

Distributed Algorithms 2025

11

Hardness of coloring

This week's goals

- **Specific technical result:**

- 3-coloring of cycles in the LOCAL model
- possible in $O(\log^* n)$ rounds (week 1)
- not possible in $o(\log^* n)$ rounds (this week)

- **General idea:**

- how to use round elimination to prove negative results in the **LOCAL** model and/or for **randomized** algorithms

Challenge & workaround

- Round elimination does not work directly in the **LOCAL** model
 - problem: **independence** vs. unique identifiers
- But we can use it to study *randomized* algorithms in the **PN** model
 - random bits are independent!
- Then results for the LOCAL model follow!

General idea:

**Randomized
round elimination**

Randomized round elimination

- **The same pair of problems:** X and $\text{re}(X)$
 - $\text{re}(X)$ does not depend on model of computing!
- **Different implications** in different models:
 - *if A is a deterministic PN-algorithm that solves X in T rounds then ...*
 - *if A is a randomized PN-algorithm that solves X in T rounds with high probability then ...*

Randomized round elimination

- A_0 : solves X , local failure probability $< p_0$
- A_1 : solves $\text{re}(X)$, local failure probability $< p_1$
- Failure probability increases $\approx$ polynomially
 - we can repeat this $\approx \mathbf{\log \log n}$ times before A_k becomes useless
 - in particular, repeating this $\approx \mathbf{\log^* n}$ times is fine!

Case study: **Coloring directed cycles**

A a
B b
C c

A bc
B ac
C ab

F G
E H
D I
C J
B K
A L

ACE GIK
BCF HIL
DEF JKL

A a
B b
C c
D d
E e
F f

ACE fdb
BCF eda
DEF cba

A a
B b
C c
D d
E e
F f

A bdf
B ade
C abdef
D abc
E abcdf
F abcde

A a
B b
C c
D d
E e
F f

A bcdef
B acdef
C abdef
D abcef
E abcdf
F abcde

$n^{O(1)}$ -coloring

0 rounds
(unique IDs)



$O(\log^* n)$ steps

3-coloring

$O(\log^* n)$
rounds

2^k -coloring

$T - 1$ rounds



$2k$ -coloring

T rounds

*Fast color
reduction*

2^k -coloring



k -coloring

*Round
elimination*

c -coloring

0 rounds



T steps

3-coloring

$T \ll \log^* n$
rounds

Proof:

**Randomized
round elimination**

Randomized round elimination

- We will use cycles as an example
- The same idea generalizes to biregular trees
 - probabilities that we get are just slightly different

Randomized round elimination in cycles

- **A_0** : local failure probability $< 1/x^3$  e.g. 0.1%
- **A_1** : form the set of *frequent labels*
 - labels that appear with probability $\geq 1/x$  e.g. 10%
- **Analysis**: focus on *lucky neighborhoods*
 - neighborhoods in which A_0 fails with probability $< 1/x^2$  e.g. 1%

Intuition

- **Before seeing anything:**
 - we know that A_0 failure rate is $< 1/x^3$
- **Gather more local information:**
 - gain more information on A_0 failure rate here
 - *"unlucky"*: much worse than average failure rate
 - ideally these are very rare
 - *"lucky"*: not much worse than average failure rate
 - ideally failure rate is still very low

New active nodes

- Assume we are in a *lucky neighborhood*
 - by definition: $P[A_0 \text{ fails}] < 1/x^2$
- Assume $[a, b]$ is a pair of *frequent labels*
 - happens here with probability $\geq 1/x \cdot 1/x = 1/x^2$
 - A_0 cannot fail here with probability $\geq 1/x^2$
 - **label pair $[a, b]$ must be feasible!**
- A_1 can fail only in unlucky neighborhoods!

Lucky neighborhoods

- **Assumption:** $P[A_0 \text{ fails}] < 1/x^3$  e.g. 0.1%
- **Definition:** $P[A_0 \text{ fails} \mid \text{unlucky}] \geq 1/x^2$  e.g. 1%
- $P[A_0 \text{ fails} \mid \text{unlucky}] \cdot P[\text{unlucky}] < 1/x^3$
- $P[\text{unlucky}] < \mathbf{1/x}$  e.g. 10%

New passive nodes

- $P[A_0 \text{ fails}] < 1/x^3$
- $P[A_0 \text{ output considered infrequent by } A_1] < \#labels \cdot \#edges \cdot 1/x$
- **Otherwise:**
 - A_0 does not fail, its outputs form a valid solution
 - A_0 outputs only labels that A_1 considers frequent
 - A_1 has to succeed in solving $re(X)$

Summary

- $P[A_0 \text{ fails}] < \mathbf{1/x^3}$
- Possible A_1 -failures:
 - $P[\text{unlucky}] < \mathbf{1/x}$
 - $P[A_0 \text{ fails}] < \mathbf{1/x^3}$
 - $P[A_0 \text{ outputs some infrequent label}] < \mathbf{\#labels \cdot \#edges \cdot 1/x}$
- $P[A_1 \text{ fails}] < \mathbf{\text{constant} \cdot 1/x}$

What works very often

- Do round elimination in deterministic PN model
 - gain intuition on how the problem behaves
- Then switch to randomized PN model
 - proper analysis of failure probabilities
- Results for deterministic & randomized LOCAL follow directly

Sinkless orientation

- **Deterministic PN:**

- not possible in $o(\log n)$ rounds (last week)
- possible in $O(\log n)$ rounds (last week)

- **Randomized PN:**

- not possible in $o(\log \log n)$ rounds (exercise)
- possible in $O(\log \log n)$ rounds (not easy)

- **Deterministic LOCAL?**