Google

Breaking Barriers and Closing Gaps for MIS and MM by Understanding Vertex Survival Probability

Seri Khoury, **Aaron Schild**
INSAIT, Google Research

Overview

Distributed Computing: The LOCAL Model

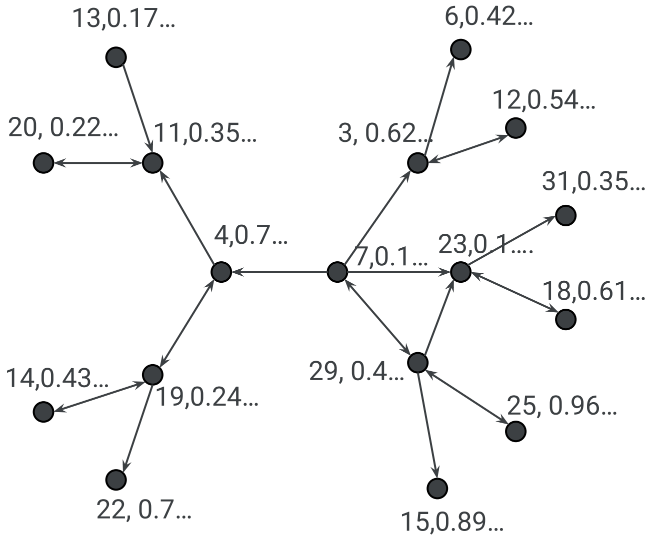

arrows: direction of message flow

- (Randomized) LOCAL model
  - n vertex graph
  - Every vertex given two things
    - O(log n) bit ID
      - Unique
    - Infinite tape of randomness
      - Tapes are independent [0,1]
  - Synchronous rounds
  - Vertex sends message to each neighbor
  - Terminate after T rounds and output solution
  - **Question**: How many rounds required?
- Sidenote: (Randomized) CONGEST model
  - Same as LOCAL, but messages are O(log n) bits

Google

Today's Problems: MM and MIS

Maximal Matching (MM)

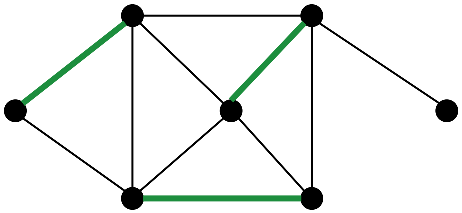

≤

(take line graph)

Maximal Independent Set (MIS)

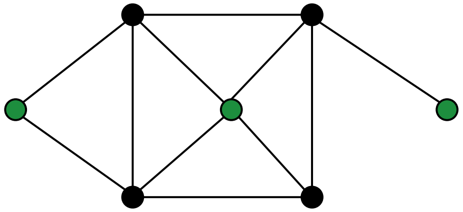

- Given: undirected graph
- Matching
  - Edges with no shared endpoints
- Maximal
  - every vertex matched/no addable edge
- Goal: find matching that is maximal

- Given: undirected graph
- Independent set (IS)
  - Nonadjacent vertices
- Maximal
  - Every vertex in or adjacent to IS
- Goal: find maximal independent set

MIS

MIS: Pre-2025 Upper and Lower Bounds

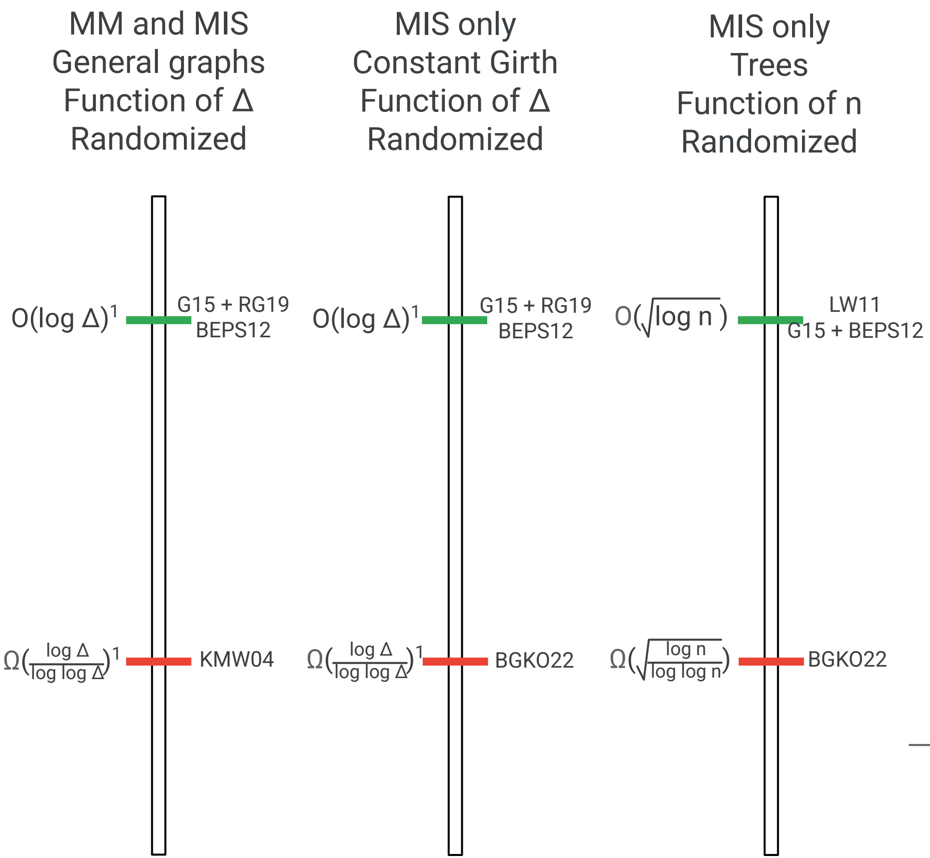

Green tick: algorithm. Red tick: lower bound.

[1]Functions of other params left out for simplicity.

n: number of vertices, Δ: max degree

- $O(\log n)$ randomized algorithm (Luby86, ABI86)
  - No improvement as function of n in 40 years
- $O(\log \Delta + \text{polyloglog } n)$ randomized (G15 + RG19)
  - Followed $O(\log^2 \Delta + \text{polyloglog } n)$ (BEPS12)
- $O(\sqrt{\log n})$ randomized (trees only) (LW11, G15 + BEPS12)
- $O(\log \alpha + \sqrt{\log n})$ randomized (G15 + RG19)
  - α is arboricity
  - g girth ➡ $n^{1/g}$ arboricity
    - Superconstant girth ➡ $o(\log n)$ rounds
    - What about constant girth?
- $O(\log^{5/3} n)$ deterministic (GG24)

- $\Omega(\min(\frac{\log \Delta}{\log \log \Delta}, \sqrt{\frac{\log n}{\log \log n}}))$ randomized (KMW04)
  - Also holds for MIS on trees (BGKO22)

Google

MIS: 2025 Upper and Lower Bounds

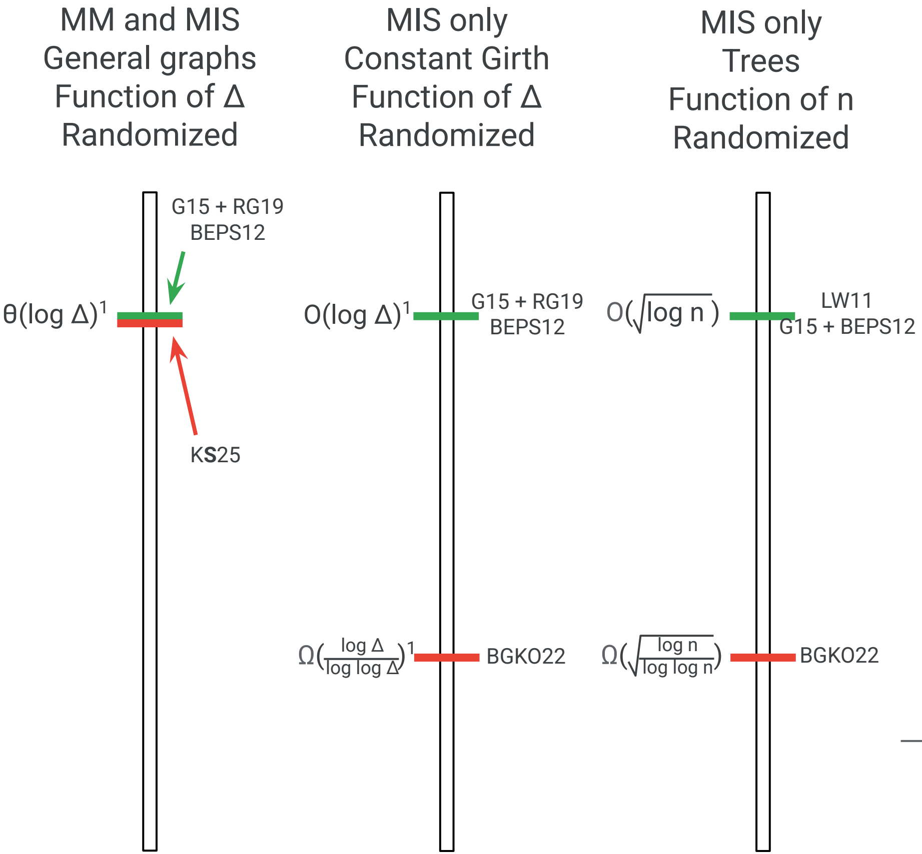

Green tick: algorithm. Red tick: lower bound.
[1]Functions of other params left out for simplicity.

n: number of vertices, Δ: max degree

- O(log n) randomized algorithm (Luby86, ABI86)
  - No improvement as function of n in 40 years
- O(log Δ + polyloglog n) randomized (G15 + RG19)
  - Followed $O(\log^2 \Delta + \text{polyloglog } n)$ (BEPS12)
- $O(\sqrt{\log n})$ randomized (trees only) (LW11, G15 + BEPS12)
- O(log α + polyloglog n) randomized (G15 + RG19)
  - α is arboricity
  - g girth ➡ $n^{1/g}$ arboricity
    - Superconstant girth ➡ o(log n) rounds
    - What about constant girth?
- $O(\log^{5/3} n)$ deterministic (GG24)

- $\Omega(\min(\frac{\log \Delta}{\log\log \Delta}, \sqrt{\frac{\log n}{\log\log n}}))$ randomized (KMW04)
  - Also holds for MIS on trees (BGKO22)
- $\Omega(\min(\log \Delta, \sqrt{\log n}))$ randomized (K**S**25)

Our Results (Randomized CONGEST Model)

1. $O(\frac{\log \Delta}{\log (\log^* \Delta)})$ round MIS on girth 7 graphs
2. $O(\sqrt{\frac{\log n}{\log (\log^* n)}})$ round MIS on trees

- These results do not help for MM
  - Line graph has short cycles
  - Fundamental by K**S**25 lower bound
- Girth result beats Luby for **constant** girth, not just superconstant girth
- Tree result **breaks conjecture** in book by Barenboim and Elkin

Implications of Our Results

Results: MIS on (1) girth ≥θ(1) in $O(\frac{\log \Delta}{\log \log^* \Delta})$ rounds (2) trees in $O(\sqrt{\frac{\log n}{\log \log^* n}})$ rounds

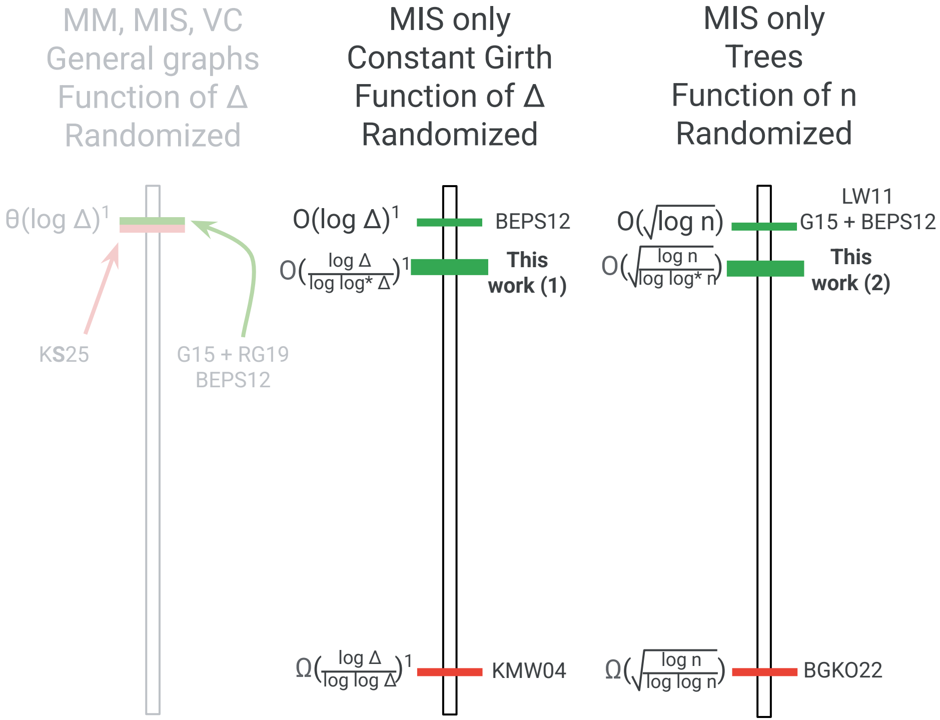

Green tick: algorithm. Red tick: lower bound.
[1]Functions of other params left out for simplicity.

1. Breaking Barriers
   a. (1) breaks O(log n) barrier for constant-girth graphs
      i. Only known for superconstant girth previously
   b. (2) breaks conjecture 11.15 in BE book
2. Separations
   a. MM strictly harder than MIS on trees
      i. Opposite is true for general graphs!
   b. MIS on trees easier than MIS on general graphs
      i. Trees strictly easier than general graphs!
3. Exciting Open Problem
   a. We don't achieve optimality
   b. Question: What's the right complexity?
      i. Has to be a weird number!

The Key Question

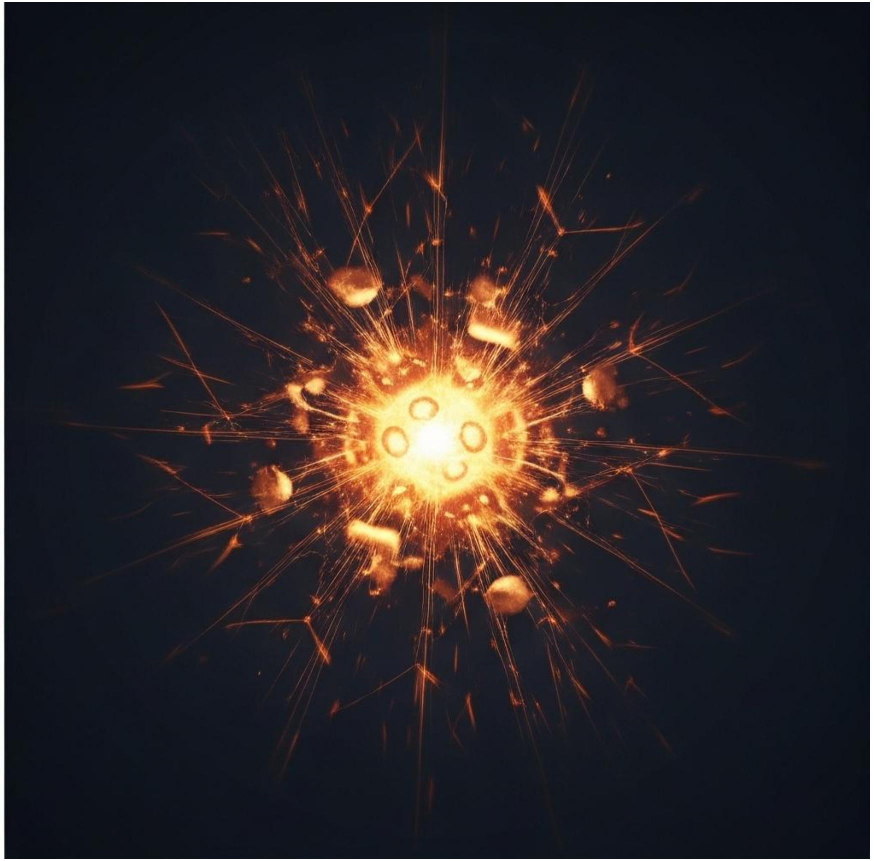

Google

A Tale of Two Probabilities

Inclusion Probability

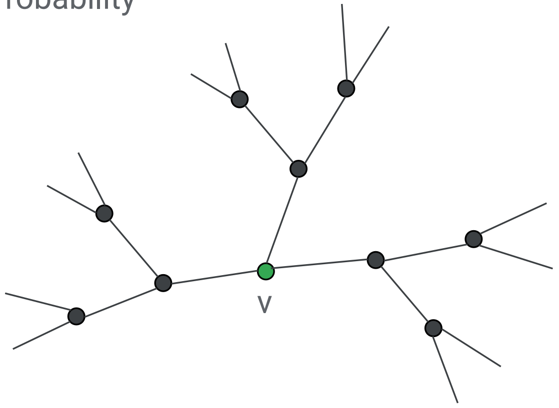

Survival Probability

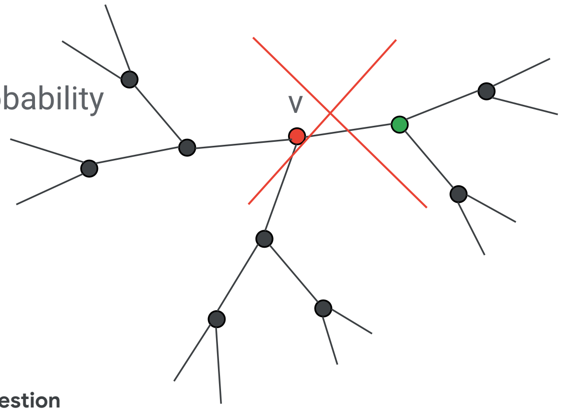

- Will think about 1 and 2-round algorithms
  - Ex: Luby repeats 1-round algo O(log n) times
- Simplifications
  - Girth Ω(1) + 2-round ➡ can focus on trees
  - G15 algorithm ➡ can focus on regular graphs
  - Independent set (IS) only, not MIS
    - Repetition yields **MIS**
  - Focus on **Δ-regular trees**
    - All vertices behave the same
- Pick arbitrary vertex v
- **Inclusion probability**
  - Probability that v is included in IS
- **Survival probability**
  - Probability that v remains after picking IS

The Implications of High Inclusion Probability

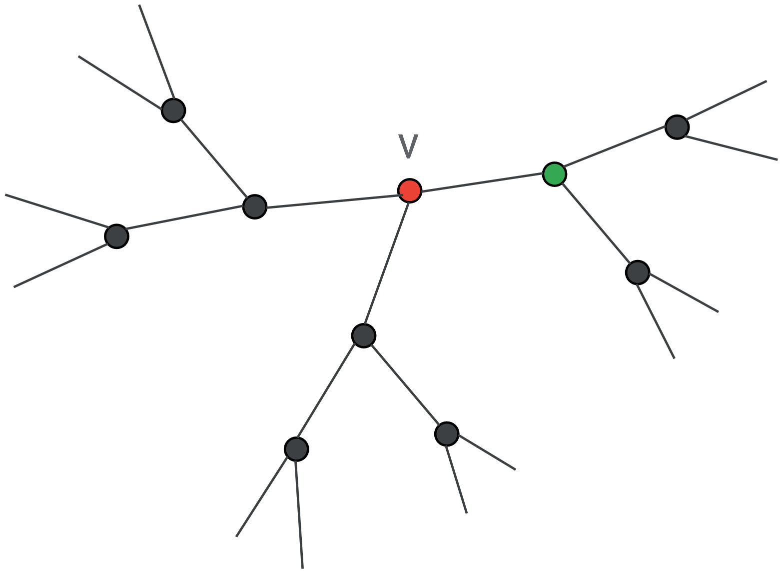

- Focus on **Δ-regular trees**
- Pick arbitrary vertex v
- **Inclusion probability (IP)**
  - Probability that v is included in IS
- **Survival probability (SP)**
  - Probability that v remains after picking IS
- High IP ➡ Low SP
  - (Roughly) IP = (k/Δ) ➡ SP = $(1 - k/\Delta)^{\Delta} \approx e^{-k}$
  - Tree structure yields independence
- Low SP ➡ fast algorithm
  - Roughly $(\log \Delta)/k$ rounds to shatter (regular)
  - $(\log \Delta)/(\log k)$ rounds for general graphs (G15)
- Goal: Find **O(1)**-round algo **with high IP**

The Key Subproblem

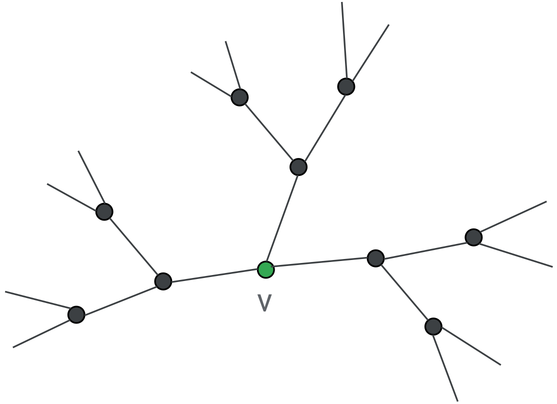

- Focus on **Δ-regular trees**
- Pick arbitrary vertex v
- **Inclusion probability (IP)**
  - Probability that v is included in IS
- High IP ➡ fast algorithm
- Goal: Find **O(1)**-round algo **with high IP**

What is the maximum inclusion probability (IP) of any r-round IS algo?

Google

Warmup: 1 Round

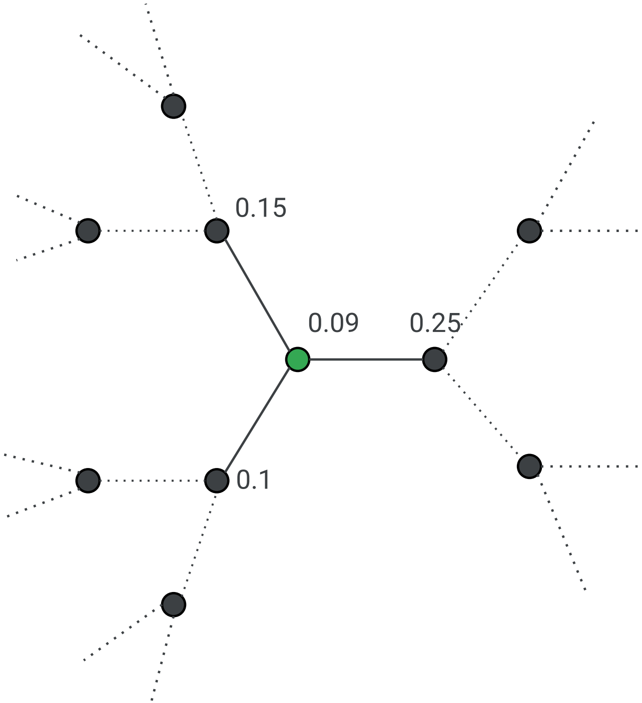

- Focus on **Δ-regular trees**
- Pick arbitrary vertex v
- **Inclusion probability (IP)**
    - Probability that v is included in IS
- Max IP of any 1-round algorithm for IS?

Google

Warmup: 1 Round

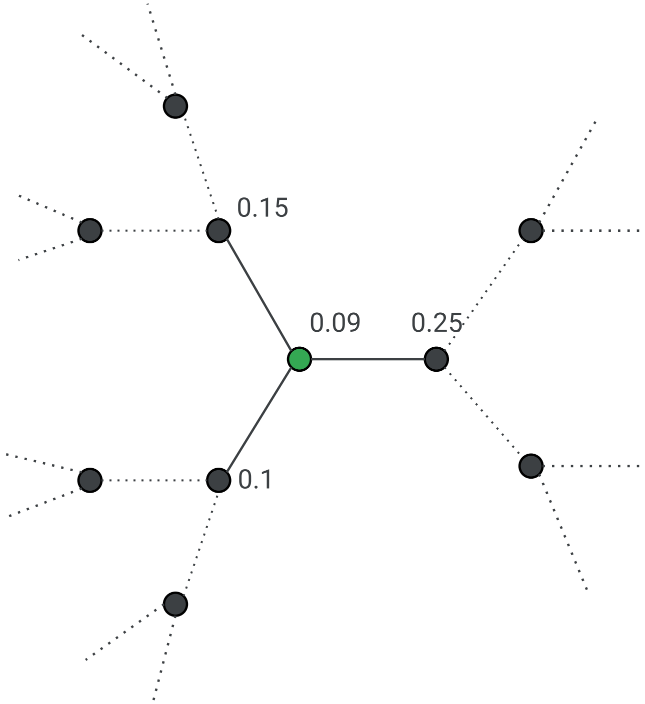

- Focus on **Δ-regular trees**
- Pick arbitrary vertex v
- **Inclusion probability (IP)**
  - Probability that v is included in IS
- Max IP of any 1-round algorithm for IS?
- Answer: $\frac{1}{\Delta + 1}$
- Algorithm: Ranking (Luby86, ABI86)
  - Each vertex picks random number
  - Joins IS if smaller than neighbors
  - v has minimum rank with desired probability

Google

1 Round Hardness

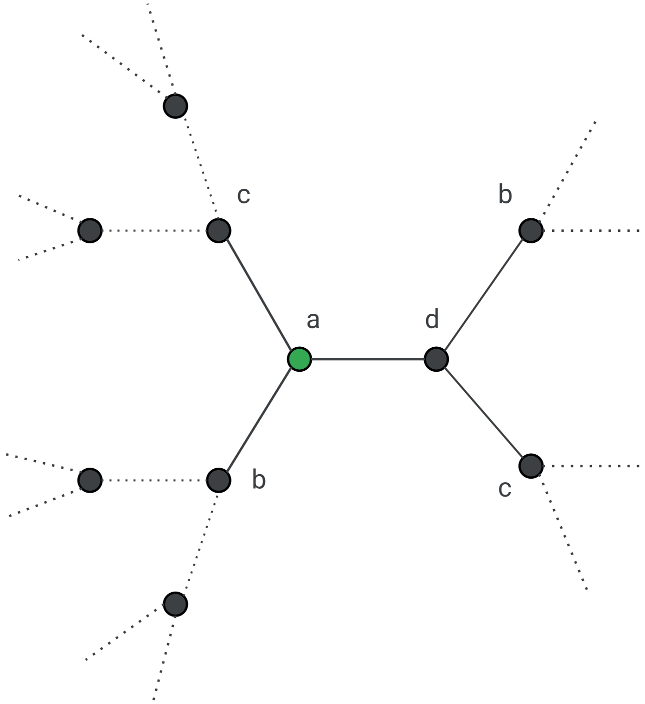

- Focus on **Δ-regular trees**
- Pick arbitrary vertex v
- **Inclusion probability (IP)**
  - Probability that v is included in IS
- Max IP of any 1-round algorithm for IS?
- Answer: $\frac{1}{\Delta + 1}$
- Algorithm: Ranking (Luby86, ABI86)
  - v has minimum rank with desired probability
- Hardness?
  - Yes, nothing higher than $\frac{1}{\Delta + 1}$
  - For any set of $\Delta + 1$ numbers, one center can accept
    - Any 2 such 1-neighborhoods can be adjacent
    - $(\Delta + 1)$-tuples are mutually exclusive

Google

What About 2 Rounds?

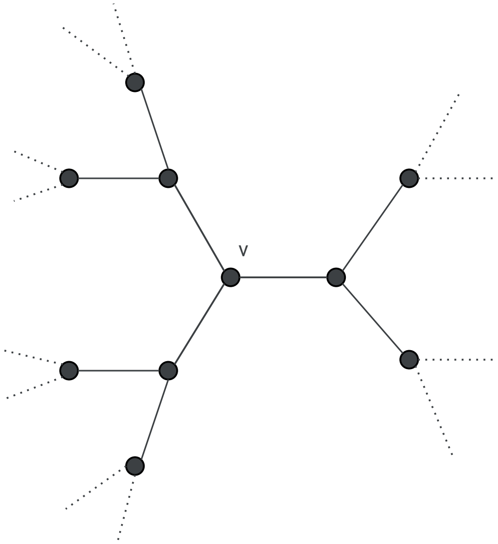

- Focus on **Δ-regular trees**
- Pick arbitrary vertex v
- **Inclusion probability (IP)**
  - Probability that v is included in IS
- Max IP of any 2-round algorithm for IS?
- One might guess roughly $\frac{2}{\Delta}$, but this is wrong!
- 2 rounds ends up being much better than 1 round 2 times

Key Lemma: There is a **2-round** algorithm for IS with inclusion probability $\boldsymbol{\Omega(\frac{\log^* \Delta}{\Delta})}$.

Sketching the Rest

Key Lemma: There is a **2-round** algorithm for IS with inclusion probability $\mathbf{\Omega(\frac{\log^* \Delta}{\Delta})}$ on **Δ-ary trees**.

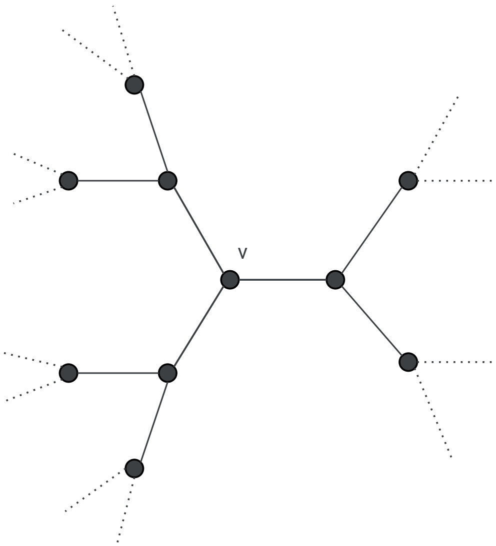

- Prove generalization to irregular graphs, following G15
  - Vertex v has "desire level" p(v)
    - Changes over rounds
    - v joins with prob p(v) if low effective degree
    - We increase this to $(\log^* \Delta)^{1-\varepsilon}$ p(v)
  - G15 analysis ➡ $O(\frac{\log \Delta}{\log \log^* \Delta}$ + polyloglog n) rounds
    - With some minor changes

Google

2 Rounds in Δ-ary trees

Marking 1-Round Algorithm (Luby86, ABI86)

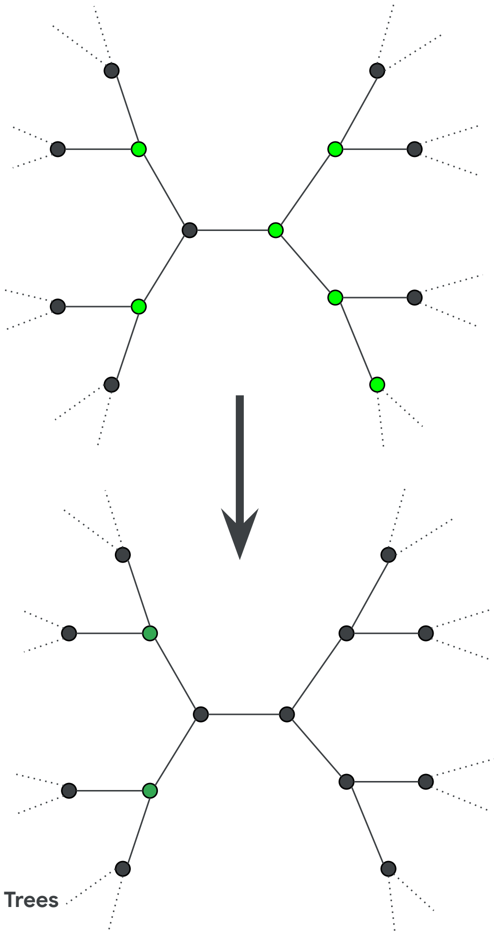

- Focus on **Δ-regular trees**
- **Inclusion probability (IP)**
  - Probability that arbitrary v is included in IS
- Mark v with probability $p = \frac{1}{\Delta}$ (or sample [0,1], accept < $\frac{1}{\Delta}$)
- Include a vertex v iff both
  - it is marked
  - no neighbor is
- Always produces independent set (IS)
- $IP = p\,(1 - p)^{\Delta} = \theta(\frac{1}{\Delta})$
- 1-round because v only needs to know neighbor marks

Google

Key Concept: i-candidates

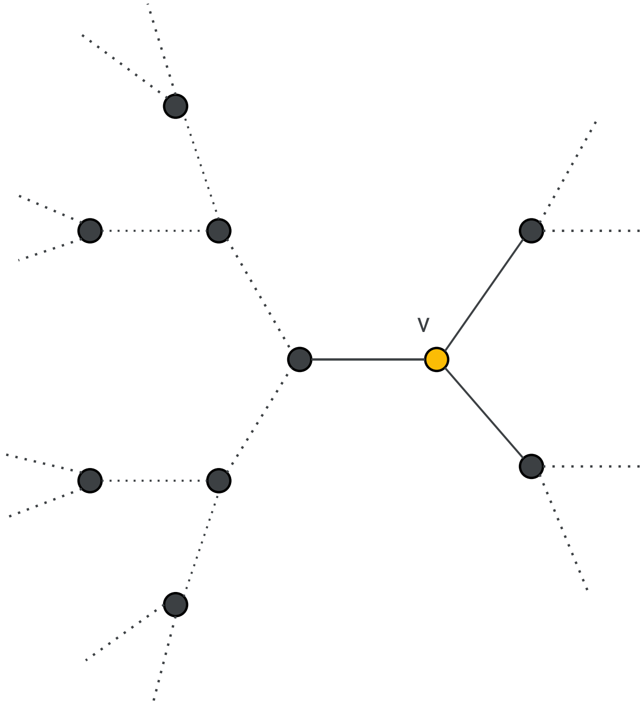

- Focus on **Δ-regular trees**
- **Inclusion probability (IP)**
  - Probability that arbitrary v is included in IS
- Sample r_v uniformly, independently from [0,1]
- ~~Include a vertex~~ v is an *i-candidate* iff both
  - ~~it is marked~~ $\frac{9a_{i+1}}{10\Delta} < r_v < \frac{a_{i+1}}{\Delta}$
  - ~~no neighbor is~~ $r_u > \frac{a_i}{\Delta}$ for any neighbor u of v
  - Intuition: v is candidate if no neighbor is *much* lower
    - Tiebreaking only necessary if candidate
- a_is: rapidly increasing sequence (set next slide)
- Candidacy determinable in 1-round
- Key Fact: different types of candidates not adjacent!

Candidacy Probability for i-candidates

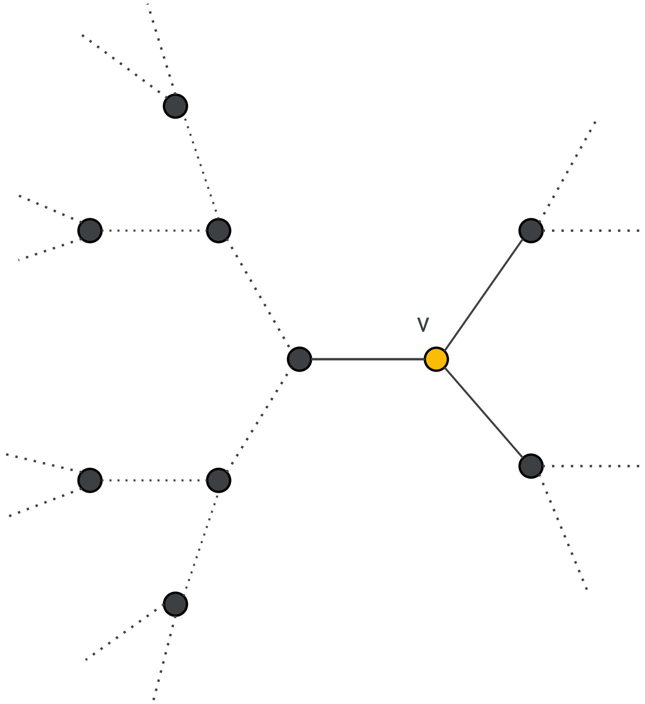

- Focus on **Δ-regular trees**
- **i-Candidacy Probability (i-CP)**
  - Probability that arbitrary v is an i-candidate
- Sample r_v uniformly, independently from [0,1]
- v is an *i-candidate* iff both
  - $\frac{9a_{i+1}}{10\Delta} < r_v < \frac{a_{i+1}}{\Delta}$
  - $r_u > \frac{a_i}{\Delta}$ for any neighbor u of v
- $\text{i-CP} = \left(\frac{a_{i+1}}{\Delta} - \frac{9a_{i+1}}{10\Delta}\right)\left(1 - \frac{a_i}{\Delta}\right)^{\Delta} \approx \frac{a_{i+1}}{10\Delta} e^{-a_i} = \mathbf{\frac{1}{10\Delta}}$
  - Set $a_{i+1} = e^{a_i}$
    - Rapidly increasing indeed
    - This is where log* Δ comes from

Google

Inclusion Probability for i-candidates

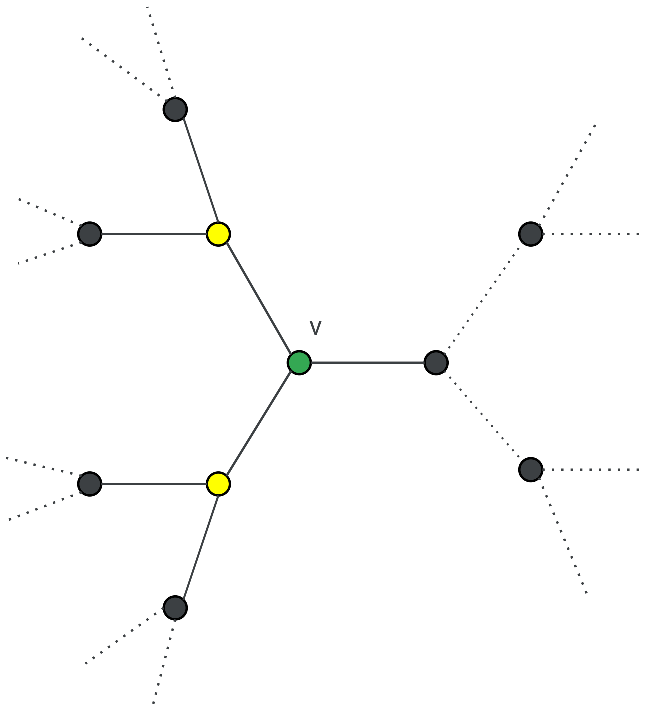

- Focus on **Δ-regular trees**
- **i-Inclusion Probability (i-IP)**
  - Probability that arbitrary v is an i-candidate **and** in IS
- v is an *i-candidate* iff both
  - $\frac{9a_{i+1}}{10\Delta} < r_v < \frac{a_{i+1}}{\Delta}$
  - $r_u > \frac{a_i}{\Delta}$ for any neighbor u of v
- i-CP ≈ $\frac{1}{10\Delta}$
- Final Algorithm: v in IS iff
  - v is a candidate
  - no neighbor is
- Pr[in IS | i-candidate] = Pr[no neighbor is candidate | i-can]
  - = Pr[no neighbor is i-can | i-can]
  - Union bound ➡ $\geq 1 - \Delta(\frac{1}{10\Delta}) > ½$
    - Actual argument deals with dependencies
- ➡ i-IP = $\Omega(\frac{1}{\Delta})$

Overall Inclusion Probability

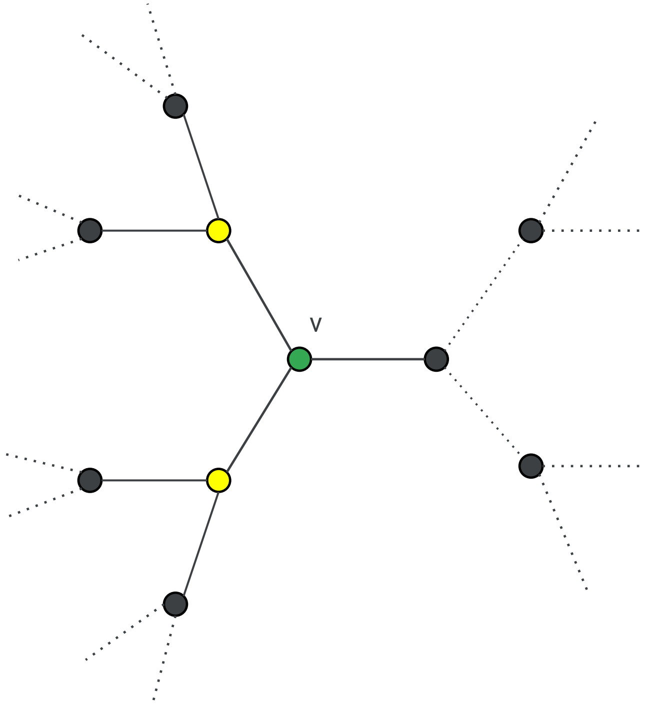

- Focus on **Δ-regular trees**
- **Inclusion Probability (IP)**
  - Probability that arbitrary v is in IS
  - i-IP : v in IS **and** v is i-candidate
- v is an *i-candidate* iff both
  - $\frac{9a_{i+1}}{10\Delta} < r_v < \frac{a_{i+1}}{\Delta}$
  - $r_u > \frac{a_i}{\Delta}$ for any neighbor u of v
- Final Algorithm: v in IS iff
  - v is a candidate
  - no neighbor is
- i-IP = $\Omega(\frac{1}{\Delta})$ for any i
- v at most one type of candidate ➡ IP = $\Omega(\frac{\log^* \Delta}{\Delta})$ ✅
  - Since there are > log* Δ buckets
- 2-rounds

Google

2-Round Summary

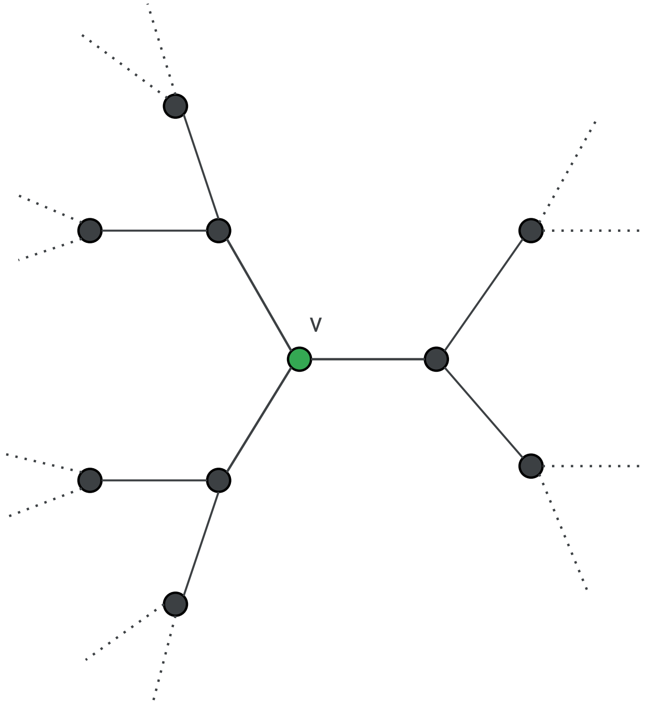

- Focus on **Δ-regular trees**
- **Inclusion Probability (IP)**
  - Probability that arbitrary v is in IS
- Found 2-round algo for IS with inclusion probability $\Omega(\frac{\log^* \Delta}{\Delta})$
- Starting point: marking version of Luby86, ABI86
- Introduce asymmetry:
  - Mark center with higher chance than neighbors
  - Not able to add to IS immediately, but is a candidate
  - Tiebreak using 2-neighborhood
- Higher Inclusion Probability ➡ **faster degree reduction**

Google

MIS Summary

Results: MIS on (1) girth ≥θ(1) in $O(\frac{\log \Delta}{\log \log^* \Delta})$ rounds (2) trees in $O(\sqrt{\frac{\log n}{\log \log^* n}})$ rounds

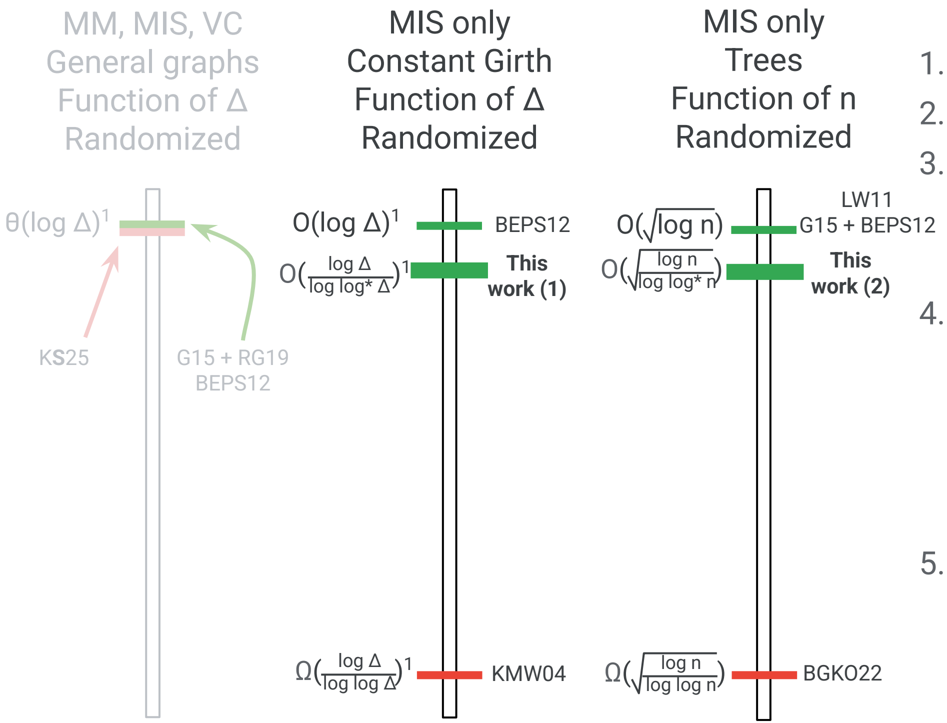

Green tick: algorithm. Red tick: lower bound.
[1]Functions of other params left out for simplicity.

1. Breaking some barriers for MIS
2. Establishing some new separations for MIS and MM
3. Techniques:
   a. New algo followed from surprising 2-round marking algo
   b. Plugs into G15 with minor modifications
4. 2-round marking algo:
   a. Instead of marking, then selecting IS nodes,
   b. Mark, then select candidates, then select IS nodes
   c. Candidate iff no neighbor is much lower (clear winner)
   d. Tiebreak only amongst candidates of same type
5. Open Questions
   a. Beat O(log n) for general graphs?
   b. Improve log* dependence to match hardness?
      i. Or is log* dependence optimal??

MM

MM and MIS: Algorithms and Lower Bounds

MM and MIS	MM only	MIS only
General graphs	Trees	Trees
Function of Δ	Function of n	Function of n
Randomized	Randomized	Randomized

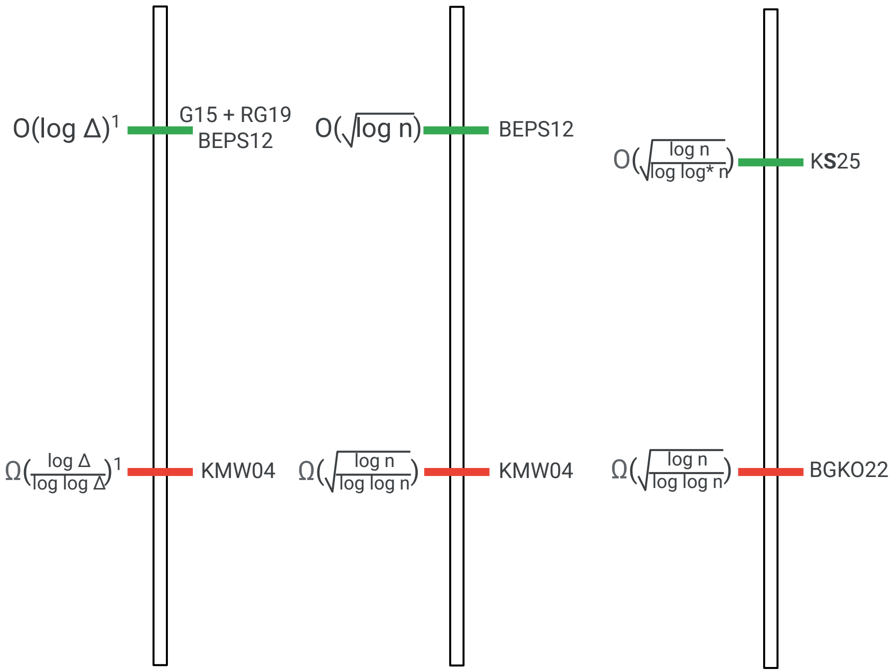

Green tick: algorithm. Red tick: lower bound.
[1]Functions of other params left out for simplicity.

n: number of vertices, Δ: max degree

- $O(\log n)$ randomized algorithm (Luby86, ABI86)
  - No improvement as function of n in 40 years
- $O(\log \Delta + \text{polyloglog } n)$ randomized (G15 + RG19)
  - Followed $O(\log^2 \Delta + \text{polyloglog } n)$ (BEPS12)
- $O(\sqrt{\log n})$ randomized (trees only) (LW11, G15 + BEPS12)
- $O(\sqrt{\frac{\log n}{\log \log^* n}})$ randomized (MIS on trees only) (K**S**25)
- $O(\log^{5/3} n)$ deterministic (GG24)

- $\Omega(\min(\frac{\log \Delta}{\log \log \Delta}, \sqrt{\frac{\log n}{\log \log n}}))$ randomized (KMW04)
  - Also holds for MIS on trees (BGKO22)
  - $\Omega(\min(\log \Delta, \sqrt{\log n}))$ claimed (KMW10)
- $\Omega(\min(\Delta, \frac{\log \log n}{\log \log \log n}))$ randomized (BBHORS19)
- $\Omega(\min(\Delta, \frac{\log n}{\log \log n}))$ deterministic (BBHORS19)

Lower Bounds: The KMW Bound (KMW04, CL20)

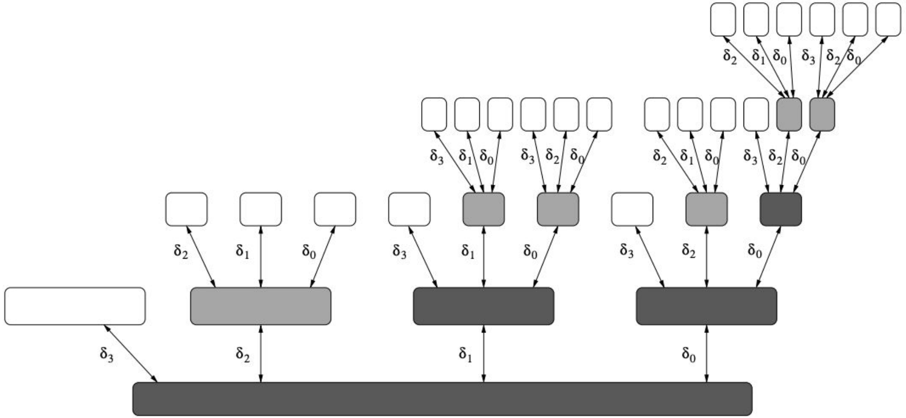

Image credit: KMW04

- $\Omega(\min(\frac{\log \Delta}{\log\log \Delta}, \sqrt{\frac{\log n}{\log\log n}}))$ randomized (KMW04)
  - Applies against O(1)-apx vertex cover (VC)
    - Also MM and MIS via reductions
- Technique: Indistinguishability
  - Let $r = \Omega(\min(\frac{\log \Delta}{\log\log \Delta}, \sqrt{\frac{\log n}{\log\log n}}))$
  - Create graph with two sets of vertices X, Y
    - X, Y have identical r-neighborhoods
      - Local property
    - Identical behavior causes large VC
      - Global property => contradiction
- Takeaways
  - Shows hardness for MM on **irregular trees**
  - Not improvable
    - O(1)-apx VC in $O(\frac{\log \Delta}{\log\log \Delta})$ time (BCS16)

Lower Bounds: Round Elimination (L92, BFHKLRSU15, BBHORS19)

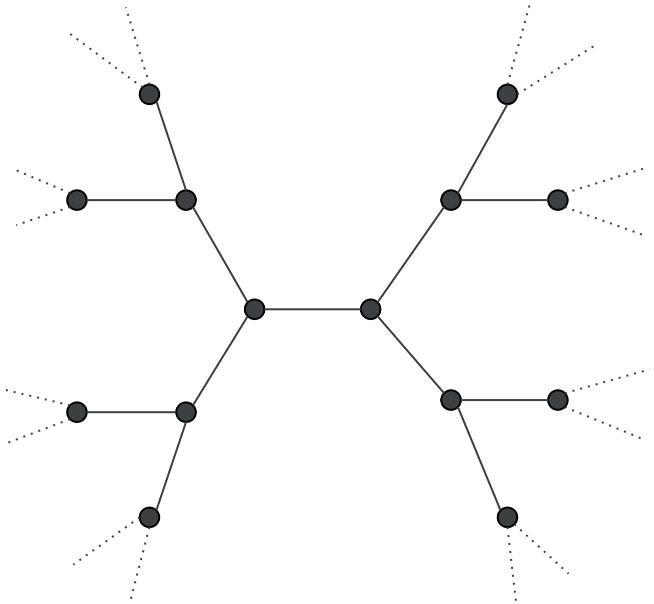

- $\Omega(\min(\Delta, \frac{\log\log n}{\log\log\log n}))$ randomized (BBHORS19)
  - Tight dependence on Δ
  - FOCS19 Best Paper Award, JACM
- Technique
  - Suppose r-round algo exists for P_0 = MM
  - Implies (r-1)-round algo for P_1, etc.
  - … Implies 0-round algo for P_r => contradiction
- Takeaways
  - Shows hardness for **regular trees**
  - More fine-grained than indistinguishability
  - Long line of excellent follow-up work
    - None break Ω(log log n) barrier
- Summary
  - Improve KMW04 with round elimination?
  - Must break the Ω(log log n) barrier

Our Result

Result: Any C **min(log Δ, $\sqrt{\text{log n}}$)**-round (randomized) LOCAL algorithm outputs a **maximal matching** with probability **at most** $\Delta^{-C'}$, even on **regular trees**.

- Applies to MIS on general graphs as well
  - By line graph reduction
- $C = 10^{-10}$, $C' = 1/1000$
- Simplified and applied to other problems by BCAO25

Function View of the LOCAL model

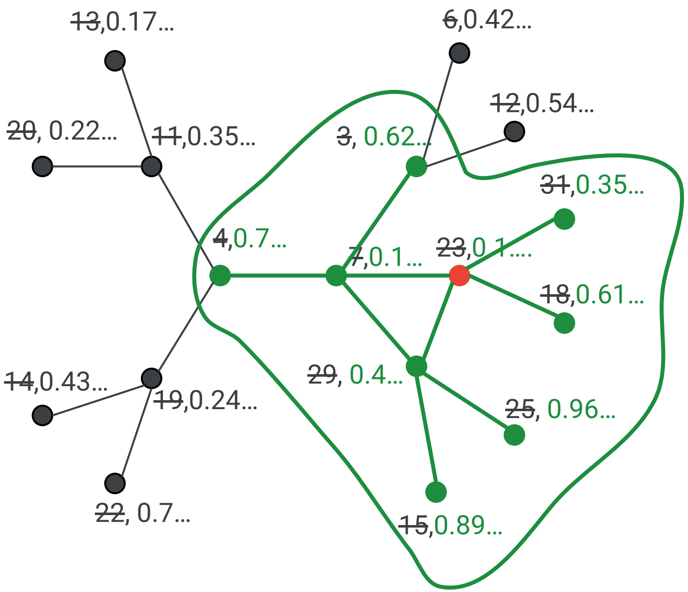

- (Randomized) LOCAL model
  - n vertex graph
  - Every vertex given two things
    - ~~O(log n) bit ID~~
      - ~~Unique~~
    - Infinite tape of randomness
      - Tapes are independent [0,1]
  - Algorithm: a function f
    - Input: r-neighborhood of a vertex
    - Output: portion of solution near vertex
    - **Question**: how small can r be?
- Key: **no unique IDs required**
  - Generate using randomness w.h.p.
- Useful view for hardness results
- LOCAL hardness implies CONGEST hardness

Implications of Our Result

Result: MM on regular trees requires $\Omega(\min(\log \Delta, \sqrt{\log n}))$ rounds.

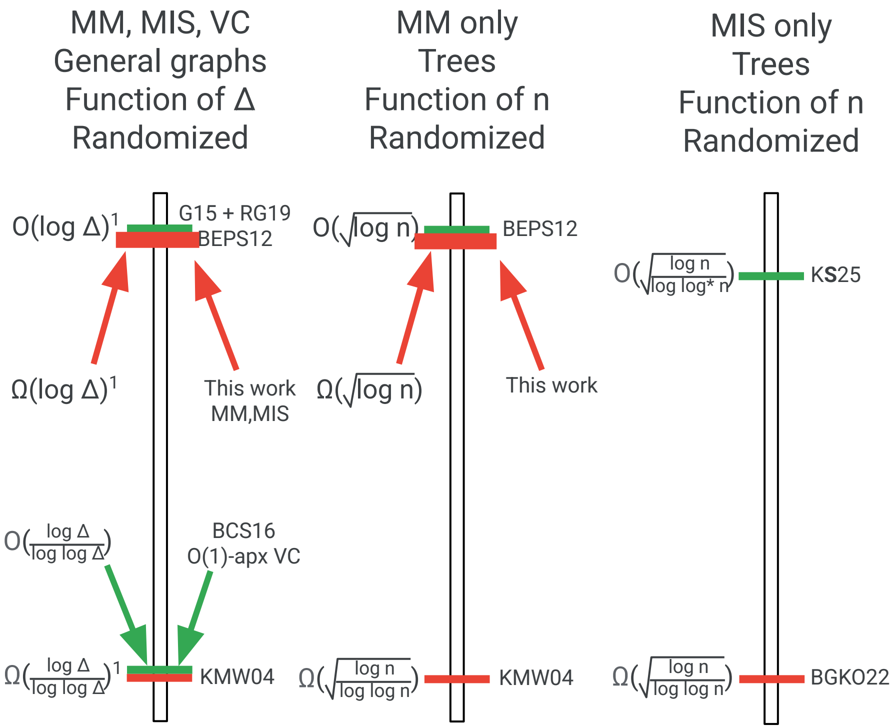

Green tick: algorithm. Red tick: lower bound.

[1]Functions of other params left out for simplicity.

1. First improvement as function of n in 20 years (since KMW04)
2. Closing Gaps (Optimality)
   a. $O(\log \Delta + \text{polyloglog } n)$ algo for MIS (G15) optimal
      i. For wide range of Δ
   b. MM on trees settled at $\Theta((\log n)^{1/2})$ (BEPS12)
3. Separations
   a. MM strictly harder than MIS on trees! (K**S**25)
      i. Opposite is true for general graphs!
      ii. KS25 breaks conjecture 11.15 in BE book
   b. MIS on trees easier than MIS on general graphs (K**S**25)
      i. Trees strictly easier than general graphs!
   c. O(1)-apx VC easier than MM/MIS (BCS16)
      i. For wide range of Δ
4. Broke $\Omega(\log \log n)$ barrier for round elimination
   a. Exponentially improved hardness for regular trees

3. MM

Setting the Scene

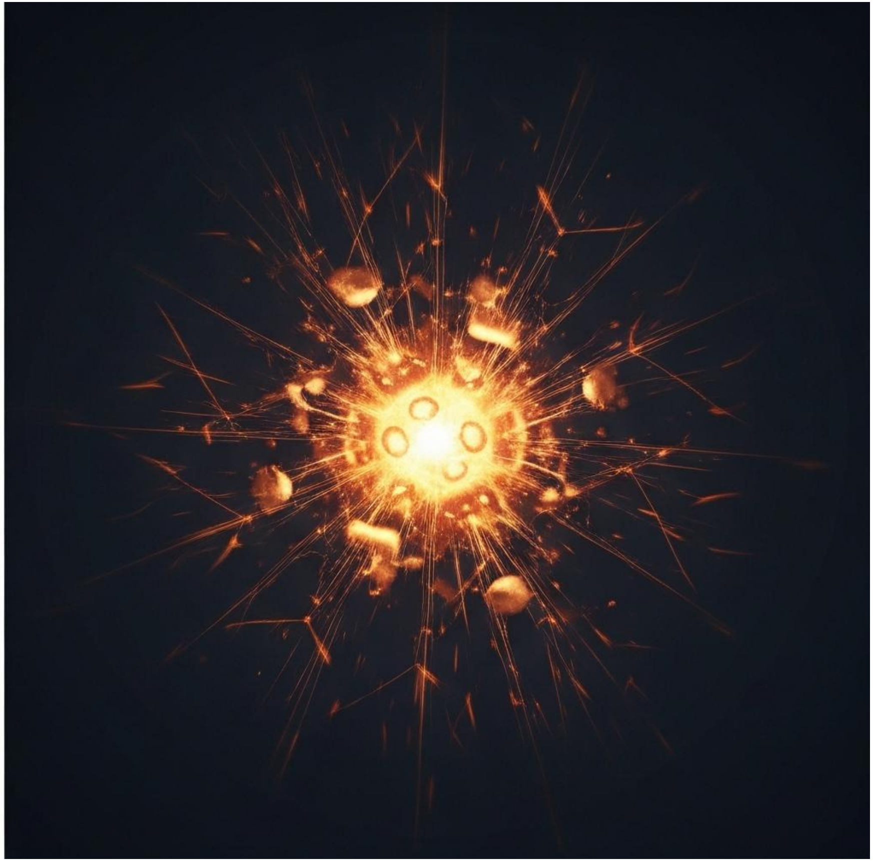

Our Hard Instance: ID Graphs

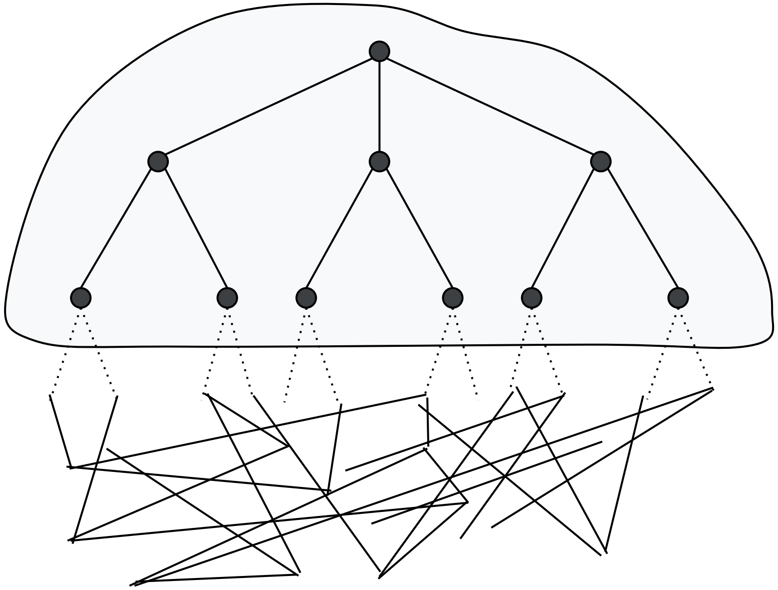

- Hard instance (ID graph) (B80, BCGGRV22)
  - Random Δ-regular graph
    - Configuration model (B80)
  - Construction
    - Effectively sampling Δ perfect matchings
    - Reality
      - n clusters of size Δ
      - Sample one perfect matching
      - Contract clusters
- Properties
  - Δ-regular
  - Girth $\Omega(\log_\Delta n)$
  - Maximum independent size $O(\frac{n \log \Delta}{\Delta})$
- Locally treelike, globally not

Google

Matching-Certified Algorithms

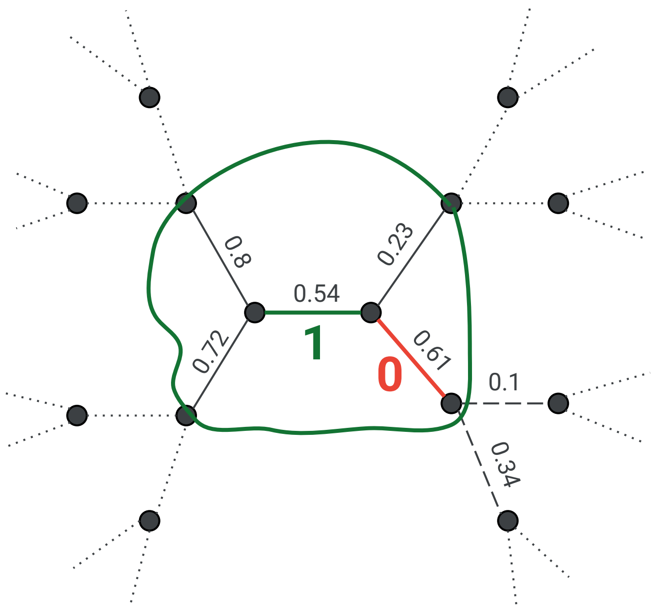

A 1-round matching-certified algorithm

- Key point of departure from round elimination
  - Decouple **maximality** from **matching**
  - Consider a matching algorithm
  - Assess its probability of achieving maximality
- Definition: r-round Matching-Certified Algorithms
  - Function that takes as input
    - r-neighborhood of edge e **in Δ-ary tree**
    - Randomness on edges
    - NOT IDs on edges
  - Returns
    - 1 if algorithm includes e, 0 otherwise
  - Matching-certified property
    - Adjacent edge r-neighborhoods can't both return 1

Lower Bounds for Matching-Certified Algos Suffice

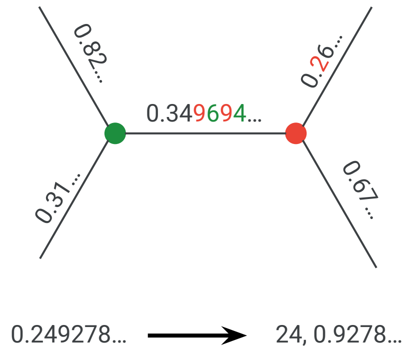

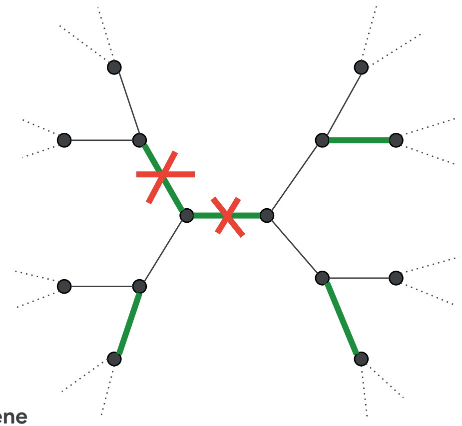

- Why can randomness be on edges, not vertices?
  - Use first O(log n) bits of ngbrs to name endpoints
  - Split subsequent randomness into blocks
  - Every other block allocated to same vert
- Why can IDs be dropped?
  - Yao Minimax
  - r-round algo with IDs + randomness => r-round algo with randomness only
    - Reserve first O(log n) bits as IDs
- What if algo doesn't always output a matching?
  - Postprocess output
    - Delete edges with shared endpoints
  - Maximal matching found with no lower prob

Matching-Certified Algorithm Example: Ranking (L86, ABI86)

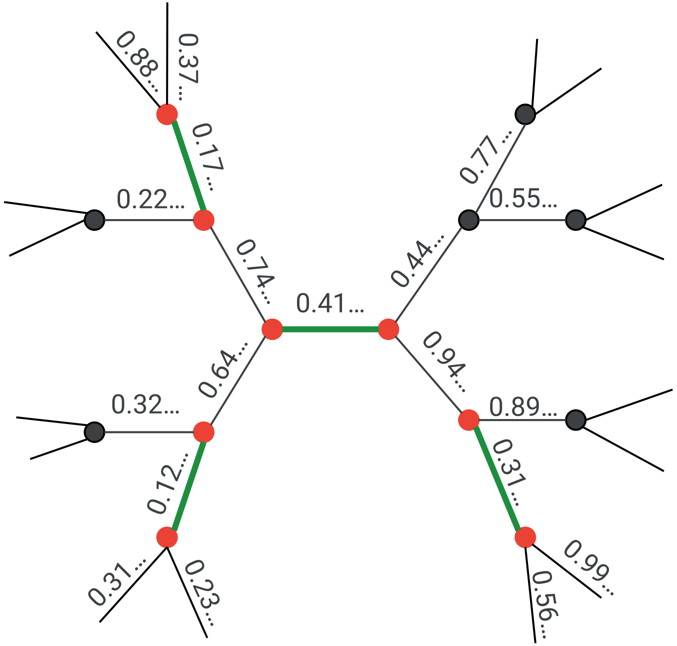

- Ranking for maximal matching
  - Every edge gets independent [0,1]
  - Edge in matching if minimum rank of incident
  - Delete matching-incident edges
  - Repeat till graph empty
- Matching-certified?
  - Yes
    - Matching found in each step
    - Deletion eliminates future conflicts
- Takeaways
  - Function of randomness only buys **symmetry**
    - Together with identical views
  - Vertices survive with **constant probability**
  - **Vertex survival probability** as a metric?

Google

Key Definition: Vertex Survival Probability (VSP)

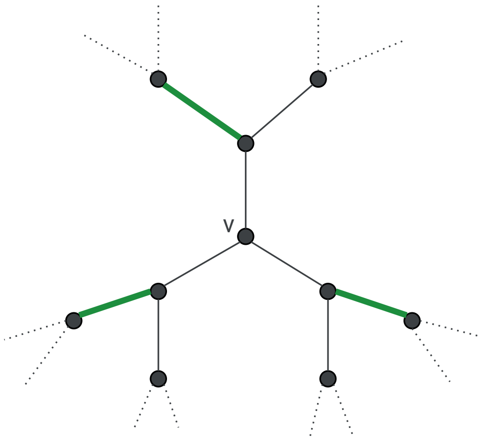

v surviving

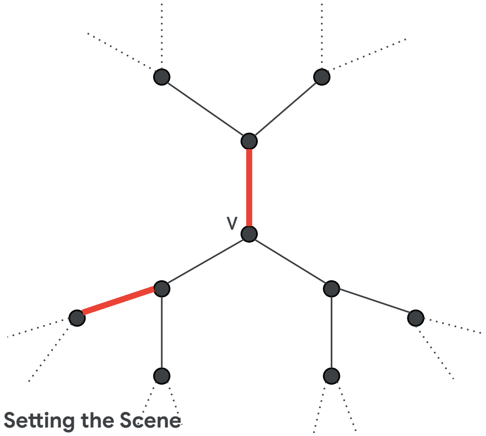

v not surviving

- VSP of r-round matching-certified (MC) algo f:
  - Pick uniformly random (r+1)-neighborhood of v
  - Δ incident edges to v (by regularity)
  - VSP(f) = Pr[f matches no incident edges to v]
- VSP(f) not function of v since v does not have ID
- Example: ranking (L86, ABI86)
  - VSP(r-round ranking) > $2^{-O(r)}$
- **Main question**: is VSP(f) > $2^{-O(r)}$ for **any** r-round MC f?
  - Yes!!

Key Theorem: Lower-Bounding Vertex Survival Probability

Key Theorem: Any r-round matching-certified algorithm f has vertex survival probability at least C^{-r}.

- Recall that matching-certified (MC) algorithms are defined only for Δ-ary trees
- $C = 10^{-80}$

Proving Key Theorem Via Round Elimination

Key (Round Elimination) Lemma:

For $r > 0$ and any **r-round** MC algo **f**, there exists an **(r-1)-round** MC algo **g** such that

$$VSP(g) \leq C\ VSP(f)$$

for some constant C.

Key Theorem:

For any r-round MC algo f, $VSP(f) \geq C^{-r}$ for constant C.

- Proof
  - Define $f_0 := f$.
  - Apply lemma r times:
    - For each $s > 0$, define
    - $f_s := g$ (obtained from f_{s-1})
  - By lemma, $VSP(f_r) \leq C^r\ VSP(f_0)$.
  - f_r is 0-round MC algo
    - $VSP(f_r) = 1$
      - Otherwise not an MC algo
  - Thus $VSP(f_0) \geq C^{-r}$ as desired. ✅

Our Results
Given Key Theorem

Google

Properties of the ID Graph

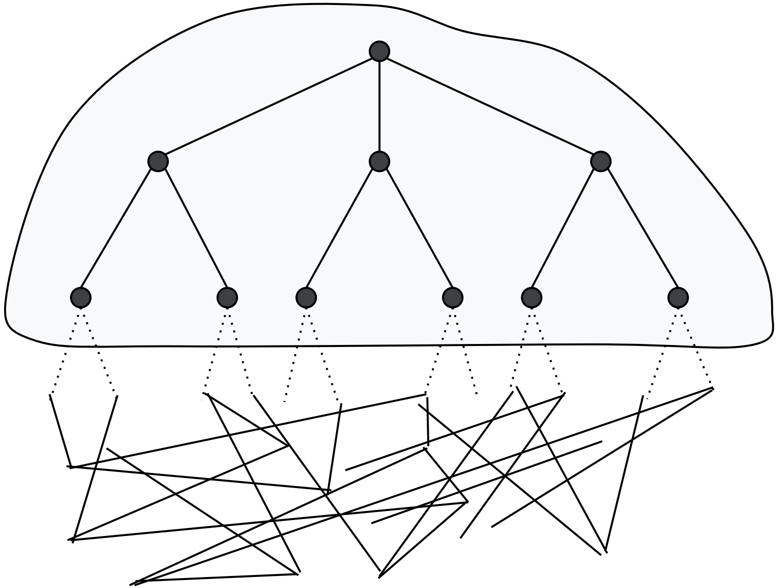

- Properties
  - Δ-regular
  - Girth $\Omega(\log_\Delta n)$
  - Maximum independent size $O(\frac{n \log \Delta}{\Delta})$
- Locally treelike, globally not

Google

Proof Sketch

Key Theorem: For any r-round MC algo f,

$VSP(f) \geq C^{-r}$

for constant C.

↓

Result: MM on regular trees requires

$\Omega(\min(\log \Delta, \sqrt{\log n}))$

rounds.

- Properties (ID Graph)
  - Δ-regular
  - Girth $\Omega(\log_\Delta n)$
  - Maximum independent size $O(\frac{n \log \Delta}{\Delta})$
- Locally treelike, globally not
  - Locally treelike => MC algorithms applicable
  - Globally not => surviving vertices not independent
    - Think $C^{-r} = 1/\Delta^{1/2}$
  - Two surviving vertices are adjacent
    - Matching not maximal! Contradiction
- Conclusion
  - Locally treelike when $r \leq O(\log_\Delta n)$ (below girth)
  - Surviving vertices not indep when $r \leq O(\log \Delta)$
    - As long as $VSP(f) \gg \frac{\log \Delta}{\Delta}$
  - Contradiction obtained if $r = \Theta(\min(\log \Delta, \log_\Delta n))$

Picture of Proof Sketch

Key Theorem: For any r-round MC algo f,
$VSP(f) \geq C^{-r}$
for constant C.

Result: MM on regular trees requires
$\Omega(\min(\log \Delta, \sqrt{\log n}))$
rounds.

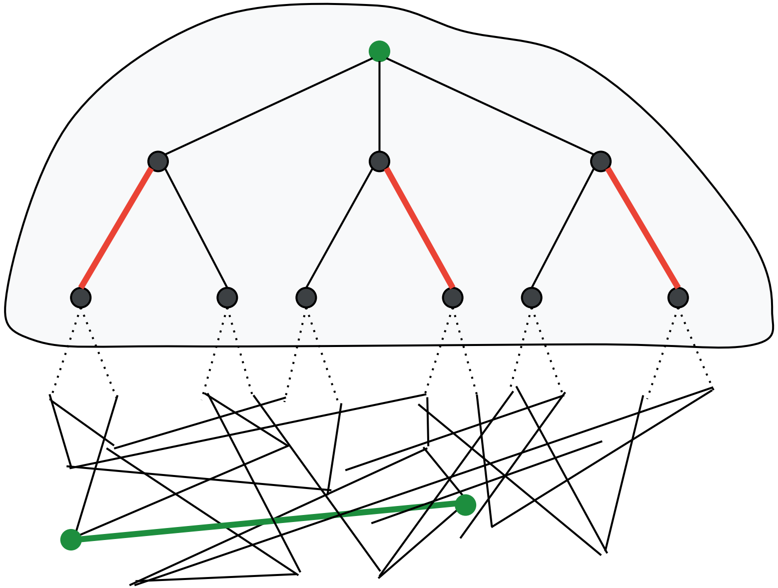

- Properties (ID Graph)
  - Δ-regular
  - Girth $\Omega(\log_\Delta n)$
  - Maximum independent size $O(\frac{n \log \Delta}{\Delta})$
- Locally treelike, globally not
  - Locally treelike => MC algorithms applicable
  - Globally not => surviving vertices not independent
    - Think $C^{-r} = 1/\Delta^{1/2}$
  - Two surviving vertices are adjacent
    - Matching not maximal! Contradiction
- Conclusion
  - Locally treelike when $r \leq O(\log_\Delta n)$ (below girth)
  - Surviving vertices not indep when $r \leq O(\log \Delta)$
    - As long as $VSP(f) \gg \frac{\log \Delta}{\Delta}$
  - Contradiction obtained if $r = \Theta(\min(\log \Delta, \log_\Delta n))$

Comparison of Approach with Standard Round Elimination (RE)

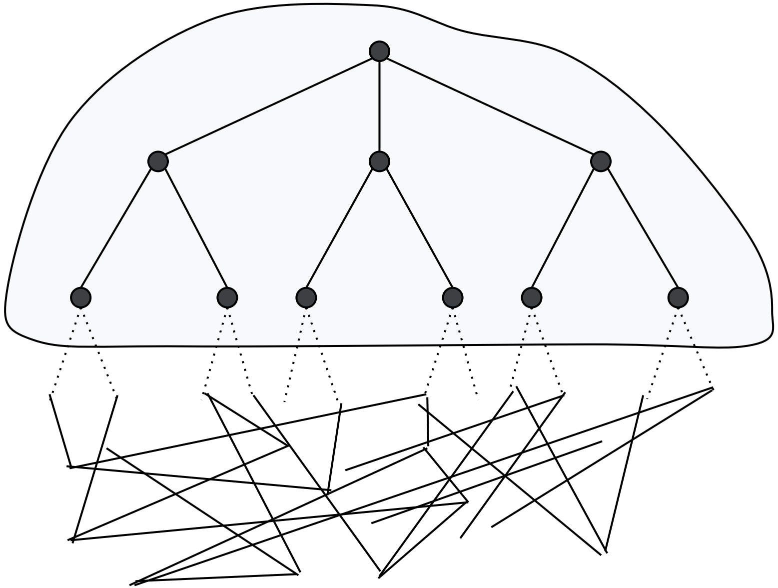

- Decoupling maximality from matching
  - LCLs model all problem constraints
    - Leads to formulation change for MM with each round elimination
    - Note: Fixed points do exist
  - Our approach: LCL with extra metric
    - Matching with VSP
      - Low rounds => high VSP
- Contradiction obtained globally
  - high VSP contradicts small max IS
  - Only local structure used for standard RE

Google

MM Summary

Result: MM on regular trees requires $\Omega(\min(\log \Delta, \sqrt{\log n}))$ rounds.

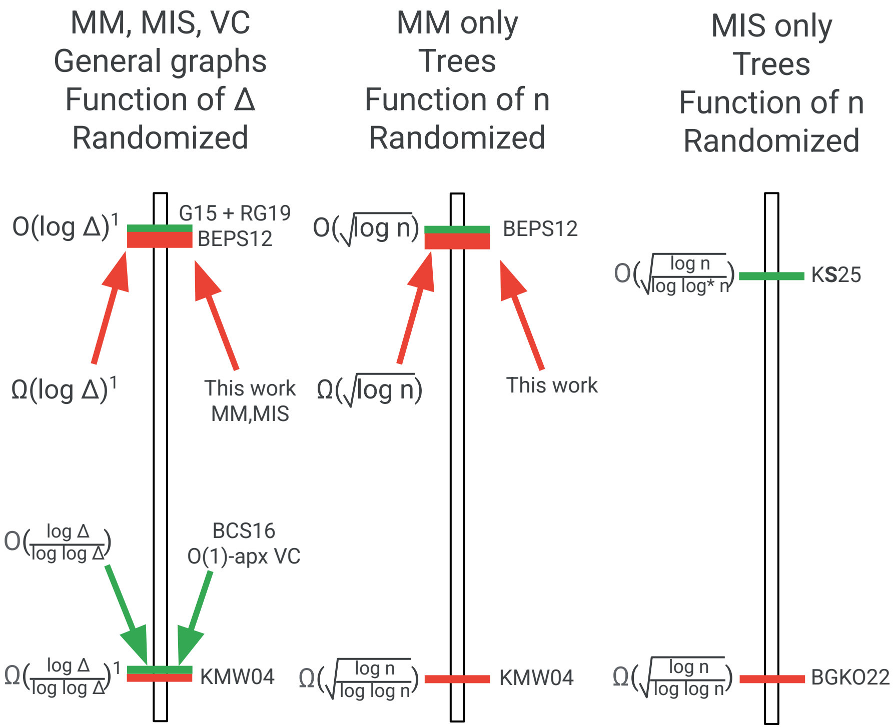

Green tick: algorithm. Red tick: lower bound.
[1]Functions of other params left out for simplicity.

- Many implications and separations for MIS, MM, VC
  - 2 separations implied by MIS algo of K**S**25
    - MIS easier than MM in trees
    - MIS on trees easier than MIS on general graphs
- Round elimination framework for randomized algorithms
  - Exponential improvement for regular graphs
- Key idea: decouple maximality from matching
- Key new vertex survival probability (VSP) metric
- Future directions
  - Round elimination framework far more general
    - Extend to those problems also?
      - BCAO25 has started this effort
    - Does decoupling work for other problems?

Conclusion

Google

Results and High-Level Takeaways

Algorithm: MIS on (1) girth ≥$\Omega(1)$ in $O(\frac{\log \Delta}{\log \log^* \Delta}$ + polyloglog n) rounds (2) trees in $O(\sqrt{\frac{\log n}{\log \log^* n}})$ rounds

Hardness: MM on regular trees requires $\Omega(\min(\log \Delta, \sqrt{\log n}))$ rounds.

1. Implications (Randomized LOCAL)
   a. Tight understanding of MM and MIS in some regimes
   b. Breaking barriers to crack open other regimes
2. Techniques
   a. Algorithms
      i. All prior randomized algorithms for MIS (Luby86, G15, …) based on same 1-round primitive
      ii. We break that barrier via **faster degree reduction**
   b. Hardness
      i. New type of round elimination framework
3. Both results obtained by deeply understanding **vertex survival probability**

Open Problems

Algorithm: MIS on (1) girth ≥$\Omega(1)$ in $O(\frac{\log \Delta}{\log \log^* \Delta}$ + polyloglog n) rounds (2) trees in $O(\sqrt{\frac{\log n}{\log \log^* n}})$ rounds

Hardness: MM on regular trees requires $\Omega(\min(\log \Delta, \sqrt{\log n}))$ rounds.

1. Better Algorithms
   a. Can log* be replaced with something larger for MIS on trees?
   b. Can faster degree reduction be used to obtain o(log n) for MIS on general graphs?
2. Other Problems
   a. These techniques enable optimal hardness/algorithms
   b. Can they be applied to other problems?
      i. BCAO25 have made progress on this

Thank you!