Tight Analysis of Parallel Randomized Greedy MIS

Manuela Fischer and Andreas Noever ETH Zurich

arbitrary graph G on n nodes



arbitrary graph G on n nodes

Independent:

set of non-adjacent nodes



arbitrary graph G on n nodes

Independent:

set of non-adjacent nodes

<u>Maximal:</u> no node can be added without violating independence



arbitrary graph G on n nodes

Independent: set of non-adjacent nodes

<u>Maximal:</u> not maximum! no node can be added without violating independence





For an arbitrary order of nodes,



For an arbitrary order of nodes, if current node is still alive,













For an arbitrary order of nodes,















For an arbitrary order of nodes, in each round, every local minimum joins MIS, and removes itself along with its neighbors.

<u>Complexity</u>: number of rounds



For an arbitrary order of nodes, in each round, every local minimum joins MIS, and removes itself along with its neighbors.



For an arbitrary order of nodes, in each round, every local minimum joins MIS, and removes itself along with its neighbors.



For an arbitrary order of nodes, in each round, every local minimum joins MIS, and removes itself along with its neighbors.



For an arbitrary order of nodes, in each round, every local minimum joins MIS, and removes itself along with its neighbors.



For an arbitrary order of nodes, in each round, every local minimum joins MIS, and removes itself along with its neighbors.



For an arbitrary order of nodes, in each round, every local minimum joins MIS, and removes itself along with its neighbors.



For an arbitrary order of nodes, in each round, every local minimum joins MIS, and removes itself along with its neighbors.



For an arbitrary order of nodes, in each round, every local minimum joins MIS, and removes itself along with its neighbors.



Randomized Parallel Greedy Algorithm

random For an arbitrary order of nodes, in each round, every local minimum joins MIS, and removes itself along with its neighbors.



<u>Complexity</u>: number of rounds

Randomized Parallel Greedy Algorithm

Randomized Parallel Greedy Algorithm

Random Graphs

Randomized Parallel Greedy Algorithm

Random Graphs

$$O\left(\frac{\log^2 n}{\log\log n}\right)$$

Coppersmith, Raghavan, Tompa [FOCS'87]
Random Graphs

$$O\left(\frac{\log^2 n}{\log\log n}\right)$$

Coppersmith, Raghavan, Tompa [FOCS'87]

 $O(\log n)$

Calkin, Frieze [RSA'90]

Random Graphs

$$O\left(\frac{\log^2 n}{\log\log n}\right)$$

Coppersmith, Raghavan, Tompa [FOCS'87]

 $O(\log n)$

Calkin, Frieze [RSA'90]



Random Graphs

$$O\left(\frac{\log^2 n}{\log\log n}\right)$$

Coppersmith, Raghavan, Tompa [FOCS'87]

 $O(\log n)$

Calkin, Frieze [RSA'90]



General Graphs

Random Graphs

$$O\left(\frac{\log^2 n}{\log\log n}\right)$$

Coppersmith, Raghavan, Tompa [FOCS'87]

 $O(\log n)$

Calkin, Frieze [RSA'90]



General Graphs

 $O(\log^2 n)$

Blelloch, Fineman, Shun [SPAA'12]

Random Graphs

$$O\left(\frac{\log^2 n}{\log\log n}\right)$$

Coppersmith, Raghavan, Tompa [FOCS'87]

 $O(\log n)$

Calkin, Frieze [RSA'90]



General Graphs



Blelloch, Fineman, Shun [SPAA'12]

 $\Omega\left(\frac{\log n}{\log \log n}\right) \qquad Calkin, Frieze [RSA'90]$

The **Randomized Parallel Greedy MIS** algorithm w.h.p. takes **O**(**log** *n*) rounds on any *n*-node graph.



The **Randomized Parallel Greedy MIS** algorithm w.h.p. takes **O**(**log** *n*) rounds on any *n*-node graph.

This is **tight**: It takes $\Omega(\log n)$ rounds.

The **Randomized Parallel Greedy MIS** algorithm w.h.p. takes **O**(**log** *n*) rounds on any *n*-node graph.

As fast as **Luby's** algorithm, but simpler.

This is **tight**: It takes $\Omega(\log n)$ rounds.

The Randomized Parallel Greedy MIS algorithm w.h.p. takes O(log n) rounds on any n-node graph.
On average, Deterministic Parallel Greedy MIS is fast.
As fast as Luby's algorithm, but simpler.

This is **tight**: It takes $\Omega(\log n)$ rounds.

The **Randomized Parallel Greedy MIS** algorithm w.h.p. takes **O**(**log** *n*) rounds on any *n*-node graph. On **average**, Deterministic Parallel Greedy MIS is fast. As fast as **Luby's** algorithm, but simpler.

This is **tight**: It takes $\Omega(\log n)$ rounds.

Also applies to $(\Delta + 1)$ -Vertex-Coloring, Maximal Matching, and $(2\Delta - 1)$ -Edge-Coloring

Proof Sketch

Proof Sketch

rounds \leq dependency chain length = $O(\log n)$









positions form dependency chain if



positions form dependency chain if

assigned nodes form a (monotonically increasing) path

alternating between MIS nodes and non-MIS nodes

MIS node is inhibitor of next non-MIS node



positions form dependency chain if

assigned nodes form a (monotonically increasing) path

alternating between MIS nodes and non-MIS nodes

MIS node is inhibitor of next non-MIS node

MIS node is inhibitor of non-MIS node if it is its lowest-rank MIS neighbor



positions form dependency chain if

assigned nodes form a (monotonically increasing) path

alternating between MIS nodes and non-MIS nodes

MIS node is inhibitor of next non-MIS node

rounds \leq dependency chain length

MIS node is inhibitor of non-MIS node if it is its lowest-rank MIS neighbor

Proof Sketch

rounds \leq dependency chain length = $O(\log n)$



bound probability that $\Omega(\log n)$ positions form dependency chain



bound probability that $\Omega(\log n)$ positions form dependency chain



bound probability that $\Omega(\log n)$ positions form dependency chain

union bound over all choices of positions

















Probability of Continuing Dependency Chain

Probability of Continuing Dependency Chain

positions form dependency chain if

assigned nodes form a path

alternating **MIS** and non-MIS nodes

MIS is inhibitor of next non-MIS node



analyze probability dependency chain continues through positions r and s
positions form dependency chain if

assigned nodes form a path

alternating **MIS** and non-MIS nodes

MIS is inhibitor of next non-MIS node



analyze probability dependency chain continues through positions r and s

positions form dependency chain if

assigned nodes form a path

alternating **MIS** and non-MIS nodes

MIS is inhibitor of next non-MIS node



positions form dependency chain if

assigned nodes form a path

alternating **MIS** and non-MIS nodes

MIS is inhibitor of next non-MIS node



```
at step step t \leq r,
```

positions form dependency chain if

assigned nodes form a path

alternating **MIS** and non-MIS nodes

MIS is inhibitor of next non-MIS node



```
at step step t \le r,
A) v must be <u>alive</u>
```

positions form dependency chain if

assigned nodes form a path

alternating **MIS** and non-MIS nodes

MIS is inhibitor of next non-MIS node



```
at step step t \leq r,
```

- A) v must be <u>alive</u>
- B) an <u>alive neighbor</u> of v must be at position s

at step $t \leq r$,

A) v must be <u>alive</u>



at step $t \leq r$,

A) v must be <u>alive</u>

B) an <u>alive neighbor</u> of v must be at position s



at step $t \leq r$,

- A) v must be <u>alive</u>
- B) an <u>alive neighbor</u> of v must be at position s





at step $t \leq r$,

- A) v must be <u>alive</u>
- B) an <u>alive neighbor</u> of v must be at position s





at step $t \leq r$,

- A) v must be <u>alive</u>
- B) an <u>alive neighbor</u> of v must be at position s





at step $t \leq r$,

A) v must be <u>alive</u>

B) an <u>alive neighbor</u> of v must be at position s

$$\mathbf{P}[\mathbf{A}] \approx \left(\mathbf{1} - \frac{d_0}{n}\right)$$





at step $t \leq r$,

A) v must be <u>alive</u>

B) an <u>alive neighbor</u> of v must be at position s

$$\mathbf{P}[\mathbf{A}] \approx \left(\mathbf{1} - \frac{d_0}{n}\right)$$





at step $t \leq r$,

A) v must be <u>alive</u>

B) an <u>alive neighbor</u> of v must be at position s

$$\mathbf{P}[\mathbf{A}] \approx \left(\mathbf{1} - \frac{d_0}{n}\right)$$





at step $t \leq r$,

A) v must be <u>alive</u>

B) an <u>alive neighbor</u> of v must be at position s

$$\mathbf{P}[\mathbf{A}] \approx \left(\mathbf{1} - \frac{d_0}{n}\right) \cdot \left(\mathbf{1} - \frac{d_1}{n}\right)$$





at step $t \leq r$,

A) v must be <u>alive</u>

$$\mathbf{P}[\mathbf{A}] \approx \left(\mathbf{1} - \frac{d_0}{n}\right) \cdot \left(\mathbf{1} - \frac{d_1}{n}\right) \cdot \left(\mathbf{1} - \frac{d_2}{n}\right)$$







at step $t \leq r$,

A) v must be <u>alive</u>

$$\mathbf{P}[\mathbf{A}] \approx \left(\mathbf{1} - \frac{d_0}{n}\right) \cdot \left(\mathbf{1} - \frac{d_1}{n}\right) \cdot \left(\mathbf{1} - \frac{d_2}{n}\right)$$







at step $t \leq r$,

A) v must be <u>alive</u>

$$\mathbf{P}[\mathbf{A}] \approx \left(\mathbf{1} - \frac{d_0}{n}\right) \cdot \cdots \cdot \left(\mathbf{1} - \frac{d_{t-1}}{n}\right)$$



degree sequence of v:
$$d_0, d_1, \dots, d_t$$



at step $t \leq r$,

A) v must be <u>alive</u>

$$P[A] \approx \left(1 - \frac{d_0}{n}\right) \cdots \left(1 - \frac{d_{t-1}}{n}\right)$$
$$P[B] \approx \frac{d_t}{n}$$





at step $t \leq r$, v must be alive A) B)

an alive neighbor of v must be at position s

P[continue]
$$\approx \prod_{i=0}^{t-1} \left(1 - \frac{d_i}{n}\right) \cdot \frac{d_t}{n}$$





at step $t \leq r$, v must be alive A) B)

an alive neighbor of v must be at position s

P[continue]
$$\approx \prod_{i=0}^{t-1} \left(1 - \frac{d_i}{n}\right) \cdot \frac{d_t}{n}$$





at step $t \leq r$, v must be alive A) B)

an alive neighbor of v must be at position s

P[continue]
$$\approx \prod_{i=0}^{t-1} \left(1 - \frac{d_i}{n}\right) \cdot \frac{d_t}{n}$$





V

at step $t \leq r$,A)v must be aliveB)an alive neighbor of v must be at position s

P[continue]
$$\approx \prod_{i=0}^{t-1} \left(1 - \frac{d_i}{n}\right) \cdot \frac{d_t}{n}$$

degree sequence of v: d_0, d_1, \dots, d_t

V

small enough for any degree sequence, even if we only look at (any) t positions



rounds \leq dependency chain length = $O(\log n)$

rounds \leq dependency chain length = $O(\log n)$

bound probability that $\Omega(\log n)$ positions form dependency chain

rounds \leq dependency chain length = $O(\log n)$

bound probability that $\Omega(\log n)$ positions form dependency chain



rounds \leq dependency chain length = $O(\log n)$

bound probability that $\Omega(\log n)$ positions form dependency chain



rounds \leq dependency chain length = $O(\log n)$

bound probability that $\Omega(\log n)$ positions form dependency chain



reserve disjoint set of positions for each position

rounds \leq dependency chain length = $O(\log n)$

bound probability that $\Omega(\log n)$ positions form dependency chain



reserve disjoint set of positions for each position

rounds \leq dependency chain length = $O(\log n)$

bound probability that $\Omega(\log n)$ positions form dependency chain



reserve disjoint set of positions for each position

use these positions to bound continuation probability of position

rounds \leq dependency chain length = $O(\log n)$

bound probability that $\Omega(\log n)$ positions form dependency chain



reserve disjoint set of positions for each position

use these positions to bound continuation probability of position

rounds \leq dependency chain length = $O(\log n)$

bound probability that $\Omega(\log n)$ positions form dependency chain



reserve disjoint set of positions for each position

use these positions to bound continuation probability of position

rounds \leq dependency chain length = $O(\log n)$

bound probability that $\Omega(\log n)$ positions form dependency chain



reserve <u>disjoint set of positions</u> for each position

use these positions to bound continuation probability of position

rounds \leq dependency chain length $= O(\log n)$

bound probability that $\Omega(\log n)$ positions form dependency chain



reserve disjoint set of positions for each position

use these positions to bound continuation probability of position

rounds \leq dependency chain length = $O(\log n)$

bound probability that $\Omega(\log n)$ positions form dependency chain



reserve disjoint set of positions for each position

use these positions to bound continuation probability of position

rounds \leq dependency chain length = $O(\log n)$

bound probability that $\Omega(\log n)$ positions form dependency chain



reserve <u>disjoint set of positions</u> for each position

use these positions to bound continuation probability of position
Recap of Proof

rounds \leq dependency chain length = $O(\log n)$

bound probability that $\Omega(\log n)$ positions form dependency chain



reserve disjoint set of positions for each position

use these positions to bound continuation probability of position

union bound over all choices of positions

Recap of Proof

rounds \leq dependency chain length = $O(\log n)$

bound probability that $\Omega(\log n)$ positions form dependency chain



reserve disjoint set of positions for each position

use these positions to bound continuation probability of position

probability of dependency chain is product of continuation probabilities

union bound over all choices of positions

The **Parallel Randomized Greedy MIS** algorithm w.h.p. takes $\Theta(\log n)$ rounds on *n*-node graphs.

The **Parallel Randomized Greedy MIS** algorithm w.h.p. takes $\Theta(\log n)$ rounds on *n*-node graphs.

As fast as Luby's algorithm, but simpler.

The **Parallel Randomized Greedy MIS** algorithm w.h.p. takes $\Theta(\log n)$ rounds on *n*-node graphs.

As fast as Luby's algorithm, but simpler.

On average, Deterministic Parallel Greedy MIS algorithm is fast.

The **Parallel Randomized Greedy MIS** algorithm w.h.p. takes $\Theta(\log n)$ rounds on *n*-node graphs.

As fast as Luby's algorithm, but simpler.

On average, Deterministic Parallel Greedy MIS algorithm is fast.

Which algorithms can be replaced by simpler ones, on average?

The **Parallel Randomized Greedy MIS** algorithm w.h.p. takes $\Theta(\log n)$ rounds on *n*-node graphs.

As fast as Luby's algorithm, but simpler.

On average, Deterministic Parallel Greedy MIS algorithm is fast.

Which algorithms can be replaced by simpler ones, on average?





