

+

+

Distributed Symmetry Breaking

Michael Elkin

Ben-Gurion University

joint w/ L. Barenboim
(PODC'08, STOC'09, PODC'10,
PODC'11, J. ACM'11)

and w/ L. Barenboim, S. Pettie and
J. Schneider (FOCS'12)

and w/ S. Pettie and H. Su (SODA'15)

+

1

- Running time = $\#$ rounds.
- All vertices wake up simultaneously.
- Vertices have unique Ids from $\{1, 2, \dots, n\} = [n]$.
- Arbitrarily large messages are allowed, though short (of size $O(\log n)$) are preferred.

Coloring

- $\Delta = \Delta(G)$ - maximum degree of a vertex in G .
- $\varphi : V \rightarrow [k]$ is a k -*vertex-coloring* if $\forall e = (u, w) \in E, \varphi(u) \neq \varphi(w)$.
- $\psi : E \rightarrow [t]$ is a t -*edge-coloring* if $\forall e, e'$ s.t. $e \cap e' \neq \emptyset, \psi(e) \neq \psi(e')$.
- In distributed setting, typically $k \geq \Delta + 1, t \geq 2\Delta - 1$.
- MIS U :
 - (1) $\forall v, w \in U, (v, w) \notin E$.
 - (2) $\forall v \notin U, \exists u \in U$ s.t. $(u, v) \in E$.
- MM M :
 - (1) $\forall e, e' \in M, e \cap e' = \emptyset$.
 - (2) $\forall e' \notin M, \exists e \in M$ s.t. $e \cap e' = \emptyset$.

- $(\Delta + 1)$ -coloring in $O(n)$ rounds is easy.

Color vertices one-by-one:

For each new vertex v there are

$\leq \Delta$ forbidden colors.

Hence there is always an available color for v in $[\Delta + 1]$.

- MIS in $O(n)$ rounds is easy too.

Initialize $U \leftarrow \emptyset$;

Treat vertices one-by-one:

For each new vertex v do:

if $\Gamma(v) \cap U = \emptyset$ then

v joins U ;

- $(2\Delta - 1)$ -edge-coloring reduces to $(\Delta + 1)$ -vertex-coloring,
MM and $(\Delta + 1)$ -vertex-coloring reduce to MIS.

+

+

Elementary Color Reduction Technique


Given an α -coloring, $\alpha > \Delta + 1$,
eliminate one color class in each round.

Vertices of color α form an independent set.

Each of them recolors itself into an available
color from $[\Delta + 1]$.

So in $\alpha - (\Delta + 1)$ rounds we get
a $(\Delta + 1)$ -coloring.

Continue with it for $\Delta + 1$ more rounds
to get an MIS.



+

4

+

+

Kuhn-Wattenhofer's (KW) Color Reduction Technique

$(\Delta + 1)$ -coloring in $O(\Delta \log \frac{\alpha}{\Delta + 1}) + \log^* n$ time.

[Kuhn, Wattenhofer (PODC'06)]

[Szegedy, Vishwanathan (STOC'92)]

- Given an α -coloring φ ,
 $\alpha = c \cdot (\Delta + 1)$,
 c is a large integer power of 2.

- $\forall i \in [c]$, let

$$U_i = \{v \mid (i - 1) \cdot (\Delta + 1) + 1 \leq \varphi(v) \leq i \cdot (\Delta + 1)\}.$$

- Pair subgraphs $G(U_1)$ with $G(U_2)$,
 $G(U_3)$ with $G(U_4), \dots$,
 $G(U_{c-1})$ with $G(U_c)$.

+

5

Consider $G(U_1 \cup U_2)$.

It is $2 \cdot (\Delta + 1)$ -colored by φ .

- Reduce the $2(\Delta + 1)$ -coloring of $G(U_1 \cup U_2)$ to get a $(\Delta + 1)$ -coloring of $G(U_1 \cup U_2)$ in $2(\Delta + 1) - (\Delta + 1) = \Delta + 1$ rounds.

In parallel, reduce the colorings of $G(U_3 \cup U_4), G(U_5 \cup U_6), \dots$

In $\Delta + 1$ rounds we get

$\frac{1}{2}\alpha$ -coloring of G .

- Keep halving the #colors by phases that last $\Delta + 1$ rounds each.

In $\log \frac{\alpha}{\Delta + 1}$ phases

(i.e., in $O(\Delta \cdot \log \alpha / \Delta)$ time)

we get $(\Delta + 1)$ -coloring.

- [Linial (FOCS'87)]:
 $O(\Delta^2)$ -coloring in $\log^* n$ time.

In conjunction with the **KW** color reduction we get $O(\Delta \log \Delta) + \log^* n$ time for $(\Delta + 1)$ -coloring.

- *Locally-iterative* means: in every round every vertex recolors itself based only on colors of its neighbors.

[Szegedy, Vishwanathan (STOC'92)]:
Any *locally-iterative* $(\Delta + 1)$ -coloring algorithm requires $\Omega(\Delta \log \Delta)$ time.

The $(\Delta + 1)$ -coloring algorithms of Linial and of Kuhn and Wattenhofer can be casted as locally-iterative.

So the **KW** is an optimal locally-iterative $(\Delta + 1)$ -coloring algorithm.

Distributed Coloring - Known Randomized Results

- $(\Delta + 1)$ -coloring, MIS and MM
in $O(\log n)$ time.

[Luby (STOC'85)],

[Alon,Babai,Itai (J.Alg.'86)],

[Israeli,Itai (IPL'86)].

$(\Delta + 1)$ -vertex-coloring

in $O(\log \Delta + \sqrt{\log n})$ time.

[Schneider,Wattenhofer (PODC'10)].

- $O(\Delta)$ -vertex-coloring in $O(\sqrt{\log n})$ time
[Kothapalli,Scheideler,Onus,
Schindelbauer (IPDPS'06)].

- $O(\Delta + \log n)$ -vertex-coloring

in $O(\log \log n)$ time,

and $O(\Delta \log^{(c)} n + \log^{1+1/c} n)$ -coloring in
 $O(f(c)) = O(1)$ time.

[Schneider,Wattenhofer (PODC'10)].

+

+

New Randomized Algorithms

[Barenboim, E., Pettie, Schneider (FOCS'12)]

- MM in $O(\log \Delta + \log^4 \log n)$ time.
- $(\Delta + 1)$ -coloring in $O(\log \Delta) + \exp\{O(\sqrt{\log \log n})\}$ time.
- $O(\Delta)$ -coloring in $\exp\{O(\sqrt{\log \log n})\}$ time.
- $\Delta^{1+\eta}$ -coloring in $O(\log^2 \log n)$ time.
- MIS in $O(\log^2 \Delta) + \exp\{O(\sqrt{\log \log n})\}$ time.
- $\Delta^{1+\eta}$ -edge-coloring in $O(\log \log n)$ time.

+


7

+

+

Basic Approach in BEPS's algorithms

- Do (roughly) $O(\log \Delta)$ "Luby steps" to break the graph into disconnected components of size $s \leq \text{polylog}(n)$.



$$|C1|, |C2|, |C3| \leq s.$$

+

8

- Use the state-of-the-art *deterministic* MIS algorithm for each component.

It completes the MIS within additional $\exp\{O(\sqrt{\log s})\} \leq \exp\{O(\sqrt{\log \log n})\}$ time.
[Panconesi, Srinivasan'92]

Using randomized subroutine within components fails because the failure probability is $1/\text{poly}(s) \approx 1/\text{polylog}(n)$.

- Works similarly for $(\Delta + 1)$ -vertex-coloring and MM problems.

For MM the second step requires just $O(\log^4 s) = O(\log^4 \log n)$ time.
[Hanckowiak, Karonski, Panconesi'99]

- Improved deterministic algorithms give rise to improved randomized ones!

+

+

New Randomized Algorithms (Cont.)

[E., Pettie, Su (SODA'15)]

$(2\Delta - 1)$ -edge-coloring
in $\exp\{O(\sqrt{\log \log n})\}$ time.

Establishes a *separation* between
MM and $(2\Delta - 1)$ -edge-coloring.

MM requires $\Omega(\sqrt{\log n})$ time.

[Kuhn, Moscibroda, Wattenhofer'10].

+

+

+

Open if such a separation exists also between MIS and $(\Delta + 1)$ -vertex-coloring.

Our algorithm (in SODA'15) gives actually $(\Delta + 1)$ -vertex-coloring in $O(\log 1/\epsilon) + \exp\{O(\sqrt{\log \log n})\}$ time for $(1 - \epsilon)$ -*locally-sparse* graphs.

Def: $G = (V, E)$ of max' degree Δ is $(1 - \epsilon)$ -*locally-sparse* if

$$\forall v, |E(\Gamma(v))| < (1 - \epsilon) \binom{\Delta}{2}$$

Line graphs are $(1/2 + o(1))$ -locally-sparse, implying our result for edge-coloring.

+

+

+

More Details about MIS Algorithm (BEPS)

Procedure Halve -
reduces the degree from Δ to $\Delta/2$.

Consists of two parts.

Part I: Runs κ Luby trials.

κ will be set

- (1) either as $\kappa = O(\sqrt{\log n})$,
- (2) or as $\kappa = O(\log \Delta)$.

The graph of remaining high-degree vertices breaks into connected components.

In case (1) these components have weak diameter $O(\sqrt{\log n})$.

In case (2) they have size $O(\Delta^4 \cdot \log n)$.

+

+

+

Procedure Halve (Cont.)

To complete computing the MIS for high-degree vertices ($deg \geq \Delta/2$) we need

(1) $O(\sqrt{\log n})$ time in case (1) via brute-force.

(2) $\exp\{O(\sqrt{\log O(\Delta^4 \log n)})\}$ time in case (2) via the state-of-the-art deterministic MIS algorithm due to [Panconesi, Srinivasan'92].

So Procedure Halve requires total

(1) $O(\sqrt{\log n})$ time in case (1).

(2) $O(\log \Delta) + \exp\{O(\sqrt{\log \Delta + \log \log n})\} = \exp\{O(\sqrt{\log \Delta + \log \log n})\}$ time in case (2).

+

Hence the entire running time is $O(\sqrt{\log n} \cdot \log \Delta)$ in case (1),
and $O(\log \Delta \cdot \exp\{O(\sqrt{\log \Delta + \log \log n})\})$
in case (2).

The latter expression is $\exp\{O(\sqrt{\log \log n})\}$,
whenever $\Delta \leq \text{polylog}(n)$.

A better bound
 $O(\log^2 \Delta + \exp\{O(\sqrt{\log \log n})\})$:
not in this talk.

+

+

Part I of Procedure Halve

Run κ Luby trials.

$$I_1 \leftarrow \emptyset, A_1 \leftarrow V.$$

I_i - the independent set before the beginning of iteration i .

$A_i \leftarrow V \setminus \hat{\Gamma}(I_i)$ - active vertices before iteration i .

$I_{\kappa+1}$ - the independent set after all the κ Luby trials.

Part II works on $G(V \setminus \hat{\Gamma}(I_{\kappa+1}))$.

This subgraph decomposes into components of either bounded weak diameter or size.

+

+

+

Procedure Halve: Part I

Run for κ iterations ($i = 1, 2, \dots, \kappa$)
on the *entire* graph G , i.e.,
on both high-degree ($\text{deg} \geq \Delta/2$)
and low-degree ($\text{deg} < \Delta/2$) vertices.

(Assume the maximum degree is at most Δ .)

1. Each $v \in A_i$ selects itself
with probability $\frac{1}{\Delta+1}$, i.a.r.
2. $I_{i+1} \leftarrow I_i \cup \{v \mid v \text{ is the only vertex in } \hat{\Gamma}(v) \text{ which selected itself}\}$

+

+

+

Procedure Halve: Analysis of Part I

$\Gamma_i(u)$ - set of active neighbors of a vertex u at the beginning of iteration i .

$$\deg_i(u) = |\Gamma_i(u)|$$

Def: A subset $S \subseteq V$ is *well-separated* if $\forall u, u' \in S, u \neq u', \text{dist}_G(u, u') \geq 5$.

Set-Survival Lemma:


Let $S \subseteq A_i$ be a well-separated set of active vertices with $\deg_i(u) \geq \Delta/2$ each.

Then

$$\mathbb{IP}(S \subseteq A_{i+1}) \leq p^{|S|},$$

for some *constant* $0 < p < 1$.

+



Neighborhoods of neighbors are disjoint.

In a particular Luby trial,
survival of a vertex u depends only
upon coins tosses of its 2-neighborhood.

Hence vertices u, u' at distance ≥ 5
from one another are independent.

+

+

Proof of Set-Survival Lemma

Sufficient to prove

$$\mathbb{IP}(u \in A_{i+1} \mid u \in A_i) \leq p .$$

Let $u = v_0, \{v_1, v_2, \dots, v_d\} = \Gamma_i(u)$.

By Lemma's assumption,

$$d = \deg_i(u) \geq \Delta/2 .$$

$$\begin{aligned} & \mathbb{IP}(\text{some vertex } \in \{v_0, \dots, v_d\} \text{ selects itself}) \\ &= 1 - \left(1 - \frac{1}{\Delta + 1}\right)^{d+1} \\ &\geq 1 - \left(1 - \frac{1}{\Delta + 1}\right)^{\Delta/2+1} > 1 - e^{-1/2} . \end{aligned}$$

+

+


+

Proof of Set-Survival Lemma (Cont.)

Conditioned upon the event

some vertex $x \in \{v_0, \dots, v_d\}$ selects itself:

j - the smallest index such that
 v_j selects itself.



$$H = \Gamma_i(v_j) \setminus \{v_0, \dots, v_j\}$$

+

$$\begin{aligned}
& \mathbf{IP}(v_j \text{ joins MIS} \mid v_j \text{ selects itself,} \\
& \quad j \text{ is the smallest}) = \\
& \mathbf{IP}(\Gamma_i(v_j) \setminus \{v_0, \dots, v_{j-1}\} \\
& \quad \text{did not select themselves}) = \\
& = \left(1 - \frac{1}{\Delta + 1}\right)^{|\Gamma_i(v_j) \setminus \{v_0, \dots, v_{j-1}\}|} \geq \\
& \geq \left(1 - \frac{1}{\Delta + 1}\right)^\Delta \geq e^{-1}.
\end{aligned}$$

$$\begin{aligned}
& \mathbf{IP}(u \text{ becomes inactive}) \geq \\
& \mathbf{IP}(\text{some } v_j \text{ selects itself}) \cdot \mathbf{IP}(v_j \in I_{i+1} \mid v_j \text{ exists}) \\
& \geq (1 - e^{-1/2}) \cdot e^{-1} .
\end{aligned}$$

$$\mathbf{IP}(u \text{ survives}) = 1 - (1 - e^{-1/2}) \cdot e^{-1} = p .$$

$$0 < p < 1$$

QED

+

+

Analysis of Proc. Halve: Long-Path Lemma, Notation

$$U = \{v \in A_{\kappa+1} \mid \deg_{\kappa+1}(u) \geq \Delta/2\}$$

Vertices that survived all Luby trials and retained high degree.

Set $\kappa = c \cdot \sqrt{\log n}$ (i.e., case (1)).

\mathcal{P} - the set of all paths (not necessarily shortest) which involve only vertices that *originally had high degree* between pairs u, v of *remote* vertices, i.e., s.t. $\text{dist}_G(u, v) \geq 5 \cdot \sqrt{\log n}$.

+

+

+

Long-Path Lemma

Lemma: Each $P \in \mathcal{P}$ contains a well-separated set $Q(P)$ of size $|Q(P)| = \sqrt{\log n}$.

Proof:

$$P = (u = u_0, u_1, \dots, u_\ell = v),$$

$$\ell \geq 5\sqrt{\log n}.$$

Set $q_1 \leftarrow u_0$.

Suppose we have already built q_1, q_2, \dots, q_j , $j \geq 1$, with pairwise distances ≥ 5 , and s.t. $\text{dist}_G(q_j, v) \geq 5$.

+

+

+


Proof of Long-Path Lemma: Cont.

$q_{j+1} \leftarrow u_i \in P$ with
the highest index s.t. $\text{dist}_G(q_j, u_i) = 5$.

(Exists, because $\text{dist}_G(q_j, v) \geq 5$.)

Claim: $\forall k, 1 \leq k \leq j,$
 $\text{dist}_G(q_k, q_{j+1} = u_i) \geq 5.$

+



Proof (of Claim):

Otherwise $\text{dist}_G(q_k, u_i) < 5$,
 but $\text{dist}_G(q_k, u_\ell) = \text{dist}_G(q_k, v) \geq 5$.

Hence $\exists i', i < i' \leq \ell$ s.t.

$\text{dist}_G(q_k, u_{i'}) = 5$, contradicting
 the maximality of the index h ,

$u_h = q_{k+1}$.

QED (Claim)

+

+

Proof of Long-Path Lemma: Cont. II

$$\text{dist}_G(q_1, v) = \ell, \text{dist}_G(q_1, q_2) = 5,$$

$$\begin{aligned} \text{dist}_G(q_2, v) + \text{dist}_G(q_1, q_2) &\geq \\ \text{dist}_G(q_1, v) &= 5\sqrt{\log n}. \end{aligned}$$

Hence

$$\text{dist}_G(q_2, v) \geq 5(\sqrt{\log n} - 1),$$

$$\text{dist}_G(q_3, v) \geq 5(\sqrt{\log n} - 2),$$

etc.

We obtain $\sqrt{\log n}$ vertices
 $q_1, q_2, \dots, q_{\sqrt{\log n}}$ in a well-separated set.

QED (of Lemma)

+

+

+

Small Components' Lemma, Preliminaries

$$\mathcal{Q} = \{Q(P) \mid P \in \mathcal{P}\}$$

The set of well-separated subsets obtained from paths of high-degree vertices whose endpoints are at distance $\geq 5\sqrt{\log n}$ one from another in G .

\mathcal{W} - the set of walks of length $5\sqrt{\log n}$.

Every $Q(P) \in \mathcal{Q}$ can be mapped into a walk $W(P) \in \mathcal{W}$.

(By connecting consecutive q_i, q_{i+1} in $Q(P)$ by an arbitrary 5-path.)

$$|\mathcal{Q}| \leq |\mathcal{W}|$$

+

Also,

$$|\mathcal{W}| \leq n \cdot \Delta^{5\sqrt{\log n}-1}.$$

(n possibilities to pick the first vertex, and Δ possibilities to pick each consecutive one.)

Hence

$$|\mathcal{Q}| \leq |\mathcal{W}| \leq n \cdot \Delta^{5\sqrt{\log n}-1}.$$

+

+

Small Components' Lemma

Lemma: After $\kappa = c\sqrt{\log n}$
Luby trials, the vertex set U
(of vertices w with $\deg_{\kappa+1}(w) \geq \Delta/2$)
induces connected components of
weak diameter $\leq 5 \cdot \sqrt{\log n}$, whp.

(Meaning that the maximum distance *in* G
between every pair of vertices w, w' in the
same component is $\leq 5\sqrt{\log n}$.)

+

+

+

Proof of Small Components' Lemma

u, v - a pair of high-degree vertices
with $\text{dist}_G(u, v) \geq 5\sqrt{\log n}$.

$P \in \mathcal{P}$ - some path between them that consists
of only high-degree vertices.

$$\mathbb{P}(Q(P) \text{ survives a single iteration}) \leq p^{|Q(P)|} .$$

The probability of the entire $Q(P)$ to survive
 κ iterations in all of which all its vertices keep
being high-degree is

$$\leq p^{|Q(P)| \cdot \kappa} = \exp\{-\Omega(\sqrt{\log n} \cdot c\sqrt{\log n})\} = n^{-c}$$

+

+

+

Proof of Small Components' Lemma, Cont.

Number of possible Q -sets is

$$|Q| \leq |W| \leq n \cdot \Delta^{5\sqrt{\log n}}.$$

For $\Delta \leq 2^{\sqrt{\log n}}$, we get a bound of $n \cdot 2^{5 \log n} = n^6$.

Probability that some $Q(P)$ survives is $\leq \frac{n^6}{n^c} = \frac{1}{n^{c-6}} = \frac{1}{\text{poly}(n)}$.

Hence, whp, every path $P \in \mathcal{P}$ which connects a pair u, v with $\text{dist}_G(u, v) \geq 5\sqrt{\log n}$ and such that all vertices of $V(P)$ are high-degree,

does not (entirely) survive, i.e., $V(P) \not\subseteq U$.

Hence, whp, only vertices u, v with $\text{dist}_G(u, v) < 5\sqrt{\log n}$ may be connected *in* U .

+

For $\Delta \geq 2^{\sqrt{\log n}}$, the Luby's bound of $O(\log n)$ is better than the bound of $O(\log \Delta \sqrt{\log n})$.

So in both cases we have

MIS in $O(\log \Delta \sqrt{\log n})$ time.

In the paper we show a much better bound of $O(\log^2 \Delta) + \exp(O(\sqrt{\log \log n}))$.

+

+

Lower Bounds vs. Upper Bounds

- $f(\Delta)$ -coloring requires $\frac{1}{2} \log^* n$ time.

[Linial (FOCS'87)]

The upper bound (BEPS)
for $(\Delta + 1)$ -coloring is
 $O(\log \Delta) + \exp\{O(\sqrt{\log \log n})\}$.

Huge gap!

- Coloring Δ -regular trees in $o(\sqrt{\Delta})$ colors requires $\omega(\log_{\Delta} n)$ time.

[Linial (FOCS'87)]

One can color unoriented forests in
 Δ^ϵ colors within $O(\log_{\Delta} n)$ time,
for an arbitrarily small $\epsilon > 0$.

[Barenboim, E. (PODC'08)] (tight).

+

- $\Omega(\log \Delta)$ and $\Omega(\sqrt{\log n})$ time is required for MIS and MM.

[Kuhn, Moscibroda, Wattenhofer],
[(PODC'04), (ArXiv'10)]

The upper bound (BEPS) for MM is $O(\log \Delta + \log^4 \log n)$.

Tight for $\log^4 \log n \leq \log \Delta \leq \sqrt{\log n}$.

For MIS the BEPS's upper bound is $O(\log^2 \Delta) + \exp\{O(\sqrt{\log \log n})\}$.

+

+

Known Deterministic Results

- $(\Delta + 1)$ -coloring and MIS in $O(\Delta^2 + \log^* n)$ time, and in $O(\Delta \log n)$ time.
[Goldberg, Plotkin, Shannon'87]
(based on [Cole, Vishkin'86])
- $O(\Delta^2)$ -coloring in $\log^* n + O(1)$ time.
[Linial'87]

Asked: can one get much fewer than Δ^2 colors in time polylogarithmic in n ?

- $(\Delta + 1)$ -coloring and MIS in $2^{O(\sqrt{\log n})}$ time. (Large messages)
[Panconesi, Srinivasan'92], based on [Awerbuch, Goldberg, Luby, Plotkin'89]

+

- MM in $O(\log^4 n)$ time.

[Hanckowiak, Karonski, Panconesi'99]

- $O(\Delta \cdot \log n)$ -edge-coloring in $O(\log^4 n)$ time.

[Czygrinow, Hanckowiak, Karonski (ESA'01)]

+

+

New Deterministic Results

- $(\Delta + 1)$ -coloring and MIS
in $O(\Delta) + \frac{1}{2} \log^* n$ time.

[Barenboim, E. (ArXiv'08, STOC'09)],
[Kuhn (SPAA'09)]

Breaks the Szegedy-Vishwanathan's
 $\Omega(\Delta \log \Delta)$ barrier.

Major Open Problem:

The lower bound is only $\frac{1}{2} \cdot \log^* n$
([Linial'87]),

while the upper bound is $O(\Delta) + \frac{1}{2} \log^* n$.

+

- (1) $\Delta^{1+\eta}$ -coloring in $O(\log \Delta \cdot \log n)$ time, for any $\eta > 0$.
- (2) $O(\Delta)$ -coloring in $O(\Delta^\epsilon \cdot \log n)$ time, for any $\epsilon > 0$.

[Barenboim, E. (PODC'10, J.ACM'11)]

Answers Linial's question in the affirmative.

(In polylogarithmic time one can get $\Delta \cdot 2^{O(\log \Delta / \log \log \Delta)}$ -coloring.)

- (1) $\Delta^{1+\eta}$ -edge-coloring in $O(\log \Delta + \log^* n)$ time, for any $\eta > 0$.
- (2) $O(\Delta)$ -edge-coloring in $O(\Delta^\epsilon + \log^* n)$ time, for any $\epsilon > 0$.

[Barenboim, E. (PODC'11)]

+

Basic Building Blocks for Further Progress

+

- **Defective coloring:**

For $(\Delta + 1)$ -coloring in $O(\Delta) + \log^* n$ time.
[Barenboim, E. (STOC'09)],
[Kuhn (SPAA'09)]

Enables one to bypass the
Szegedy-Vishwanathan's barrier
of $\Omega(\Delta \log \Delta)$ for
locally-iterative algorithms.

- **Arbdefective coloring:**

For $\Delta^{1+\eta}$ -coloring
in $O(\log \Delta \cdot \log n)$ deterministic time.
[Barenboim, E. (PODC'10, J.ACM'11)]

Answering in the affirmative
Linial's open question.

+

+

$(\Delta + 1)$ -Coloring in
 $O(\Delta) + \log^* n$ Time
(Defective Coloring)

[Burr, Jacobson'85], [Harary, Jones'86]

[Cowen, Cowen, Woodall'86]

Def: The *defect* of a vertex v wrt coloring φ is the number of neighbors $u \in \Gamma(v)$ with $\varphi(u) = \varphi(v)$.

Def: The *defect* d of a k -coloring φ is the maximum defect of a vertex wrt φ . φ is called a *d -defective k -coloring*.

Thm: [Lovasz'66]

$\forall G, \forall p$ there exists

a $\lfloor \Delta/p \rfloor$ -defective p -coloring of G .

+

+

+

Proof of Lovasz's Thm

φ - an arbitrary p -coloring.

(Not necessarily legal or Δ/p -defective.)


while $\exists v$ with $\text{defect}(v) > \Delta/p$ *do*

{

$\varphi(v) \leftarrow$ the color used by

min. $\#$ neighbors of v ;

}



$\Delta = 5,$
 $p = 2,$
 there exists
 a color used
 by $2 < 5/2$
 neighbors

+

ME_i - the total #monochromatic edges before iteration i starts.

$$ME_{i+1} = ME_i - \text{defect}(v) + \lfloor \frac{\Delta}{p} \rfloor < ME_i.$$

But $0 \leq ME_i \leq |E|$, and so within a finite number of iterations this process terminates.

+

+

Distributed Counterparts of Lovasz's Theorem

Thm: [Barenboim, E. (STOC'09)]

$\forall G, \forall p \lfloor \Delta/p \rfloor$ -defective $O(p^2)$ -coloring of G
can be computed in $O(\Delta^\epsilon) + \frac{1}{2} \log^* n$ time,
 $\forall \epsilon > 0$.

Thm: [Kuhn (SPAA'09)]

$\forall G, \forall p \lfloor \Delta/p \rfloor$ -defective $O(p^2)$ -coloring of G
can be computed in $O(\log^* \Delta) + \frac{1}{2} \log^* n$ time.

Open: can one efficiently achieve
a linear (in Δ) product of defect and #colors?

Partial answer: for *edge*-coloring
it is possible.

Also, for vertex-coloring of graphs
with bounded independence.

[Barenboim, E. (PODC'11)]

+

+

+

$(\Delta + 1)$ -Coloring Algorithm

- Compute $O\left(\frac{\Delta}{\log \Delta}\right)$ -defective $\log^2 \Delta$ -coloring of G in $o(\Delta) + O(\log^* n)$ time.
($p = \log \Delta$)

- Each color class induces a subgraph with maximum degree $\Delta' = O\left(\frac{\Delta}{\log \Delta}\right)$.

Subgraphs are vertex-disjoint.

- In parallel, compute $(\Delta' + 1)$ -coloring in each of the $\log^2 \Delta$ subgraphs in $O(\Delta' \log \Delta' + \log^* n) = O(\Delta + \log^* n)$ time, using **KW** algorithm.

- Overall we get $O((\Delta' + 1) \log^2 \Delta) = O(\Delta \log \Delta)$ -coloring φ of the entire original graph.

(Using distinct palettes.)

+

- Invoke KW iterative procedure.

Given α -coloring it returns
 $(\Delta + 1)$ -coloring in $O(\Delta \cdot \log \frac{\alpha}{\Delta})$ time.

For $\alpha = \Delta \log \Delta$,
the time is $O(\Delta \log \log \Delta)$.

Overall running time is
 $O((\Delta + 1) \cdot \log \log \Delta + \log^* n) + o(\Delta)$.

This is a self-improving scheme!

Now we have $(\Delta + 1)$ -coloring algorithm that
runs in $O(\Delta \log \log \Delta + \log^* n)$ time.

- Compute $O\left(\frac{\Delta}{\log \log \Delta}\right)$ -defective
 $(\log \log \Delta)^2$ -coloring
in $o(\Delta) + O(\log^* n)$ time.
- $\Delta' = \frac{\Delta}{\log \log \Delta}$.

Compute $(\Delta' + 1)$ -coloring of each subgraph in

$O(\Delta' \log \log \Delta' + \log^* n) = O(\Delta + \log^* n)$ time.

- Combine these colorings into an $O(\Delta \log \log \Delta)$ -coloring of G (in zero time).
- Reduce the $O(\Delta \cdot \log \log \Delta)$ -coloring via **KW** iterative procedure into a $(\Delta + 1)$ -coloring within $O(\Delta \cdot \log^{(3)} \Delta + \log^* n)$ additional time.

Overall we get $(\Delta + 1)$ -coloring in $O(\Delta \cdot \log^{(3)} \Delta + \log^* n)$ time.



Repeating this argument $\log^* \Delta$ times we get $(\Delta + 1)$ -coloring in $O(\Delta + \log^* n)$ time.

+

+

A tradeoff (an application)

$\forall t$, $O(\Delta \cdot t)$ -coloring in $O(\Delta/t + \log^* n)$ time.
 (Interpolates between Linial's $O(\Delta^2)$ -coloring in $\log^* n$ time, and our $(\Delta + 1)$ -coloring in $O(\Delta + \log^* n)$ time.)

- Compute (Δ/t) -defective $O(t^2)$ -coloring in $O(\log^* n)$ time.
- We get $O(t^2)$ vertex-disjoint subgraphs, each with $\Delta' \leq \Delta/t$.

Compute $(\Delta' + 1)$ -coloring of each, in parallel, in

$$O(\Delta' + \log^* n) = O(\Delta/t + \log^* n),$$

using the last result for $(\Delta' + 1)$ -coloring.

- Combine the colorings in zero time to get $O(t^2 \cdot \Delta') = O(\Delta \cdot t)$ -coloring, in total $O(\Delta/t + \log^* n)$ time.

+

34

+

+

Open Questions

1. A $(\Delta + 1)$ -coloring or an MIS in deterministic polylogarithmic time?

Or at least $O(\Delta)$ -coloring.

Currently we have $\Delta \cdot 2^{O(\frac{\log \Delta}{\log \log \Delta})}$ -coloring.

2. A $\Delta^{2-\epsilon}$ -coloring in sublogarithmic deterministic time?

3. A $(\Delta + 1)$ -coloring in $o(\Delta) + \frac{1}{2} \log^* n$ time?
Or a lower bound?

Currently we have $O(\Delta) + \frac{1}{2} \log^* n$ time.

+

4. Δ/p -defective $O(p)$ -coloring in deterministic polylogarithmic time?
(Known for edge-coloring, and for vertex-coloring of graphs with bounded neighborhood independence.)
5. $(2a + 1)$ -coloring faster than in $O(a^2 \log n)$ time?
 $(2 + \eta) \cdot a$ -coloring faster than in $O(a \log n)$ time?

We know

$(2 + \eta)^{1/\epsilon} a$ -coloring in $O(a^\epsilon \cdot \log n)$ time,
and $a^{1+\eta}$ -coloring in $O(\log a \cdot \log n)$ time.

There is also a *lower bound* of $\Omega\left(\frac{\log n}{\log a}\right)$
for $O(a^2)$ -coloring.

So unlike graphs with bounded degree,
for graphs of bounded arboricity one
cannot hope for sublogarithmic time.

6. MIS or MM in randomized $o(\log n)$ time, for *all* values of Δ (or a)?

7. Randomized MIS in planar graphs in $o(\log^{2/3} n)$ time?
Or a lower bound?

+

+

Thank you!!

+