LOCAL Model & LCL Problems

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:

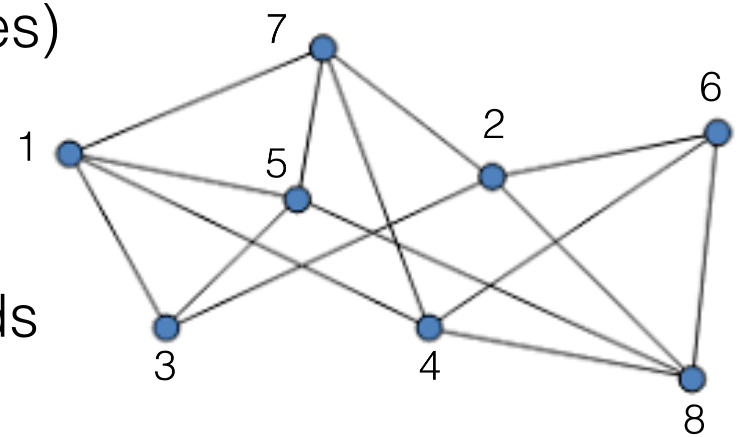

1. Sends a message to each neighbor
2. Receives a message from each neighbor
3. Performs individual computation (same algorithm for all nodes)

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: (1) Gather all data in the t-ball around it; (2) Simulate and compute the solution.

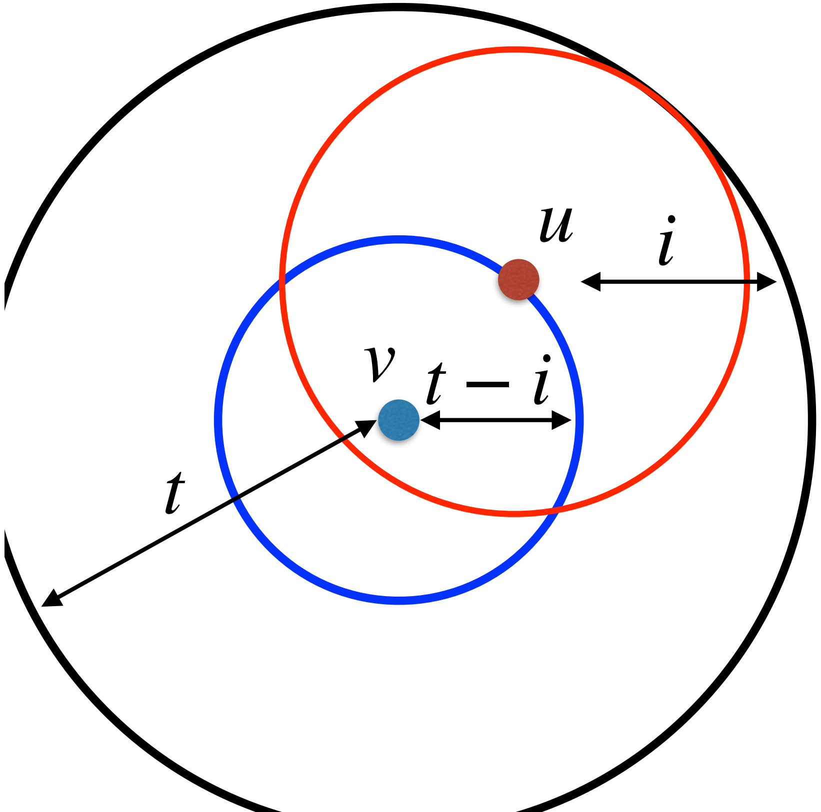

For every $i = 1,\ldots,t$
it suffices for node v to simulate the i-th round of all nodes in $B_G(v, t-i)$
$= \{u \in V(G) \mid \mathrm{dist}_G(u,v) \leq t-i\}$

Four classical problems

MIS

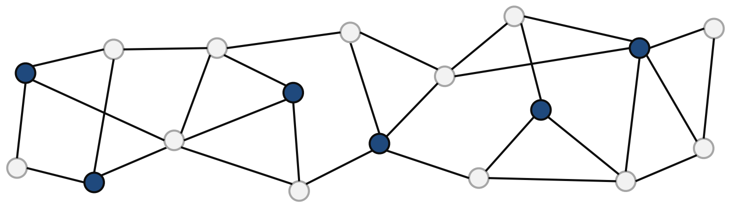

-Vertex Coloring

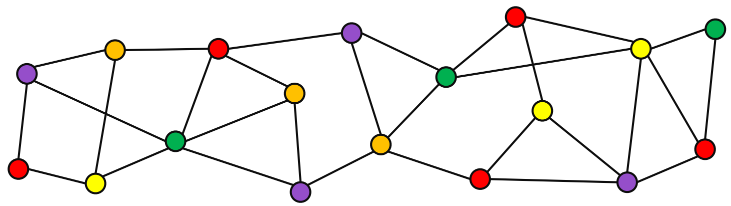

Maximal Matching

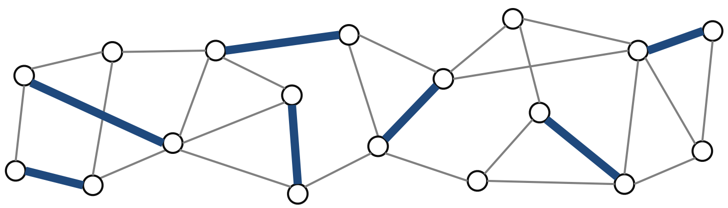

-Edge Coloring

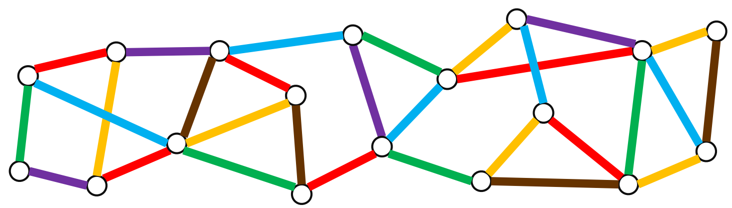

Reduction

- (Δ+1)-coloring ➜ MIS
  in Δ rounds by maximizing {1}
- MIS ➜ (Δ+1)-coloring
  in MIS$((\Delta+1)n, 2\Delta)$ rounds by simulation

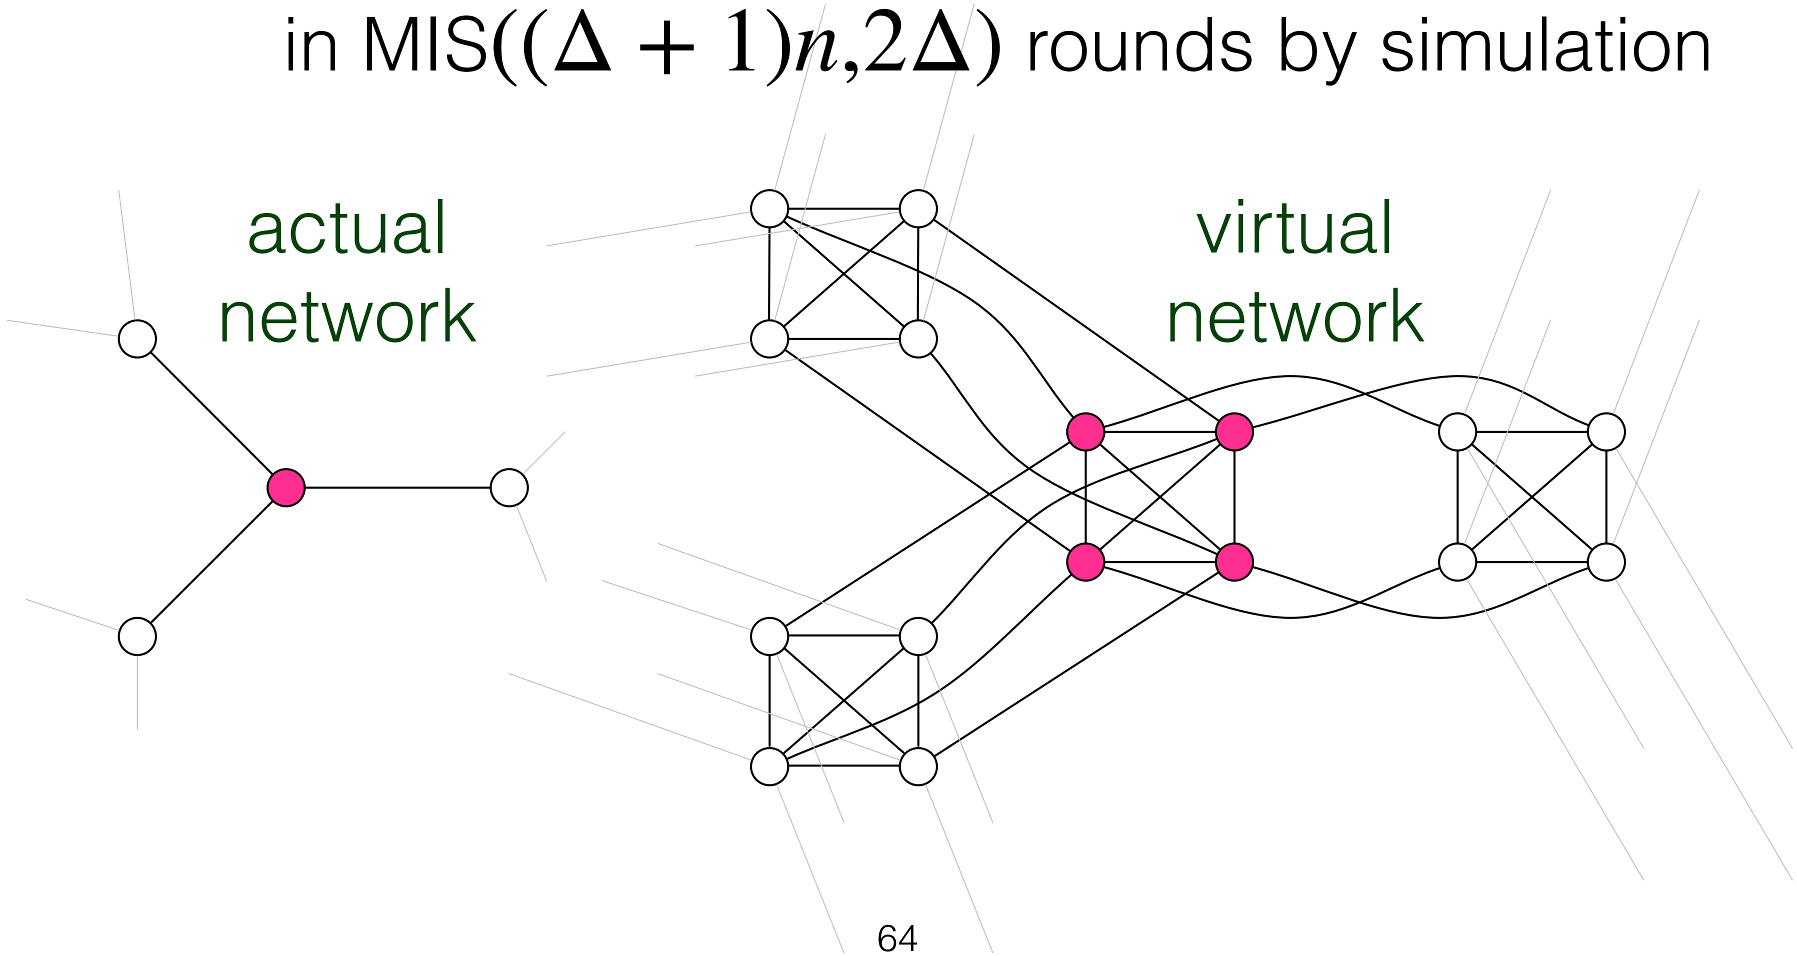

Claim 1. At most one node of each clique in the MIS

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

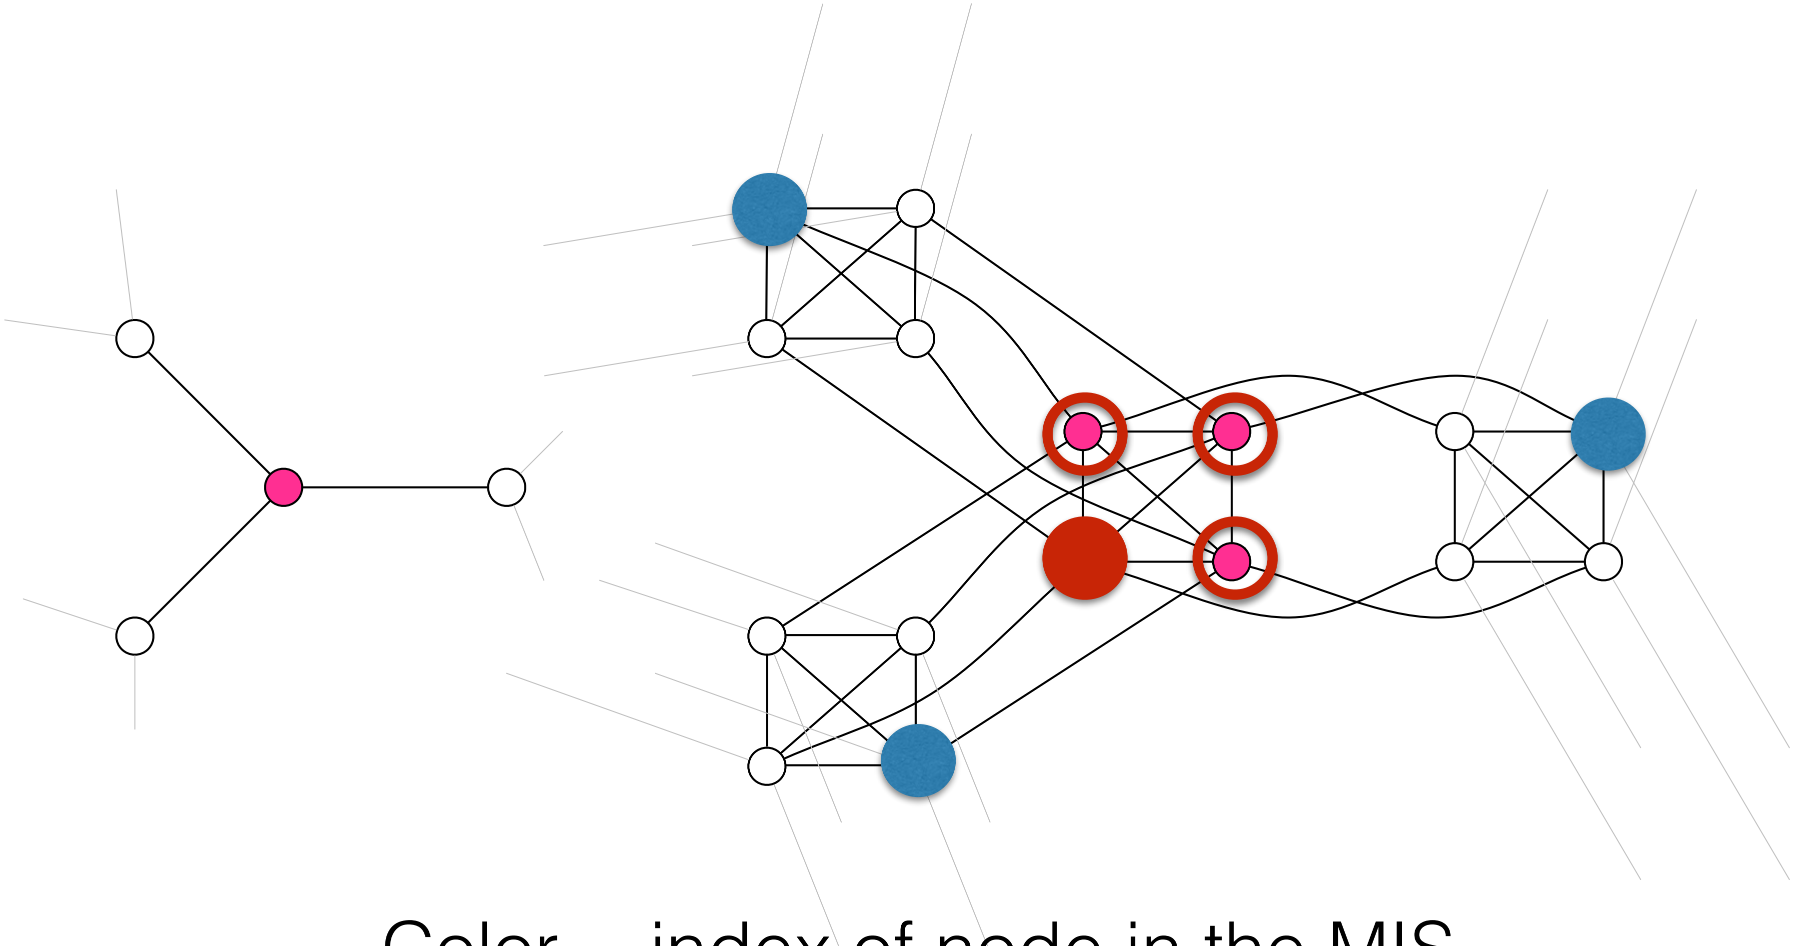

Color = index of node in the MIS

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)) $\iff$ e and e' are incident in G

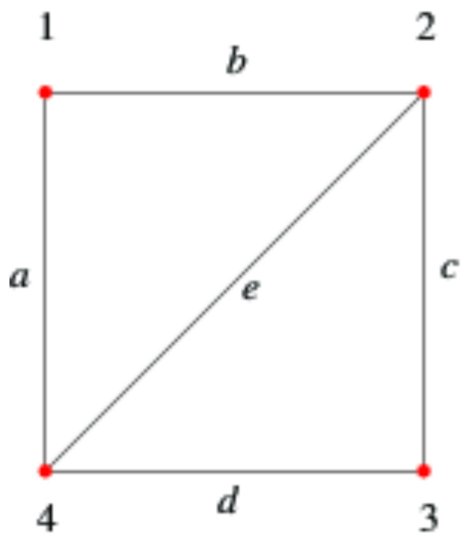

G

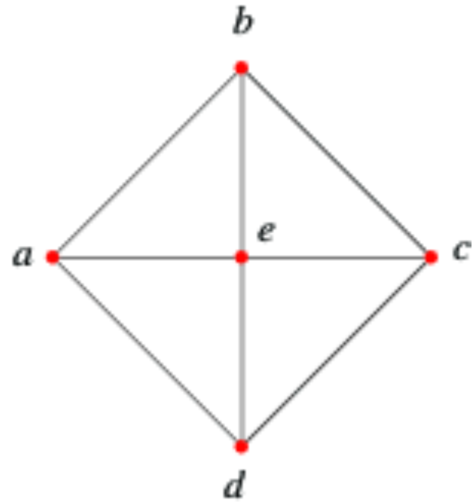

L(G)

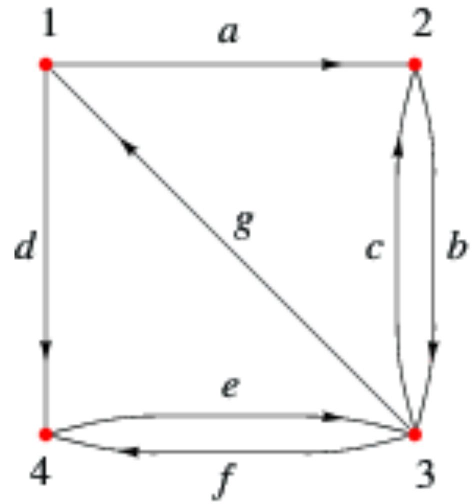

G

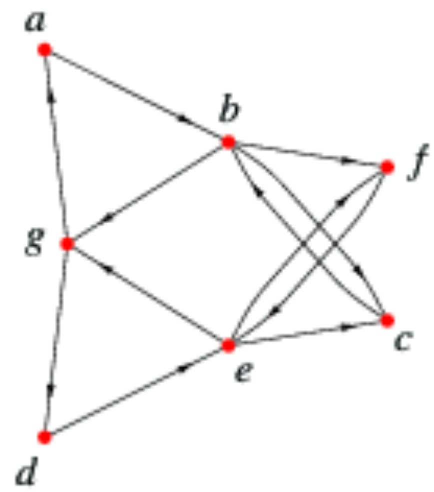

L(G)

Reductions

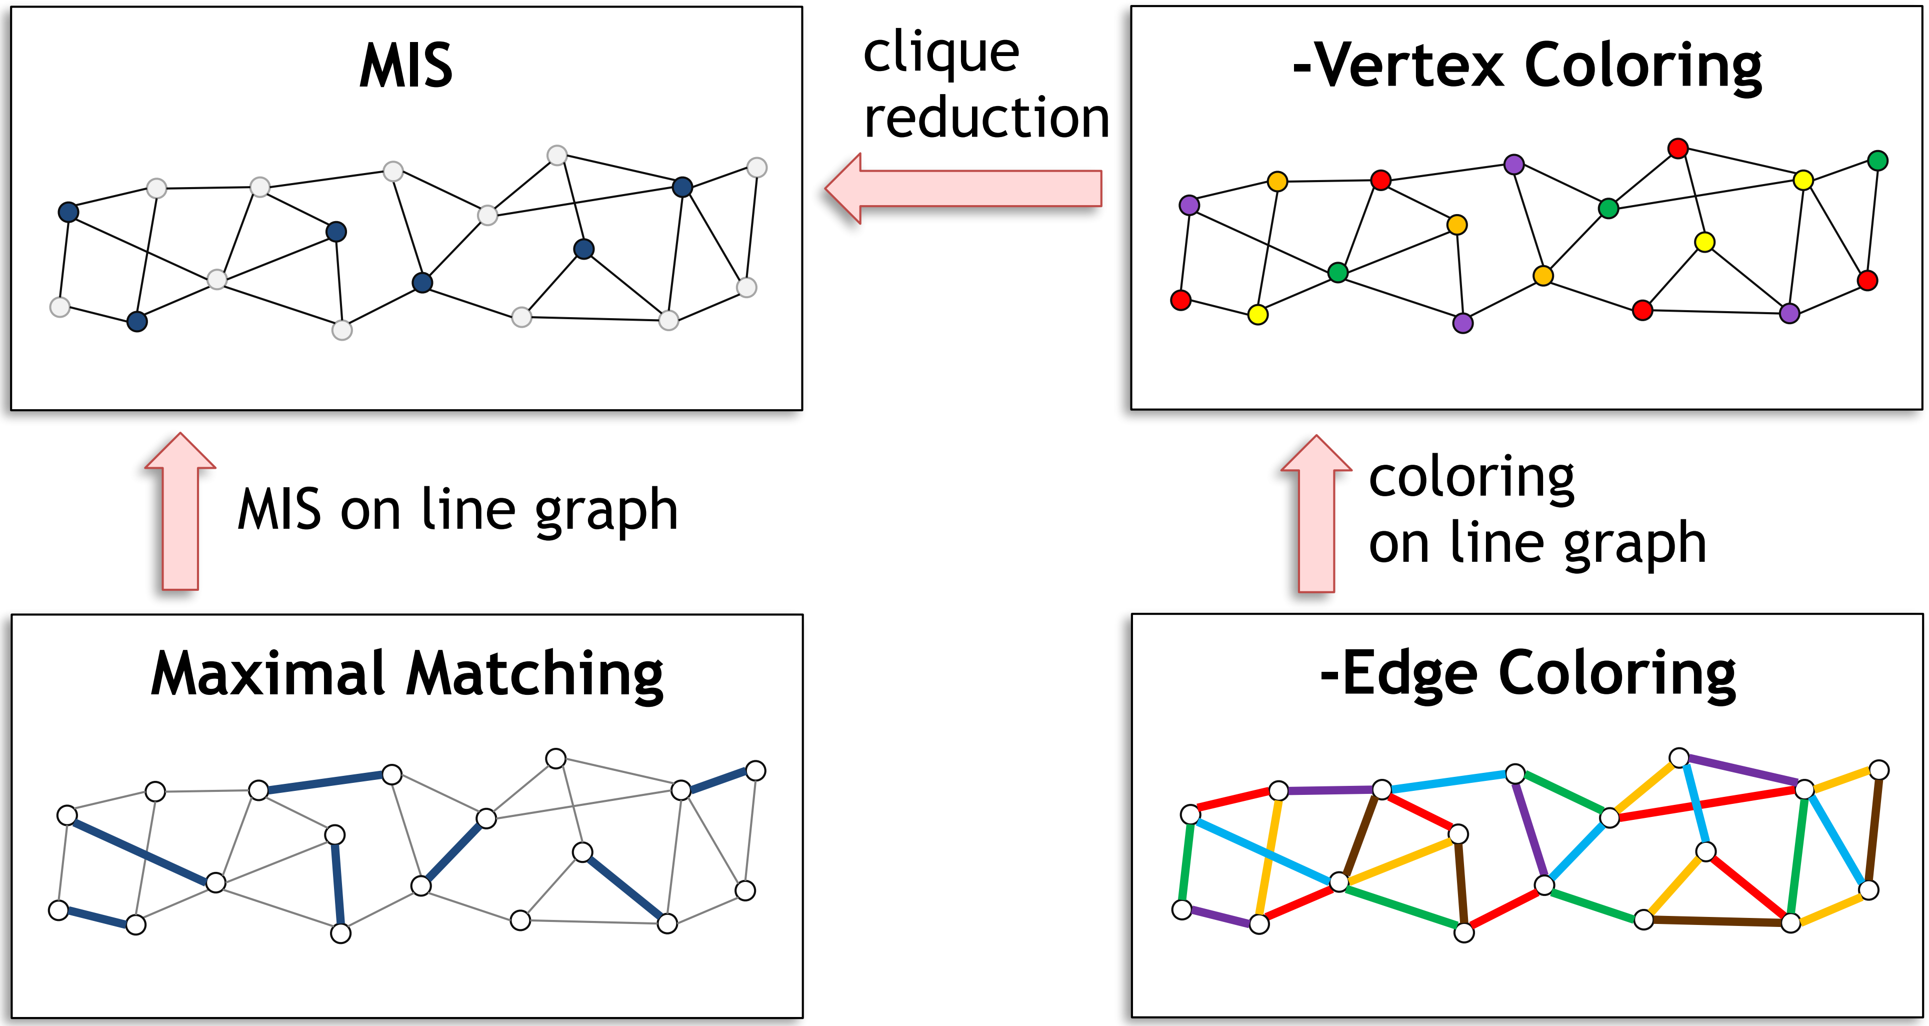
MIS
clique
reduction
-Vertex Coloring
MIS on line graph
coloring
on line graph
Maximal Matching
-Edge Coloring

Randomized algorithm for MIS

Algorithm (Luby, 1985)

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

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

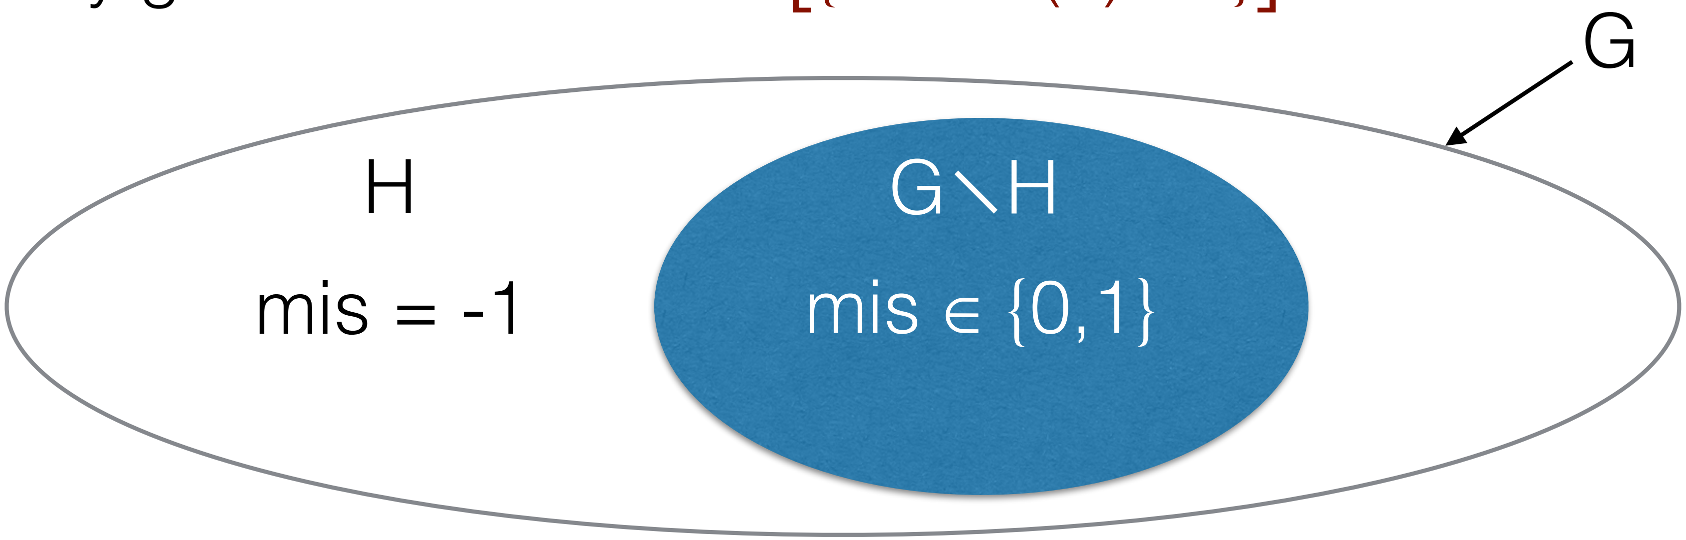

Trick: enforcing an order between nodes:

$v \succ u \iff deg_H(v) > deg_H(u)$
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 $v \in N(u)$
    free(u) ← $\nexists\, v \in N(u)$ such that $v \succ u$ and join(v)=true
    if (join(u) = true **and** free(u) = true) **then** mis(u) ← 1
    exchange mis with every $v \in N(u)$
    if (mis(u) = -1 **and** $\exists v \in N(u)$ mis(v)=1) **then** mis(u) ← 0
    exchange mis with every $v \in N(u)$

Round-Complexity of Luby's Algorithm

Remark A very similar algorithm was independently discovered by Alon, Babai, and Itai (1986).

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

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 <u>large</u> 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 degH(u) = 0 then mis(u) ← 1
else join(u) ← true with proba 1/(2 degH(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

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

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

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

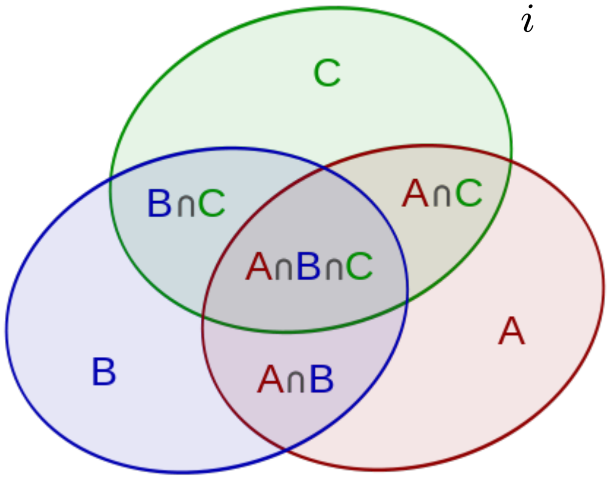

$$\begin{aligned}
\Pr[mis(u) \neq -1] &\geq \Pr[\exists v \in S : mis(v) = 1] \\
&\geq \sum_{v \in S} \Pr[mis(v) = 1] - \sum_{v,w \in S, v \neq w} \Pr[mis(v) = 1 \wedge mis(w) = 1].
\end{aligned}$$

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

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)

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, $\deg^+(u) \geq 2 \deg^-(u)$

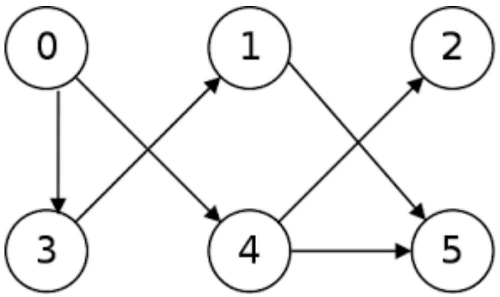

Out-degree **In-degree**

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

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

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

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

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

$\Longrightarrow \quad \displaystyle\sum_{u \text{ grand}} \deg^-(u) \ge \frac{m}{2}$ $\Longrightarrow$ at least m/2 large edges

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

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

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

Let $p = \Pr[X \le ½\, \mathbb{E}X]$

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

$$\Longrightarrow \quad p \le \frac{m - \mathbb{E}X}{m - \frac{1}{2}\mathbb{E}X} \le \frac{m - \frac{1}{2}\mathbb{E}X}{m} \le 1 - \frac{1}{144}.$$

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

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

Step 4

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

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

If $Y_i = 1$ then #edges divided by $\alpha = \dfrac{144}{143}$.

Remark: $Y \geq \log_\alpha |E(G)| = \log_{\frac{144}{143}} |E(G)| \Longrightarrow$ termination.

How big should be k?

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

Step 4 (continued)

We have $\mathbb{E}Y = kq$, so, with $\delta = 1/2$, we get

$$\Pr[Y \le kq/2] \le e^{-kq/8}$$

For $k = c \log_\alpha n$, we get

$$\Pr[Y \le cq \log_\alpha n/2] \le e^{-cq \log_\alpha n/8}$$

For $c = 4/q$ we get $cq \log_\alpha n/2 > \log_\alpha |E(G)|$
and $cq \ge 8 \ln \alpha$.

$$\Pr[Y < \log_\alpha |E(G)|] \le e^{-cq \log_\alpha n/8} = 1/n^{\frac{cq}{8 \ln \alpha}} \le 1/n$$

Thus, for $k = 4 \cdot 144 \cdot \log_{\frac{144}{143}} n$, we get
$\Pr[Y < \log_{\frac{144}{143}} |E(G)|] \le 1/n$

Thus Luby's algorithm terminates in $O(\log n)$ rounds with probability at least $1 - 1/n$. ❑

Locally Checkable Labeling

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

Definition (Naor and Stockmeyer, 1995) An LCL in $\mathscr{F}_\Delta$ 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 $k \geq \Delta+1$.

Solving an LCL Problem

Let $\Delta \geq 2$, and let L be an LCL in $\mathscr{F}_\Delta$

The LCL problem associated to L consists, for all nodes of any graph $G \in \mathscr{F}_\Delta$, to compute a label such that the collection of labels results in good balls centered at every node of G.

The definition can be generalized with inputs given to the nodes, in addition to their IDs.

The Limited Power of Randomized Algorithms

$\Pr[\text{success}] \geq 1 - 1/n$

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

For any LCL problem, its randomized (Monte Carlo) complexity on instances of size n is **at least** its deterministic complexity on instances of size $\sqrt{\log n}$

One needs to design better deterministic algorithms for improving the performances of randomized algorithms!

Proof

- Let Π be an LCL problem
- Let $\text{Det}_\Pi(n, \Delta)$ and $\text{Rand}_\Pi(n, \Delta)$ be the LOCAL deterministic and randomized complexity of Π for instance of size n and maximum degree Δ.
- We show that $\text{Det}_\Pi(n, \Delta) \leq \text{Rand}_\Pi(2^{n^2}, \Delta)$
- We assume that, initially, each node v knows its ID, as well as n and Δ, with $\text{ID}(v) \in \{0,1\}^{c \log n}$, for some $c \geq 1$.
- Let $\mathcal{G}_{n,\Delta}$ be the family of n-node graphs with maximum degree Δ, and nodes with IDs on at most $c \log n$ bits.

- Let $\mathscr{R}$ be an optimal randomized algorithm for Π
- Our aim is to construct a deterministic algorithm $\mathscr{D}$ for Π
- Assume that, in $\mathscr{G}_{n,\Delta}$, algorithm $\mathscr{R}$
  - performs in $t(n, \Delta)$ rounds, and
  - uses $r(n, \Delta)$ random bits at each node.
- Note that $|\mathscr{G}_{n,\Delta}| \leq 2^{\binom{n}{2}+cn\log n} \ll 2^{n^2}$. Let $N = 2^{n^2}$.
- For any $G \in \mathscr{G}_{n,\Delta}$ let $G' = G \cup H$ with $H \in \mathscr{G}_{N-n,\Delta}$
- In G', nodes are given (N, Δ) as input
- We have: $\Pr[\mathscr{R}$ fails in $G'] \leq 1/N$

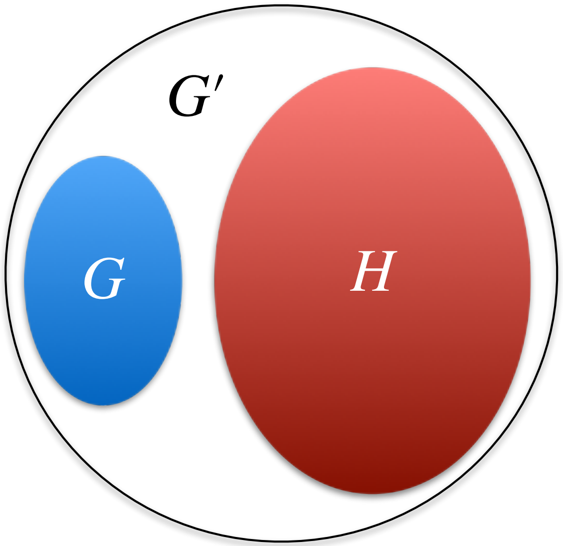

- Let us consider any function ϕ of the form

$$\phi : \{0,1\}^{c \log n} \to \{0,1\}^{r(N,\Delta)}$$

- Let $\mathscr{R}[\phi]$ be the deterministic algorithm equal to $\mathscr{R}$ with the fixed random strings determined by ϕ, i.e., node v uses $\mathscr{R}$ with bit-string $\phi(\mathrm{ID}(v))$.

- $\mathscr{R}[\phi]$ runs in $t(N, \Delta)$ rounds in G

- We said that ϕ is *bad* if $\mathscr{R}[\phi]$ fails on some $G \in \mathscr{G}_{n,\Delta}$ as part of a larger graph G'

$$\Pr_{\phi \sim unif}[\phi \text{ bad}] \le \sum_{G \in \mathcal{G}_{n,\Delta}} \Pr_{\phi \sim unif}[\mathcal{R}[\phi] \text{ errs on } G]$$

$$= \sum_{G \in \mathcal{G}_{n,\Delta}} \Pr[\mathcal{R} \text{ errs on } G \text{ with input } (N, \Delta)]$$

$$\le |\mathcal{G}_{n,\Delta}|/N < 1$$

It follows that there exists $\phi^\star : \{0,1\}^{c \log n} \to \{0,1\}^{r(N,\Delta)}$ good

Deterministic algorithm $\mathcal{D}$: on input (n, Δ), every node v

- computes $N = 2^{n^2}$, $t(N, \Delta)$ and $r(N, \Delta)$
- performs $\mathcal{R}[\phi^\star]$ for $t(N, \Delta)$ rounds

By construction $\mathcal{D}$ never errs in $\mathcal{G}_{n,\Delta}$ ❏

End Lecture 4