. $(x_2,\ldots,x_{2t+1},y)$
3. $(x_3,\ldots,x_{2t+1},y,z)$
4. $(x_4,\ldots,x_{2t+1},y,z,x_1)$
5. $(x_5,\ldots,x_{2t+1},y,z,x_1,x_2)$

$\vdots$

2t+1. $(x_{2t+1},y,z,x_1,\ldots,x_{2t-2})$

2t+2. $(y,z,x_1,\ldots,x_{2t-2},x_{2t-1})$

2t+3. $(z,x_1,\ldots,x_{2t-1},x_{2t})$

odd cycle

⬇

$\chi(G_{t,n}) > 2$

❑

Step 2

Lemma $t < \frac{1}{2} \log^* n - O(1) \Rightarrow \chi(G_{t,n}) > 3$

Proof is technical (uses line graphs)[1]

But worth reading!

[1]Other proofs use Ramsey theory.

A simpler proof of Linial's lower bound

Proof (Laurinharju & Suomela, 2014)

$\mathcal{A}$ is a k-ary c-coloring function if

1. $\mathcal{A}(x_1,x_2,\ldots,x_k) \in \{1,2,\ldots,c\}$ for all $1 \leq x_1 < x_2 < \ldots < x_k \leq n$
2. $\mathcal{A}(x_1,x_2,\ldots,x_k) \neq A(x_2,x_3,\ldots,x_{k+1})$ for all $x_k < x_{k+1} \leq n$

Claim 0. t-tound algorithm $\mathcal{A}$ for 3-coloring C_n

➥ $\mathcal{A}$ is (2t+1)-ary 3-coloring function

Claim 1. If $\mathcal{A}$ is a 1-ary c-coloring function then $c \geq n$.

Claim 2. If $\mathcal{A}$ is a k-ary c-coloring function, then there is a (k–1)-ary 2^c-colouring function $\mathcal{B}$.

$$\mathcal{B}(x_1,x_2,\ldots,x_{k-1}) = \{\ \mathcal{A}(x_1,x_2,\ldots,x_{k-1},x_k) : x_k > x_{k-1}\ \}$$

For contradiction, let $1 \leq x_1 < x_2 < \ldots < x_k \leq n$ with

$$\mathcal{B}(x_1,x_2,\ldots,x_{k-1}) = \mathcal{B}(x_2,\ldots,x_{k-1},x_k)$$

Let $c = \mathcal{A}(x_1,x_2,\ldots,x_{k-1},x_k)$.

➥ $c \in \mathcal{B}(x_1,x_2,\ldots,x_{k-1})$ ➡ $c \in \mathcal{B}(x_2,\ldots,x_{k-1},x_k)$

➥ $\exists\ x_{k+1} > x_k : c = \mathcal{A}(x_2,\ldots,x_k,x_{k+1})$ ➡ $\mathcal{A}$ is not k-ary c-coloring function. ∎

Theorem Any 3-coloring algorithm for C_n performs in $\Omega(\log^* n)$ rounds.

Proof Let $\mathcal{A}$ be a t-tound algorithm for 3-coloring C_n

$\Rightarrow \mathcal{A}$ is $(2t+1)$-ary 3-coloring function (by Claim 0)

$\Rightarrow \mathcal{A}$ is $(2t)$-ary 2^3 -coloring function (by Claim 2)

$\Rightarrow \mathcal{A}$ is $(2t-1)$-ary $2^{(2)3}$-coloring function

$\Rightarrow \mathcal{A}$ is $(2t-2)$-ary $2^{(3)3}$-coloring function

$\vdots$

$\Rightarrow \mathcal{A}$ is (1)-ary $2^{(2t)3}$-coloring function

$\Rightarrow 2^{(2t)3} \geq n$ (by Claim 1)

$\Rightarrow t \geq \frac{1}{2} \log^* n - 1.$ ❑

$(\Delta+1)$-coloring arbitrary graphs

- Best lower bound (Linial, 1992)

$\Omega(\log^* n)$ rounds

- Best upper bound (Panconesi & Srinivasan, 1992)

$2^{O(\sqrt{\log n})}$ rounds

Gap open for a quarter of a century!

(Δ+1)-coloring arbitrary graphs

Complexity as $f(n)+g(\Delta)$

Theorem (Linial, 1992)
There is a $(\Delta+1)$-coloring algorithm performing in $O(\log^* n) + \tilde{O}(\Delta^2)$ rounds.

Theorem (F., Heinrich, Kosowski, 2016)
There is a $(\Delta+1)$-coloring algorithm performing in $O(\log^* n) + \tilde{O}(\sqrt{\Delta})$ rounds.

$O(\Delta^2)$-coloring

Theorem (Linial, 1992) $O(\Delta^2)$-coloring in $\log^* n+O(1)$ rounds

Lemma For all $k > \Delta \geq 2$, there exists $J = \{S_1,\ldots,S_k\}$ where

$$S_i \subseteq \{1,\ldots, 5 \lceil \Delta^2 \log k \rceil \} \text{ for } i=1,\ldots,k$$

such that, for every $\Delta+1$ sets $S_{i0}, S_{i1},\ldots, S_{i\Delta}$ in J, we have

$$S_{i0} \not\subseteq \cup_{j=1,\ldots,\Delta} S_{ij}.$$

Algorithm: Init: $k = n$ and color(u)=ID(u)

Each round: color range $[1,k]$ reduced to $[1,5 \lceil \Delta^2 \log k \rceil]$

color(u) = c ➡ u has set S_c

New color: smallest $x \in S_c \setminus \cup_{i=1,\ldots,\Delta} S_{color(vi)}$. ❑

Locally Iterative Algorithm

Theorem [L. Barenboim, M. Elkin, U. Goldenberg (2017)
There exists a locally iterative algorithm for $(\Delta+1)$-coloring, performing in $O(\log^* n + \Delta)$ rounds.

Proof. Compute $O(\Delta^2)$-coloring in $\log^* n+O(1)$ rounds.

Assume for simplicity a $(\Delta+1)^2$-coloring with $\Delta + 1 = p$ prime.

Represent color $c_0(v) = (a_v,b_v)$ where $a_v,b_v \in GF(p)$.

- if $\nexists\, u \in N(v)$, with $b_u = b_v$ then v adopts $(0,b_v)$ as final color;
- otherwise, v recolors itself as $(a_v, b_v + a_v)$.

The following two properties hold:

- Recoloring preserves proper coloring
- After $2p + 1 = 2(\Delta + 1) + 1$ rounds, all nodes have finalized their color.

❑

Locally Checkable Labeling

Let $\mathcal{F}_\Delta$ be the set of all (connected) graphs with maximum degree Δ.

Definition (Naor and Stockmeyer, 1995) An LCL in $\mathcal{F}_\Delta$ is specified by a finite set of labels, and a finite set of labeled balls with maximum degree Δ, called good balls.

Examples:

- k-coloring, k-edge-coloring
- maximal independent set (MIS)
- maximal matching
- Etc.

Focus is on LCL tasks solvable sequentially by a greedy algorithm selecting nodes in arbitrary order, like, e.g., k-coloring for $k \geq \Delta+1$.

Maximal Independent Set

- (Δ+1)-coloring ➜ MIS in Δ rounds by maximizing {1}
- MIS ➜ (Δ+1)-coloring by simulation

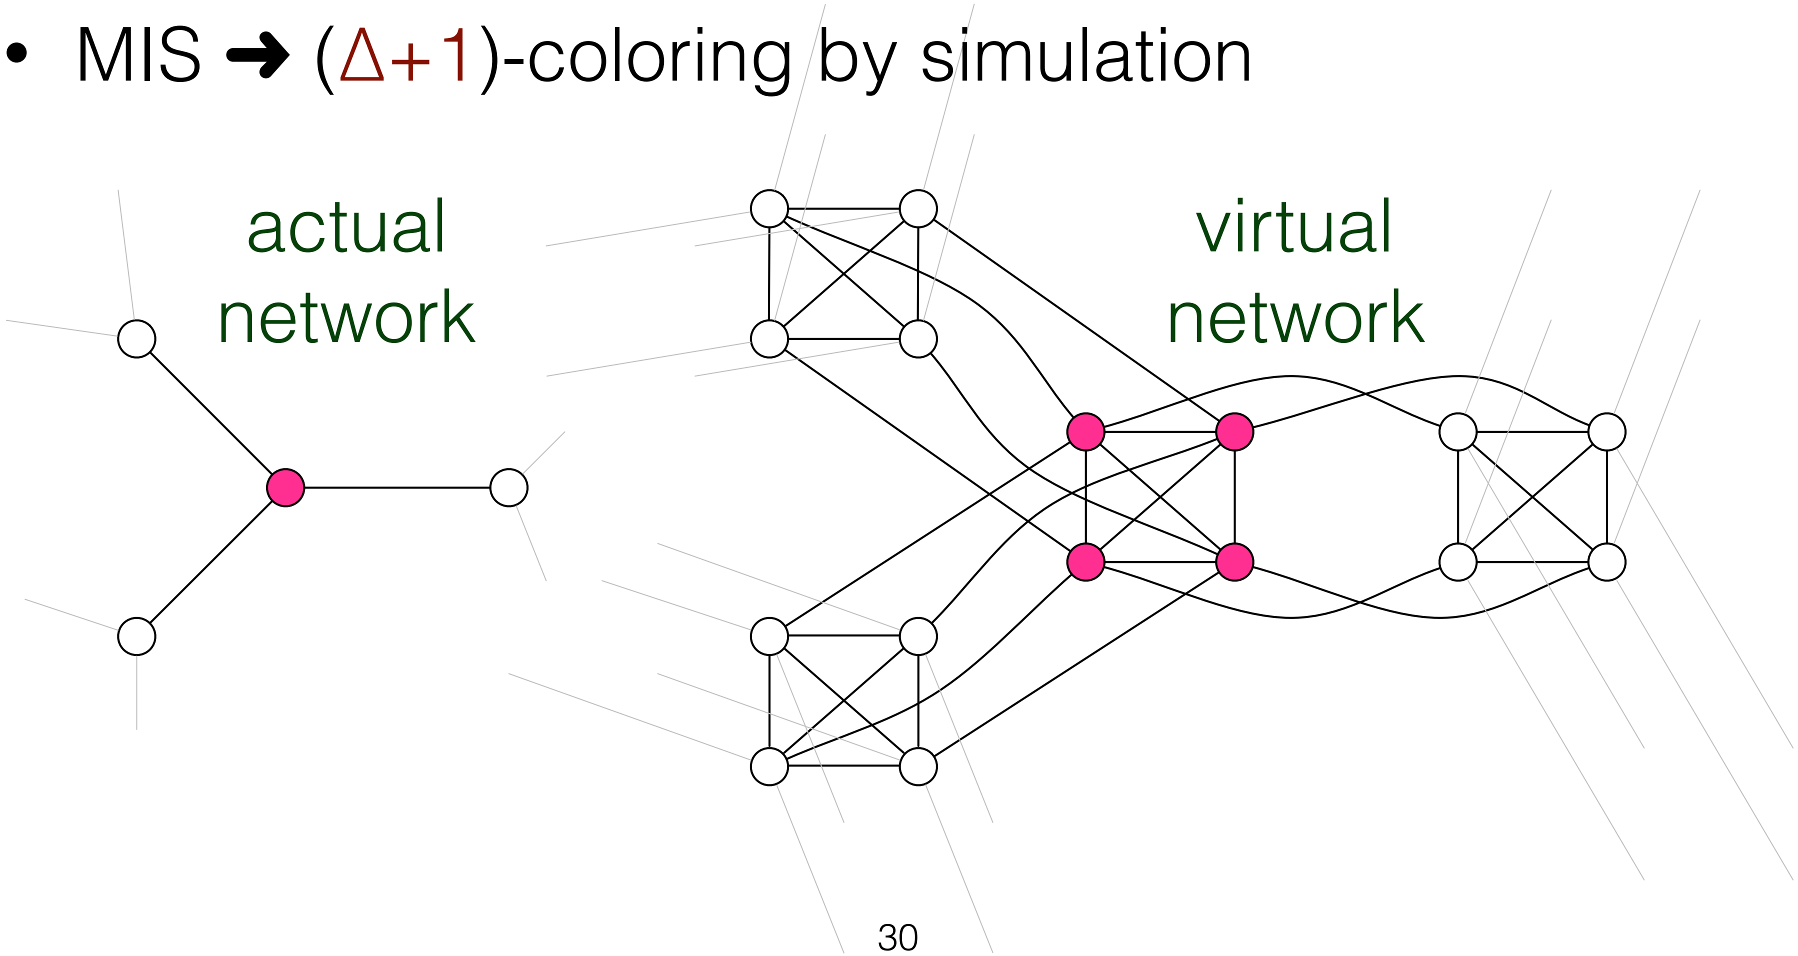

Claim 1. At most one node of each clique in the MIS

Claim 2. At least one node of each clique in the MIS

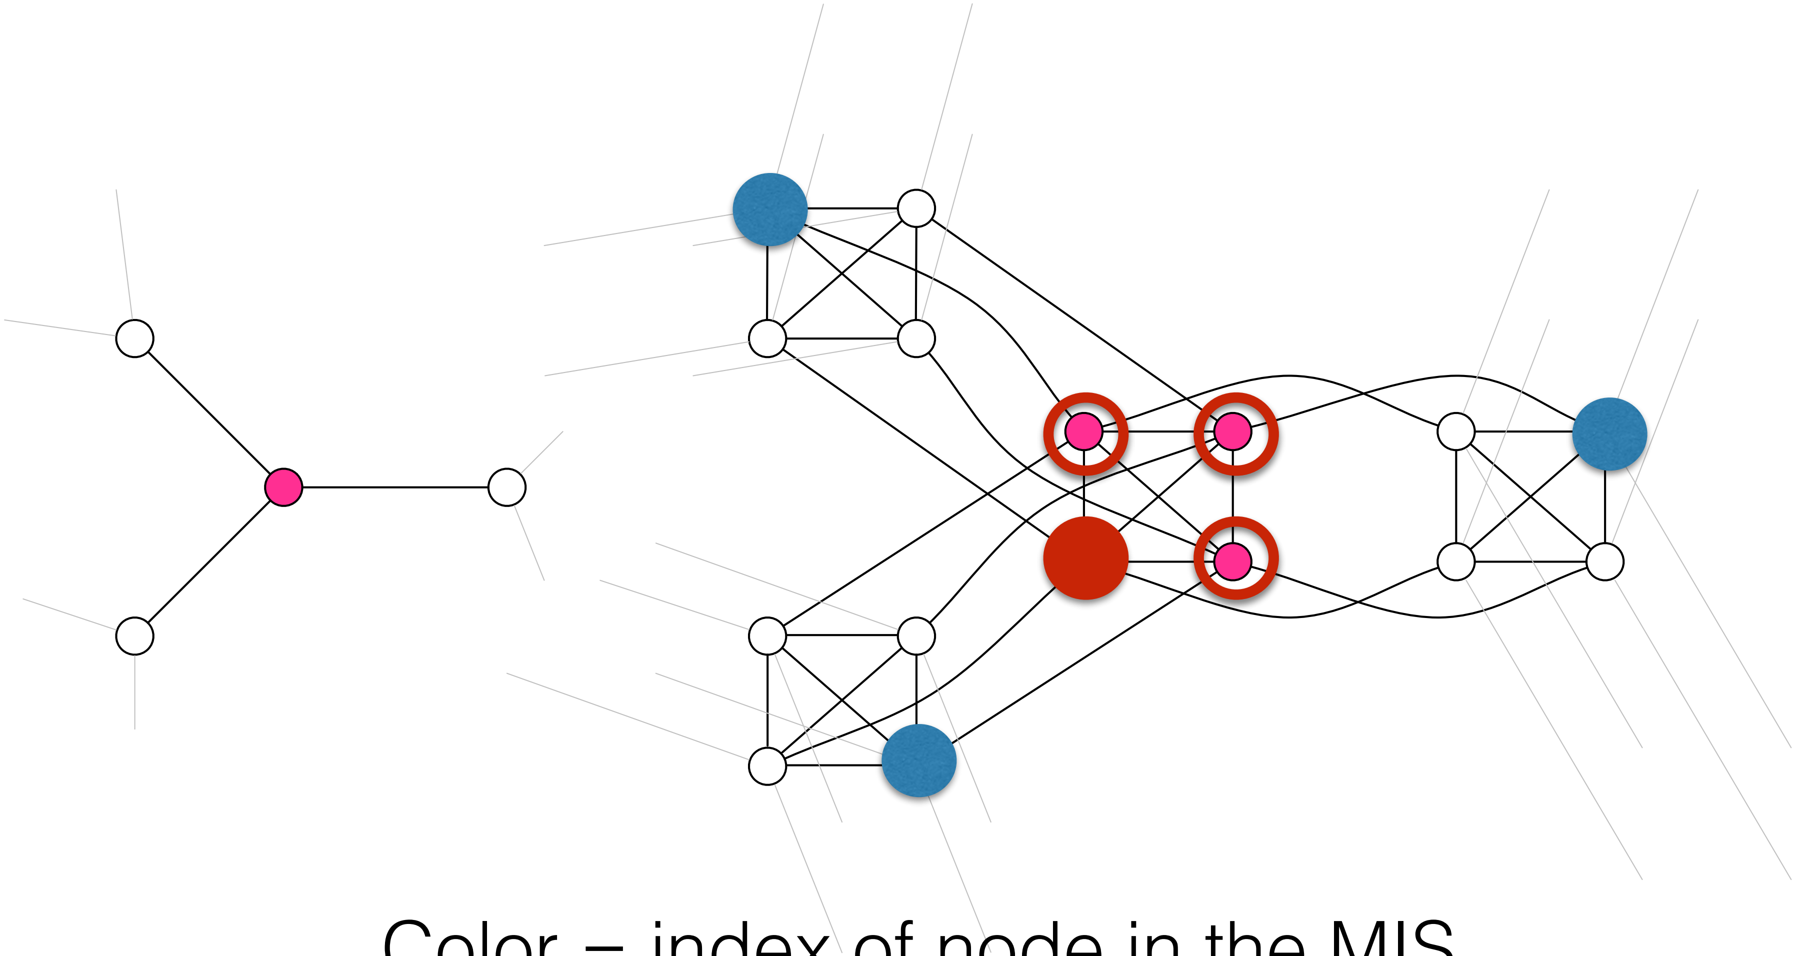

Color = index of node in the MIS

Line Graphs

Definition The line graph of a graph G is the graph L(G) such that

- V(L(G)) = E(G)
- {e,e'} ∈ E(L(G)) ⟺ e and e' are incident in G

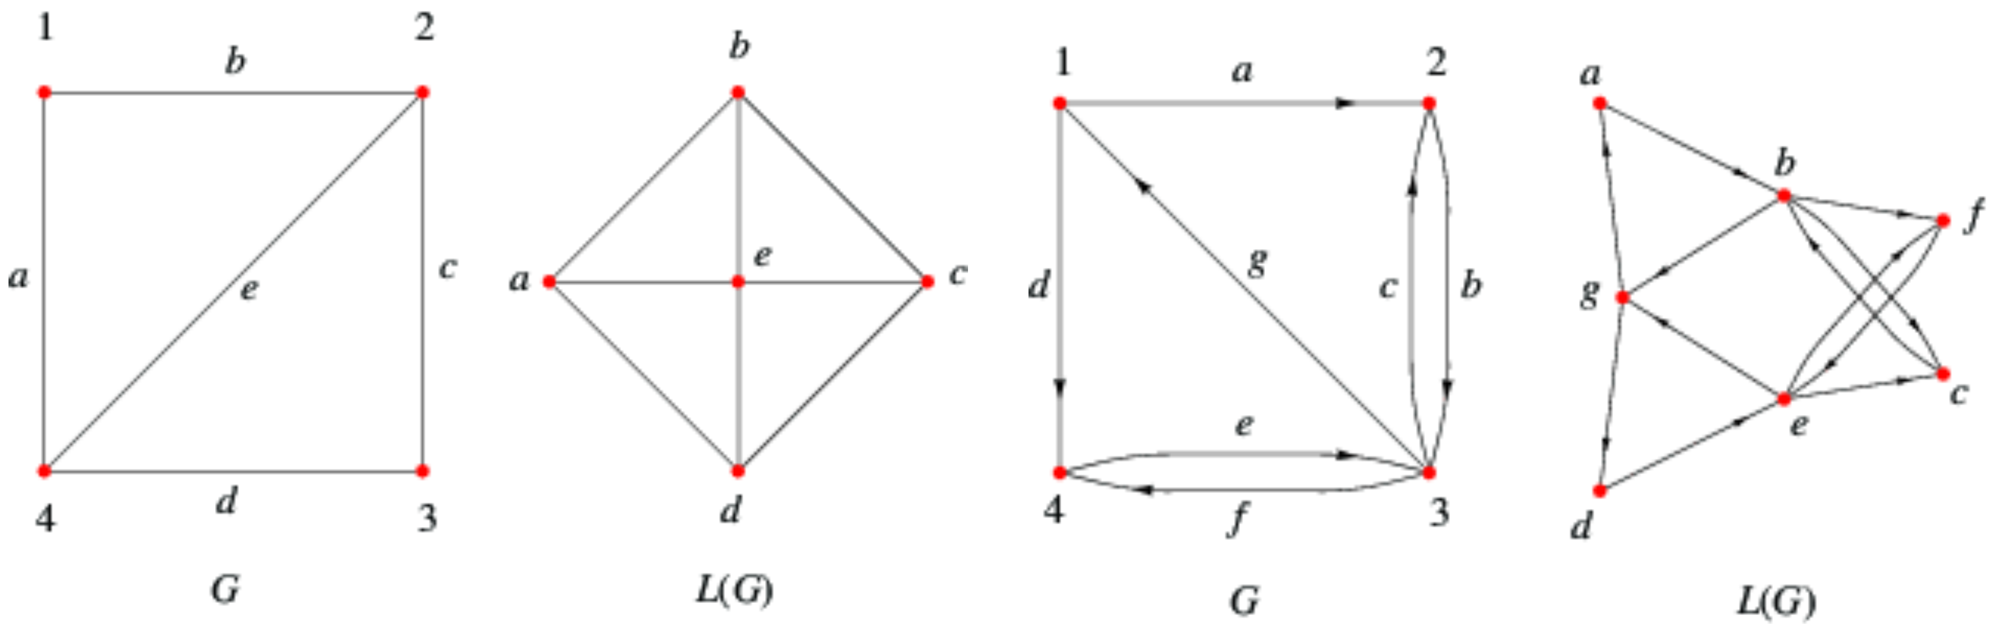

Four classical problems

MIS

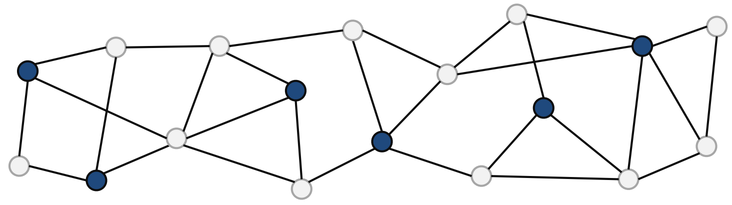

-Vertex Coloring

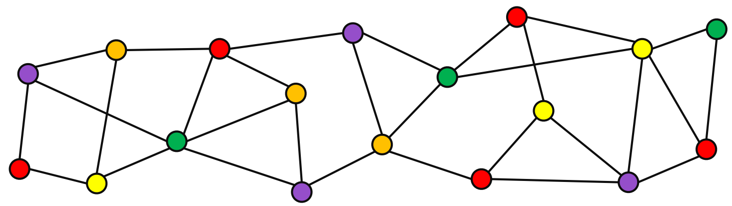

MIS on line graph

-coloring
on line graph

Maximal Matching

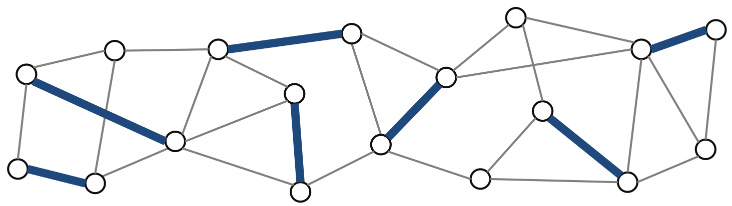

-Edge Coloring

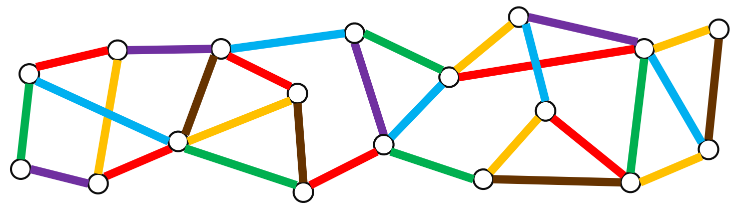

Round Complexity

	MIS	(Δ+1)-coloring
Deterministic	$2^{\sqrt{\log(n)}}$ Panconesi, Srinivasan (1992)	$2^{\sqrt{\log(n)}}$ Panconesi, Srinivasan (1992)
Randomized	$2^{\sqrt{\log\log(n)}}+O(\log \Delta)$ Ghaffari (2016)	$2^{\sqrt{\log\log(n)}}$ Chang, Li, Pettie (2018)
	Maximal Matching	**(2Δ-1)-edge-coloring**
Deterministic	$O(\log^3 n)$ Fischer (2017)	$O(\log^6 n)$ Ghaffari, Fisher, Kuhn (2017) Ghaffari, Harris, Kuhn (2018)
Randomized	$O(\log^3 \log n)+O(\log \Delta)$ Barenboim, Elkin, Pettie, Schneider (2012)	$O(\log^6 \log n)$ Elkin, Pettie, Su (2015)

Lower Bounds

	MIS and Maximal Matching	(Δ+1)-coloring and (2Δ-1)-edge-coloring
Deterministic and Randomized	$\Omega(\min\{ \log \Delta / \log\log \Delta, \sqrt{\log n / \log\log n} \})$ Kuhn, Moscibroda, Wattenhofer (2004)	$\Omega(\log^* n)$ Linial (1987) Naor (1990)

Randomized Algorithms

Randomized algorithm for (Δ+1)-coloring

Algorithm (Barenboim and Elkin, 2013) for node u

while uncolored **do**

$\mathscr{C}$ = {colors previously adopted by neighbors}

pick $\ell(u)$ at random in {0,1,…,Δ+1} - $\mathscr{C}$

- 0 is picked w/ probability ½
- $\ell(u) \in$ {1,…,Δ+1} - $\mathscr{C}$ is picket w/ proba $1/(2(\Delta+1-|\mathscr{C}|))$

if $\ell(u) \neq 0$ **and** $\ell(u) \notin$ {colors picked by neighbors}

then adopt $\ell(u)$ as my color

else remain uncolored

inform neighbors of status

1 round

1 round

Definition A sequence $(\mathcal{E}_n)_{n\geq 1}$ of events holds with high probability (whp) whenever $\Pr[\mathcal{E}_n] = 1 - O(1/n^c)$ for some constant $c > 0$.

Theorem (Barenboim and Elkin, 2013) The $(\Delta+1)$-coloring algorithm takes, w.h.p., $O(\log n)$ rounds.

A and B independent
$\Leftrightarrow \Pr[A \wedge B] = \Pr[A] \cdot \Pr[B]$

Recall:

« A given B holds » or
« A conditioned to B »

- $\Pr[A|B] = \Pr[A \wedge B] / \Pr[B] \Rightarrow \Pr[A \wedge B] = \Pr[A|B] \cdot \Pr[B]$ ($\wedge$: and)

- $\Pr[A] = \Pr[A|B] \cdot \Pr[B] + \Pr[A|\neg B] \cdot \Pr[\neg B]$

- Union bound: $\Pr[A \vee B] \leq \Pr[A] + \Pr[B]$ ($\vee$: or)

$$\Pr[\exists\ s \in S : s \models \mathcal{P}] = \Pr[(s_1 \models \mathcal{P}) \vee (s_2 \models \mathcal{P}) \vee \ldots \vee (s_m \models \mathcal{P})]$$

Claim For every node u, at any round, Pr[u terminates] ≥ ¼

$$\begin{aligned}\Pr[u \text{ termine}] &= \Pr[\ell(u) \neq 0 \text{ et aucun } v \in N(u) \text{ satisfait } \ell(v) = \ell(u)] \\ &= \Pr[\forall v \in N(u), \ell(v) \neq \ell(u) \mid \ell(u) \neq 0] \cdot \Pr[\ell(u) \neq 0] \\ &= \frac{1}{2} \cdot \Pr[\forall v \in N(u), \ell(v) \neq \ell(u) \mid \ell(u) \neq 0]\end{aligned}$$

$$\begin{aligned}\Pr[\ell(v) = \ell(u) \mid \ell(u) \neq 0] &= \Pr[\ell(v) = \ell(u) \mid \ell(u) \neq 0 \wedge \ell(v) = 0]\ \Pr[\ell(v) = 0] \\ &\quad + \Pr[\ell(v) = \ell(u) \mid \ell(u) \neq 0 \wedge \ell(v) \neq 0]\ \Pr[\ell(v) \neq 0] \\ &= \Pr[\ell(v) = \ell(u) \mid \ell(u) \neq 0 \wedge \ell(v) \neq 0]\ \Pr[\ell(v) \neq 0] \\ &\leq \frac{1}{2}\ \Pr[\ell(v) = \ell(u) \mid \ell(u) \neq 0 \wedge \ell(v) \neq 0] \\ &= \frac{1}{2}\ \frac{1}{\Delta + 1 - |C(u)|}.\end{aligned}$$

$$\Pr[\exists v \in N(u) : \ell(v) = \ell(u) \mid \ell(u) \neq 0] \leq (\Delta - |C(u)|)\frac{1}{2(\Delta + 1 - |C(u)|)} < \frac{1}{2}$$

■

O(log n) rounds w.h.p.

Pr[u does not terminate in k ln(n) rounds]

$\leq (3/4)^{k \ln(n)} = n^{-k \ln(4/3)}$

Pr[∃u that does not terminate in k ln(n) rounds] $\leq n^{1-k \ln(4/3)}$

Let $c > 1$, by choosing $k = (1+c)/\ln(4/3)$, we get:

Pr[all nodes terminates after $(1+c)/\ln(4/3) \ln(n)$ rounds]

$\geq 1-1/n^c$.

❏

Randomized algorithm for MIS

Algorithm (Luby, 1986)

mis(u) $\in$ {-1,0,1} = {undecided, not in MIS, in MIS}

At any given round: $H = G\big[\{u : mis(u)=-1\}\big]$

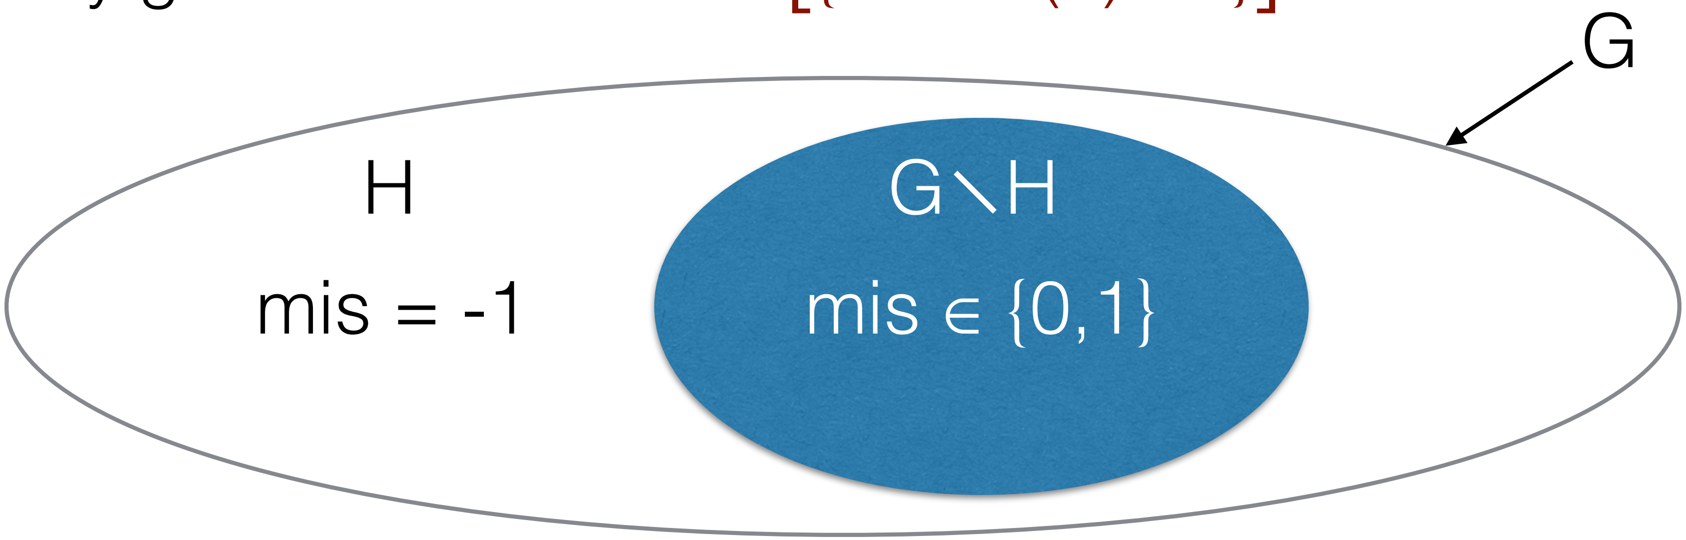

Trick: enforcing an order between nodes:

$v \succ u \iff deg_H(v) > deg_H(u)$

$\quad$ or $(deg_H(v) = deg_H(u)$ and $ID(v) > ID(u))$

Luby's algorithm

One phase of the algorithm for node u with mis(u) = -1

if $deg_H(u) = 0$ **then** mis(u) ← 1
else join(u) ← true with proba $1/(2\ deg_H(u))$, false otherwise
exchange join with every v ∈ N(u)
free(u) ← ∄ v ∈ N(u) such that v ≻ u and join(v)=true
if (join(u) = true **and** free(u) = true) **then** mis(u) ← 1
exchange mis with every v ∈ N(u)
if (mis(u) = -1 **and** ∃v∈N(u) mis(v)=1) **then** mis(u) ← 0
exchange mis with every v ∈ N(u)

Luby's algorithm terminates in O(log n) rounds, w.h.p.

Structure of the proof:

1. $\Pr[mis(u) = 1] \geq 1/(4\ deg_H(u))$
2. For a set $\mathcal{N}$ of nodes,

$$u \in \mathcal{N} \Rightarrow \Pr[u \text{ terminates}] \geq 1/36$$

3. For a large set $\mathcal{E}$ of edges,

$$e \in \mathcal{E} \Rightarrow \Pr[e \text{ removed from H}] \geq 1/36$$

4. Use concentration result (Chernoff bound) to get w.h.p.

Step 1

$$\begin{aligned}
\Pr[mis(u) \neq 1 \mid join(u)] &= \Pr[\exists v \in N(u) : v \succ u \wedge join(v) \mid join(u)] \\
&= \Pr[\exists v \in N(u) : v \succ u \wedge join(v)] \\
&\leq \sum_{v \in N(u) : v \succ u} \Pr[join(v)] \\
&= \sum_{v \in N(u) : v \succ u} \frac{1}{2\deg(v)} \\
&\leq \sum_{v \in N(u) : v \succ u} \frac{1}{2\deg(u)} \\
&\leq \frac{\deg(u)}{2\deg(u)} \\
&\leq \frac{1}{2}
\end{aligned}$$

if $\deg_H(u) = 0$ **then** mis(u) ← 1
else join(u) ← true with proba $1/(2 \deg_H(u))$
exchange join with every v ∈ N(u)
free(u) ← ∄ v ∈ N(u) such that v ≻ u and join(v)=true
if (join(u) = true **and** free(u) = true) **then** mis(u) ← 1
exchange mis with every v ∈ N(u)
if (mis(u) = -1 **and** ∃v∈N(u) mis(v)=1) **then** mis(u) ← 0
exchange mis with every v ∈ N(u)

$$\Pr[mis(u) = 1] = \Pr[mis(u) = 1 \mid join(u)] \cdot \Pr[join(u)]$$

$$\Pr[mis(u) = 1] \geq \frac{1}{2} \cdot \frac{1}{2\deg(u)} = \frac{1}{4\deg(u)}.$$

Step 2

A node u is large if $\sum_{v \in N(u)} \frac{1}{2\deg(v)} \geq \frac{1}{6}$

Claim: u large ⇒ Pr[u terminates] ≥1/36

- True if ∃v∈N(u) : $\deg_H(v) \leq 2$
- ∀v∈N(u), if $\deg_H(v) \geq 3$ then $\frac{1}{2\deg(v)} \leq \frac{1}{6}$

$\Longrightarrow$ ∃ S ⊆ N(u) : $\frac{1}{6} \leq \sum_{v \in S} \frac{1}{2\deg(v)} \leq \frac{1}{3}$

$$\Pr[E_1 \vee E_2 \vee \cdots \vee E_r] = \sum_i \Pr[E_i] - \sum_{i \neq j} \Pr[E_i \wedge E_j] + \sum_{i \neq j \neq k} \Pr[E_i \wedge E_j \wedge E_k] - \ldots$$
$$\cdots + (-1)^{r+1} \Pr[E_1 \wedge \cdots \wedge E_r].$$

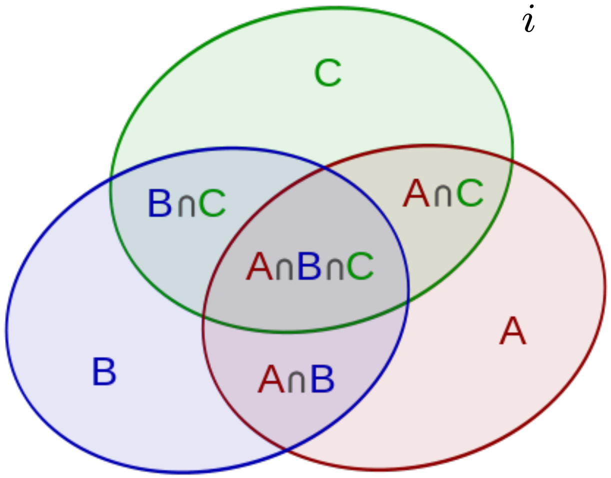

$$\Pr[mis(u) \neq -1] \geq \Pr[\exists v \in S : mis(v) = 1]$$
$$\geq \sum_{v \in S} \Pr[mis(v) = 1] - \sum_{v,w \in S, v \neq w} \Pr[mis(v) = 1 \wedge mis(w) = 1].$$

$$\Longrightarrow \Pr[mis(u) \neq -1] \geq \sum_{v \in S} \Pr[mis(v) = 1] - \sum_{v,w \in S, v \neq w} \Pr[join(v) \wedge join(w)]$$
$$\geq \sum_{v \in S} \Pr[mis(v) = 1] - \sum_{v \in S} \sum_{w \in S} \Pr[join(v)] \cdot \Pr[join(w)]$$
$$\geq \sum_{v \in S} \frac{1}{4\deg(v)} - \sum_{v \in S} \sum_{w \in S} \frac{1}{2\deg(v)} \cdot \frac{1}{2\deg(w)}$$
$$\geq \left(\sum_{v \in S} \frac{1}{2\deg(v)} \right) \left(\frac{1}{2} - \sum_{w \in S} \frac{1}{2\deg(w)} \right)$$
$$\geq \frac{1}{6} \left(\frac{1}{2} - \frac{1}{3} \right) = \frac{1}{36}.$$

if $\deg_H(u) = 0$ **then** mis(u) ← 1
else join(u) ← true with proba $1/(2 \deg_H(u))$
exchange join with every v ∈ N(u)
free(u) ← ∄ v ∈ N(u) such that v ≻ u and join(v)=true
if (join(u) = true **and** free(u) = true) **then** mis(u) ← 1
exchange mis with every v ∈ N(u)
if (mis(u) = -1 **and** ∃v∈N(u) mis(v)=1) **then** mis(u) ← 0
exchange mis with every v ∈ N(u)

Step 3

An edge $e=\{u,v\}$ is large if u or v is large

For $e = \{u,v\}$ with $u \prec v$, orient the edge $u \rightarrow v$

Claim For every small node u, $\deg^+(u) \geq 2 \deg^-(u)$

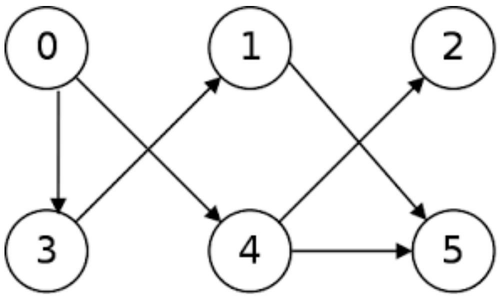

Out-degree

In-degree

Indeed: $\deg^+(u) < 2 \deg^-(u) \implies \deg(u) < 3 \deg^-(u)$

$S = \{v \in N(u) : \deg(v) \leq \deg(u)\}$

$|S| \geq \deg^-(u) \implies |S| \geq |N(u)|/3$

$$\sum_{v \in N(u)} \frac{1}{2\deg(v)} \geq \sum_{v \in S} \frac{1}{2\deg(v)} \geq \sum_{v \in S} \frac{1}{2\deg(u)} \geq \frac{\deg(u)}{3} \cdot \frac{1}{2\deg(u)} = \frac{1}{6} \quad \blacksquare$$

Let m = |E(H)|
We have: $$\sum_{u \text{ petit}} \deg^-(u) \leq \frac{1}{2} \sum_{u \text{ petit}} \deg^+(u) \leq \frac{m}{2}$$

$$\implies \sum_{u \text{ grand}} \deg^-(u) \geq \frac{m}{2} \implies$$ at least m/2 large edges

X_e = Bernouilli variable equal to 1 if e is removed from H

For e large, $\Pr[X_e=1] \geq 1/36 \implies \mathbb{E}X_e \geq 1/36$

$X = \sum_{e \text{ large}} X_e \implies \mathbb{E}X = \sum_{e \text{ large}} \mathbb{E}X_e \geq m/72$

Let $p = \Pr[X \leq ½\, \mathbb{E}X]$

$$\mathbb{E}X = \sum_{x=0}^{m} x \;\Pr[X=x] = \sum_{x=0}^{\frac{1}{2}\mathbb{E}X} x \;\Pr[X=x] + \sum_{x=\frac{1}{2}\,\mathbb{E}X+1}^{m} x \;\Pr[X=x] \leq \tfrac{1}{2}\, p\, \mathbb{E}X + (1-p)m$$

$$\implies p \leq \frac{m - \mathbb{E}X}{m - \frac{1}{2}\mathbb{E}X} \leq \frac{m - \frac{1}{2}\mathbb{E}X}{m} \leq 1 - \frac{1}{144}.$$

Let $\mathscr{E}$ = « at least m/144 edges are removed from H »

$\Pr[\mathscr{E}] \geq 1/144$

Step 4

Let $Y_1, Y_2, \ldots, Y_k$ be Bernouilli variables w/ parameter $q = 1/144$

Let $Y = Y_1 + Y_2 + \ldots + Y_k$

Remark: Let $\alpha = 144/143$. If $Y \geq \log_\alpha |E(G)|$ then termination.

Chernoff Inequality: $\forall\ \delta \in\ ]0,1[$, $\Pr[Y \leq (1-\delta)\mathbb{E}Y] \leq e^{-\frac{1}{2}\delta^2\mathbb{E}Y}$.

We have $\mathbb{E}Y = kq$, so, with $\delta = ½$, we get $\Pr[Y \leq \frac{kq}{2}] \leq e^{-\frac{kq}{8}}$

For $k = c \log_\alpha n$, we get $\Pr[Y \leq \frac{cq \log_\alpha n}{2}] \leq e^{-\frac{cq \log_\alpha n}{8}}$

Let $c = 4/q \implies ½\, c\, q \log_\alpha n \geq \log_\alpha |E(G)|$ and $cq \geq 8 \ln(\alpha)$.

$$\implies e^{-\frac{cq \log_\alpha n}{8}} = \frac{1}{n^{\frac{cq}{8 \ln \alpha}}} \leq \frac{1}{n}. \quad \implies \Pr[Y \leq \log_\alpha m] \leq \frac{1}{n}.$$

Thus Luby's algorithm terminates in O(log n) rounds w.h.p.

❑

Deterministic ⬌ Randomized

Network Decomposition

Definition A (d,c)-decomposition of an n-node graph G = (V, E) is a partition of V into clusters such that each cluster has diameter at most d and the cluster graph is properly colored with colors 1, . . . , c.

Theorem [Linial and Saks (1993)]
Every graph has a (O(log n),O(log n))-decomposition, and such a decomposition can be computed by a randomized algorithm in $O(\log^2 n)$ rounds in the LOCAL model.

Theorem [Panconesi and Srinivasan (1992)]
A $(2^{O(\sqrt{\log n})}, 2^{O(\sqrt{\log n})})$-decomposition can be computed deterministically in $2^{O(\sqrt{\log n})}$ rounds in the LOCAL model.

Impact on coloring and MIS

Lemma Given a (d,c)-decomposition, $(\Delta+1)$-coloring and MIS can be solved in O(cd) rounds in the LOCAL model.

Proof

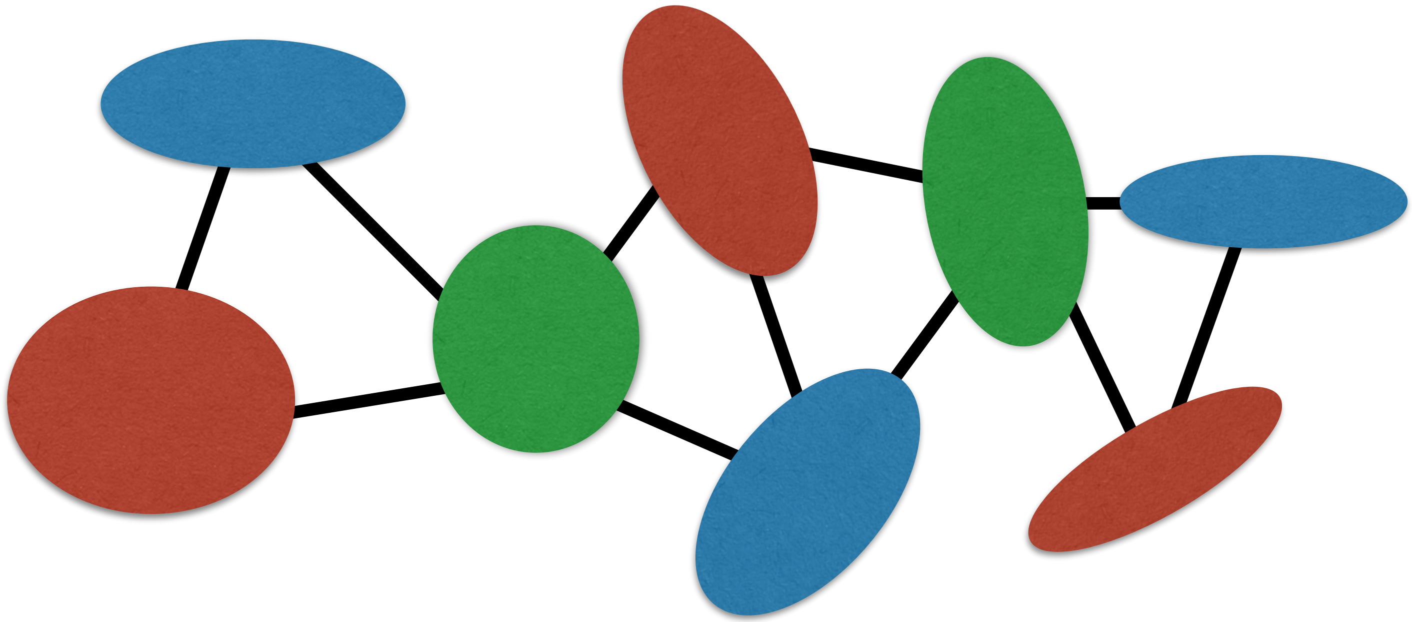

Proceed in c phases, each of O(d) rounds ❏

Theorem [V. Rozhon and M. Ghaffari (2019)]
A (O(log n),O(log n))-decomposition can be computed deterministically in $O(\log^{O(1)} n)$ rounds in the LOCAL model.

Corollary (Δ+1)-coloring and MIS can be deterministically solved in $O(\log^{O(1)} n)$ rounds in the LOCAL model.

SLOCAL Model

M. Ghaffari, F Kuhn, Y. Maus (2017)

- Sequential variant of the LOCAL model:
  - nodes are considered sequentially, one by one
  - the current node computes its output based solely on the states of the nodes in the ball of radius t around it
- LOCAL(t) = {problems solvable in t rounds}
- SLOCAL(t) = {problem solvable with balls of radius t}
- P-LOCAL = LOCAL($\log^{O(1)} n$)
- P-SLOCAL = SLOCAL($\log^{O(1)} n$)

Completeness Results

In the LOCAL model, a problem Q is t-reducible to another problem P if

t-round algorithm for P ⇒ t-round algorithm for Q.

P is P-SLOCAL-complete if P ∈ P-SLOCAL, and any Q ∈ P-SLOCAL is $O(\log^{O(1)} n)$ -reducible to P.

Theorem [M. Ghaffari, F Kuhn, Y. Maus (2017)] Computing a $(O(\log^{O(1)} n), O(\log^{O(1)} n))$-decomposition is P-SLOCAL-complete.

Corollary P-LOCAL = P-SLOCAL.

Derandomization

For Locally Checkable Labeling (LCL) problems:

Theorem [M. Naor and L. Stockmeyer (1992)]
LOCAL(O(1)) = RLOCAL(O(1))

Theorem [L. Feuilloley and P. F. (2015)]
LOCAL(O(1)) = RLOCAL(O(1)) also for randomly locally checkable problems.

Theorem [V. Rozhon and M. Ghaffari (2019)]
P-LOCAL = P-RLOCAL.

Randomized Algorithms using Shattering

Pick ● or ○ u.a.r.

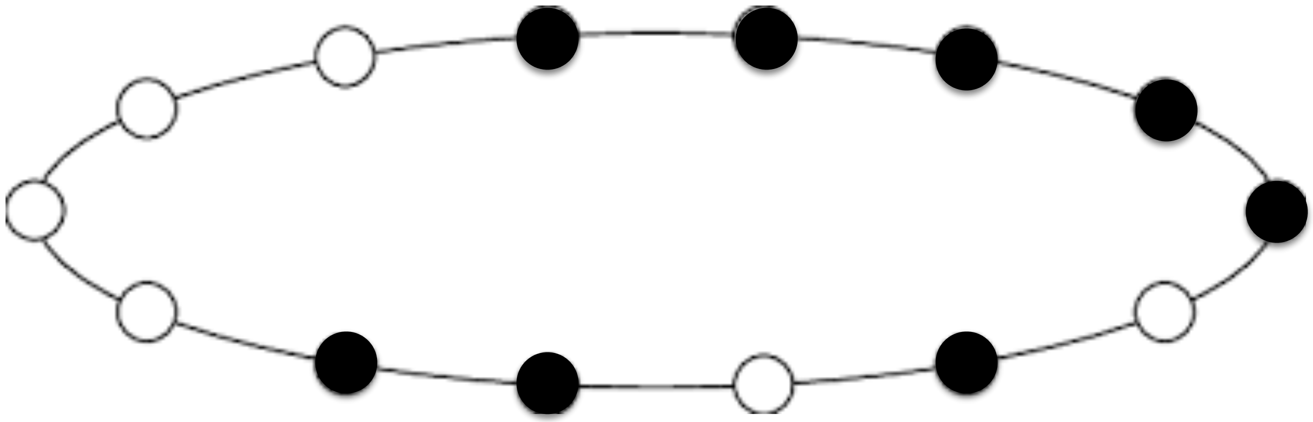

W.h.p., max length monochromatic interval ≤ O(log n)

3-coloring or MIS: #rounds ≈ **Det**(O(log n))

Graph Shattering

1. Shatter the graph using randomization
2. Complete each piece deterministically

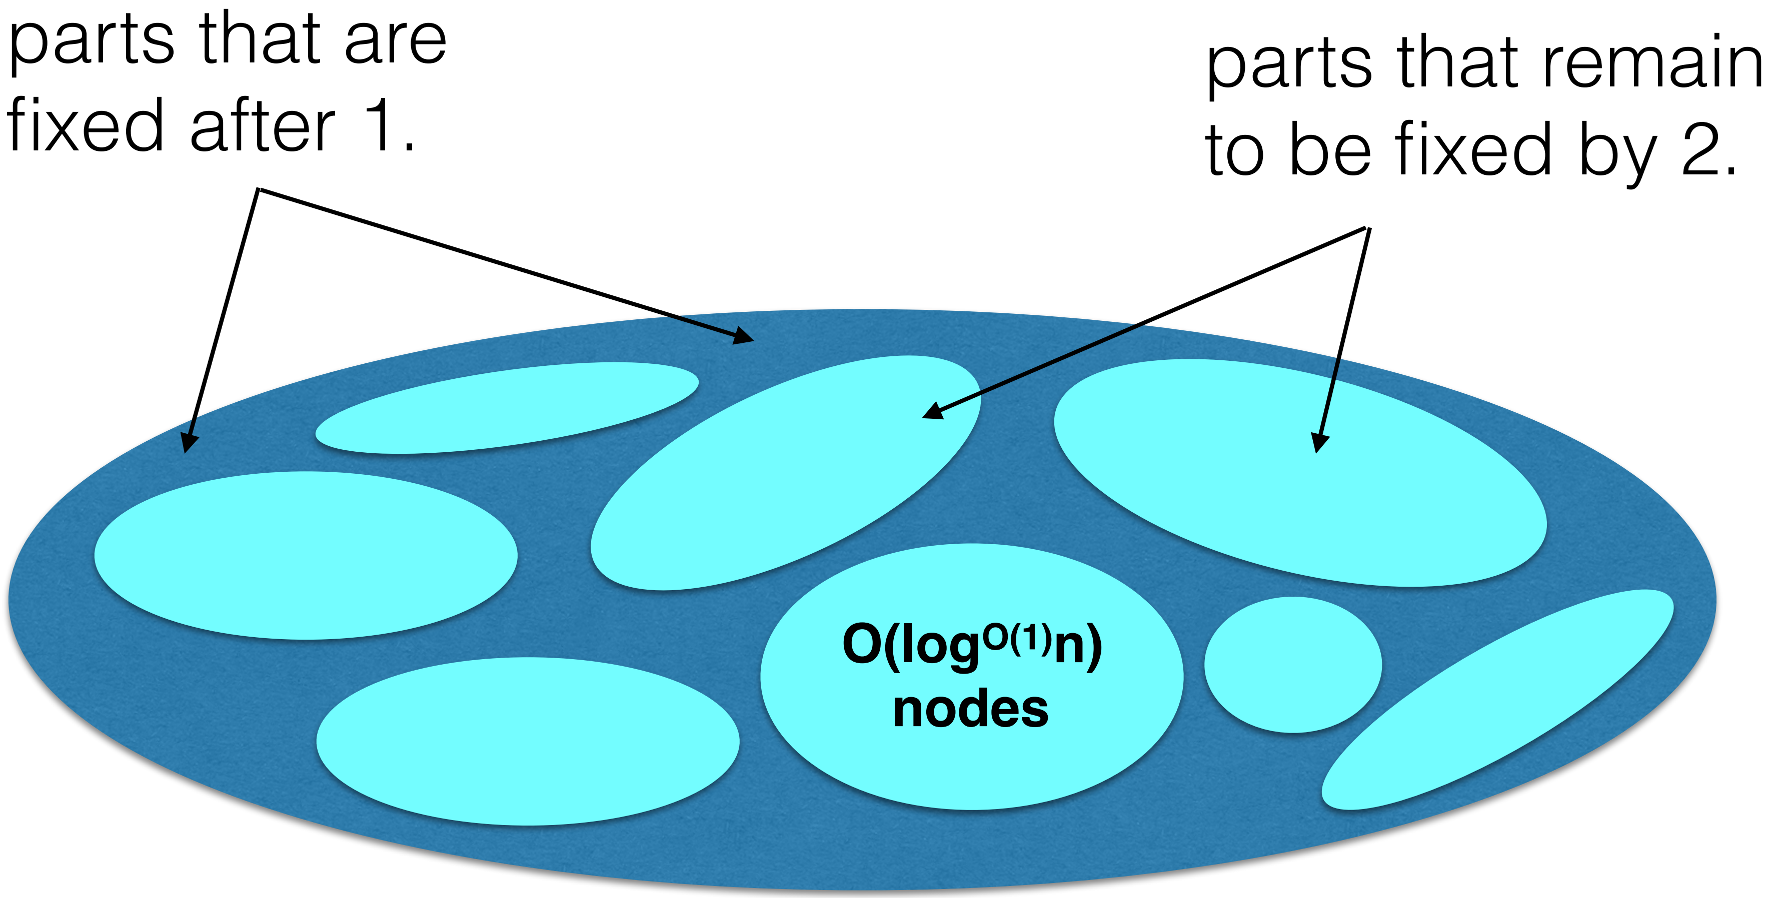

$$\textbf{Rand}(n) \approx \textbf{Det}(O(\log^{O(1)} n))$$

Deterministic lower bounds

⬇

Randomized lower bounds

Theorem [Y.-J. Chang, T. Kopelowitz, S. Pettie (2016)]

For any LCL problem in the LOCAL model, its randomized complexity on instances of size n is at least its deterministic complexity on instances of size $\sqrt{\log n}$.

Conclusion: one needs to design better deterministic algorithms for improving the performances of randomized algorithms!

Concluding remarks

Round Complexity

	MIS	(Δ+1)-coloring
Deterministic	$O(\log^{O(1)} n)$ Rozhon, Ghaffari (2019)	$O(\log^{O(1)} n)$ Rozhon, Ghaffari (2019)
Randomized	$O(\log^{O(1)} \log n) + O(\log \Delta)$	$O(\log^{O(1)} \log n)$
	Maximal Matching	**(2Δ-1)-edge-coloring**
Deterministic	$O(\log^3 n)$ Fisher (2017)	$O(\log^6 n)$ Ghaffari, Fisher, Kuhn (2017) Ghaffari, Harris, Kuhn (2018)
Randomized	$O(\log^3 \log n) + O(\log \Delta)$ Barenboim, Elkin, Pettie, Schneider (2012)	$O(\log^6 \log n)$ Elkin, Pettie, Su (2015)

Lower Bounds

	MIS and Maximal Matching	$(\Delta+1)$-coloring and $(2\Delta-1)$-edge-coloring
Deterministic and Randomized	$\Omega(\min\{ \log \Delta / \log\log \Delta, \sqrt{\log n / \log\log n} \})$ Kuhn, Moscibroda, Wattenhofer (2004)	$\Omega(\log^* n)$ Linial (1987) Naor (1990)

Open problems

- Improve the constants (i.e., the degrees of the polylog)
- Close the gaps between lower and upper bounds
- Is $(\Delta+1)$-coloring solvable in $O(\log^* n)$ rounds?