The Best-of-Three Voting on Dense Graphs

Nan Kang[1] Nicolás Rivera[2]

[1]Department of Informatics
King's College London

[2]Cambridge Computer Laboratory
University of Cambridge

07/Feb/2019

Contents

Best-of-k protocol

- Consider a graph $G = (V, E)$ with $|V| = n$, in which every vertex has an initial opinion.
- At each time step, every vertex randomly samples k neighbours with replacement, and adopts the majority opinion.
- (if no majority: wait or picks a random popular opinion.)
- Consensus time? Reflects initial majority?

Best-of-k protocol

- Consider a graph $G = (V, E)$ with $|V| = n$, in which every vertex has an initial opinion.
- At each time step, every vertex randomly samples k neighbours with replacement, and adopts the majority opinion.
- (if no majority: wait or picks a random popular opinion.)
- Consensus time? Reflects initial majority?

Best-of-k protocol

- Consider a graph $G = (V, E)$ with $|V| = n$, in which every vertex has an initial opinion.
- At each time step, every vertex randomly samples k neighbours with replacement, and adopts the majority opinion.
- (if no majority: wait or picks a random popular opinion.)
- Consensus time? Reflects initial majority?

Best-of-k protocol

- Consider a graph $G = (V, E)$ with $|V| = n$, in which every vertex has an initial opinion.
- At each time step, every vertex randomly samples k neighbours with replacement, and adopts the majority opinion.
- (if no majority: wait or picks a random popular opinion.)
- Consensus time? Reflects initial majority?

Example: Best-of-1

Initially, each vertex is assigned a colour of either red or blue.

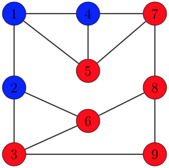

In each step, every vertex adopts the opinion of a random neighbour.

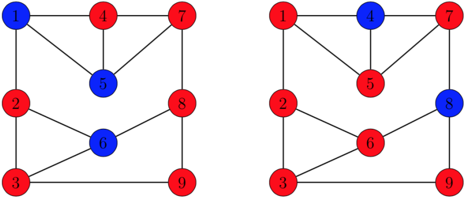

(a) Possible opinions of each node after the first step. $\Longrightarrow$ (b) Possible pinions of each node after the second step.

Example: Best-of-3

Initially, each vertex is assigned a colour of either red or blue.

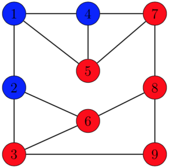

In each step, every vertex adopts majority opinion of 3 random neighbours.

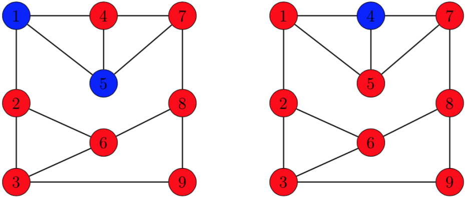

(a) Opinions of each node after the first step. $\Longrightarrow$ (b) Opinions of each node after the second step.

Previous work: $k = 1$

Consensus time under Best-of-1 protocol: (voter model)

- non-bipartite graphs
- Pr(consensus to *OpinionA*) is proportional to $\sum_v d_v$, where v has *OpinionA*.
- $\Theta(n)$ w.h.p in K_n.

  [Yehuda Hassin and David Peleg. *Distributed probabilistic polling and applications to proportionate agreement.* 2001.]

Previous work: $k = 1$

Consensus time under Best-of-1 protocol: (voter model)

- non-bipartite graphs
- Pr(consensus to *OpinionA*) is proportional to $\sum_v d_v$, where v has *OpinionA*.
- $\Theta(n)$ w.h.p in K_n.

  [Yehuda Hassin and David Peleg. *Distributed probabilistic polling and applications to proportionate agreement.* 2001.]

Previous work: $k = 1$

Consensus time under Best-of-1 protocol: (voter model)

- non-bipartite graphs
- Pr(consensus to *OpinionA*) is proportional to $\sum_v d_v$, where v has *OpinionA* .
- $\Theta(n)$ w.h.p in K_n .

  [Yehuda Hassin and David Peleg. *Distributed probabilistic polling and applications to proportionate agreement.* 2001.]

Previous work: $k = 2$

Consensus time under Best-of-2 protocol:

- Converge to majority under appropriate conditions.
- $O(\log n)$ w.h.p in expanders.

  [Colin Cooper, Robert Els asser, Tomasz Radzik, Nicolas Rivera, and Takeharu Shiraga. *Fast consensus for voting on general expander graphs.* 2015.]

Previous work: $k = 2$

Consensus time under Best-of-2 protocol:

- Converge to majority under appropriate conditions.
- $O(\log n)$ w.h.p in expanders.

  [Colin Cooper, Robert Els asser, Tomasz Radzik, Nicolas Rivera, and Takeharu Shiraga. *Fast consensus for voting on general expander graphs.* 2015.]

Previous works: $k = 3$

Consensus time under Best-of-3 protocol:

- $O(\log n)$ w.h.p in K_n with more than two opinions.

  [Luca Becchetti, Andrea Clementi, Emanuele Natale, Francesco Pasquale, Riccardo Silvestri, and Luca Trevisan. *Simple dynamics for plurality consensus.* 2014.]

- $O(\log n)$ w.h.p in expanders.

  [Colin Cooper, Tomasz Radzik, Nicola s Rivera, and Takeharu Shiraga. *Fast plurality consensus in regular expanders.* 2016.]

Previous works: $k = 3$

Consensus time under Best-of-3 protocol:

- $O(\log n)$ w.h.p in K_n with more than two opinions.

  [Luca Becchetti, Andrea Clementi, Emanuele Natale, Francesco Pasquale, Riccardo Silvestri, and Luca Trevisan. *Simple dynamics for plurality consensus.* 2014.]

- $O(\log n)$ w.h.p in expanders.

  [Colin Cooper, Tomasz Radzik, Nicola s Rivera, and Takeharu Shiraga. *Fast plurality consensus in regular expanders.* 2016.]

Previous works: $k \geq 5$

Consensus time under Best-of-5 protocol:

- $O(\log\log n)$ w.h.p in almost all graphs with a given degree sequence;
- $O(\log n)$ w.h.p in d-regular graphs, $d \geq 5$;
- $O(\log\log n)$ in $G_{n,p}$ w.h.p with $p = O(\log n/n)$.

  [Mohammed Amin Abdullah, Moez Draief. *Consensus on the initial global majority by local majority polling for a class of sparse graphs.* 2013]

Why Best-of-3?

- Best-of-1 is slow, not a desired model when consensus to majority is required.
- Best-of-2 and 3 take $O(\log n)$, while Best-of-5 takes $O(\log \log n)$ from previous work.
- So, we want to close the gap between Best-of-3 and 5. Fast consensus time and low cost.

Why Best-of-3?

- Best-of-1 is slow, not a desired model when consensus to majority is required.
- Best-of-2 and 3 take $O(\log n)$, while Best-of-5 takes $O(\log \log n)$ from previous work.
- So, we want to close the gap between Best-of-3 and 5. Fast consensus time and low cost.

Why Best-of-3?

- Best-of-1 is slow, not a desired model when consensus to majority is required.
- Best-of-2 and 3 take $O(\log n)$, while Best-of-5 takes $O(\log \log n)$ from previous work.
- So, we want to close the gap between Best-of-3 and 5. Fast consensus time and low cost.

Illustration: initialisation

(synchronous, two-party)
At the beginning, each vertex is randomly assigned a colour of either red or blue.

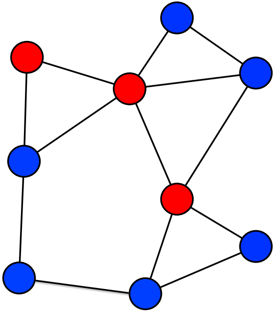

Figure: Step 0

In each step, every vertex samples three random neighbours, and assumes the majority colour.

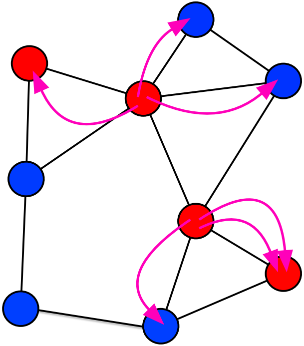
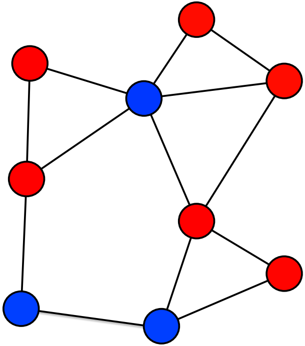

Figure: Step 1

Illustration: sampling (step 2)

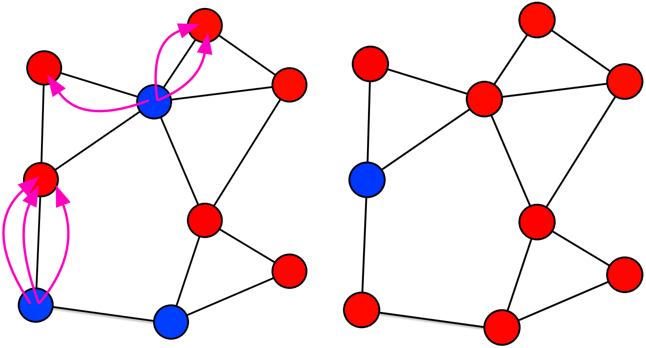

Figure: Step 2

Illustration: sampling (step 3)

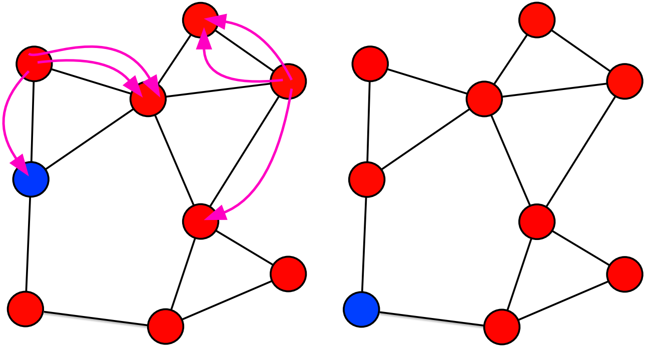

Figure: Step 3

Illustration: sampling (step 4)

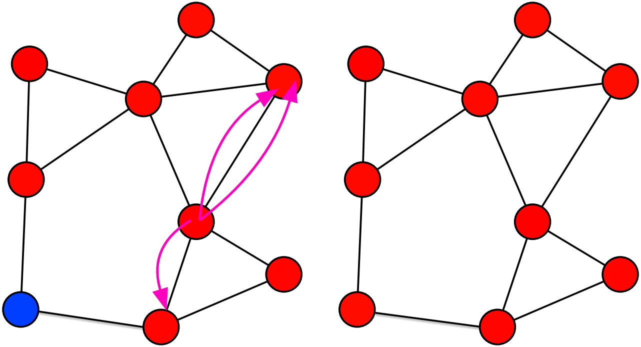

Figure: Step 4

Main results

Thereom.
Consider the Best-of-Three protocol on a graph G of n vertices. Initially, each vertex of G is assigned a colour **R** with probability $\frac{1}{2}+\delta$, and **B** with probability $\frac{1}{2}-\delta$, where $\delta \in (0, \frac{1}{2})$.

- If G is a graph with minimum degree $d = n^{\Omega(1/\log_2 \log_2 n)}$, and $\delta = \log_2 n^{-O(1)}$, then with probability $1 - O(1/n)$, every vertex of G has colour **R** after $O(\log_2 \log_2 n) + O\left(\log_2\left(\delta^{-1}\right)\right)$ timesteps.
- Particularly, if G is a complete graph and $\delta = \log_2 n^{-O(1)}$, with probability $1 - O(1/n)$, every vertex of G has colour **R** after $\frac{21}{16}\log_2 \log_2 n + \frac{16}{5}\log_2\left(\delta^{-1}\right)$ timesteps.

Main results

Thereom.
Consider the Best-of-Three protocol on a graph G of n vertices. Initially, each vertex of G is assigned a colour **R** with probability $\frac{1}{2} + \delta$, and **B** with probability $\frac{1}{2} - \delta$, where $\delta \in (0, \frac{1}{2})$.

- If G is a graph with minimum degree $d = n^{\Omega(1/\log_2 \log_2 n)}$, and $\delta = \log_2 n^{-O(1)}$, then with probability $1 - O(1/n)$, every vertex of G has colour **R** after $O(\log_2 \log_2 n) + O\left(\log_2 \left(\delta^{-1}\right)\right)$ timesteps.
- Particularly, if G is a complete graph and $\delta = \log_2 n^{-O(1)}$, with probability $1 - O(1/n)$, every vertex of G has colour **R** after $\frac{21}{16} \log_2 \log_2 n + \frac{16}{5} \log_2 \left(\delta^{-1}\right)$ timesteps.

Main results

Thereom.
Consider the Best-of-Three protocol on a graph G of n vertices. Initially, each vertex of G is assigned a colour **R** with probability $\frac{1}{2}+\delta$, and **B** with probability $\frac{1}{2}-\delta$, where $\delta \in (0, \frac{1}{2})$.

- If G is a graph with minimum degree $d = n^{\Omega(1/\log_2 \log_2 n)}$, and $\delta = \log_2 n^{-O(1)}$, then with probability $1 - O(1/n)$, every vertex of G has colour **R** after $O(\log_2 \log_2 n) + O\left(\log_2\left(\delta^{-1}\right)\right)$ timesteps.
- Particularly, if G is a complete graph and $\delta = \log_2 n^{-O(1)}$, with probability $1 - O(1/n)$, every vertex of G has colour **R** after $\frac{21}{16}\log_2 \log_2 n + \frac{16}{5}\log_2\left(\delta^{-1}\right)$ timesteps.

Idea of the structure: Pseudo-tree

For an arbitrary vertex v of G, we investigate the process of v updating its opinion in a reverse chronological order.

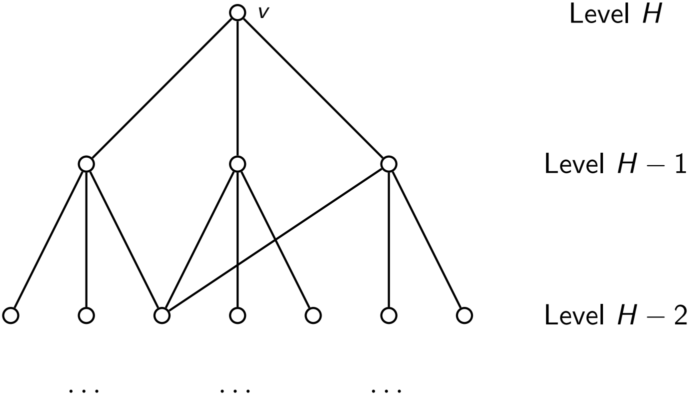

Figure: Updating process of a vertex: pseudo-tree.

Clashes in the process

- random sampling $\longrightarrow$ clashes $\longrightarrow$ dependency

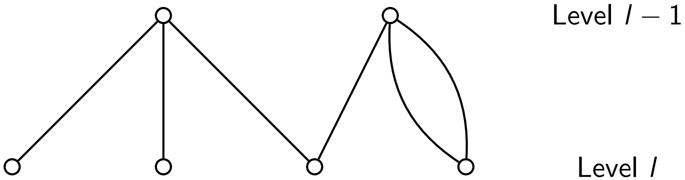

Figure: Example of clashes.

- clash: in the pseudo-tree, a vertex at level l has more than one parent at level $l-1$.
- no clash $\longrightarrow$ ternary tree $\longrightarrow$ independent opinions

Clashes in the process

- random sampling $\longrightarrow$ clashes $\longrightarrow$ dependency

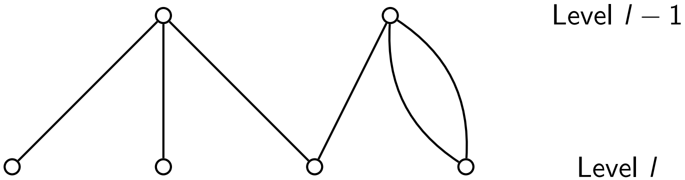

Figure: Example of clashes.

- clash: in the pseudo-tree, a vertex at level l has more than one parent at level $l-1$.
- no clash $\longrightarrow$ ternary tree $\longrightarrow$ independent opinions

Clashes in the process

- random sampling $\longrightarrow$ clashes $\longrightarrow$ dependency

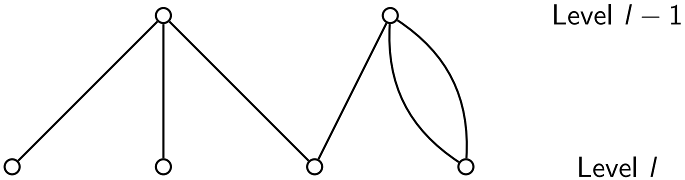

Figure: Example of clashes.

- clash: in the pseudo-tree, a vertex at level l has more than one parent at level $l-1$.
- no clash $\longrightarrow$ ternary tree $\longrightarrow$ independent opinions

Best-of-3 in graphs with min degree $d = n^\alpha$

1. Choose an arbitrary vertex, and construct a pseudo-tree its opinion-updating process.
2. To deal with the dependency resulting from clashes:

- Level 0 to T: Sprinkling model
- Level T to H: Duplicating model

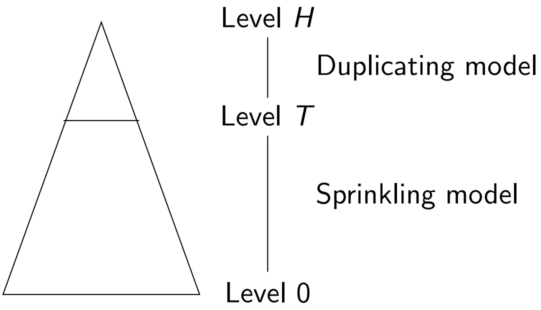

Sprinkling model: idea

- From level 0 to T : bottom to top, left to right. (add extra blues)

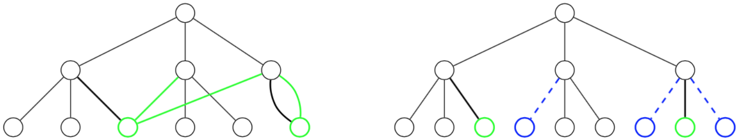

(a) Part of the original tree with clashes. $\Longrightarrow$ (b) New model with extra blue nodes.

Figure: Construction of the Sprinkling model.

- $\mathbb{P}\{\text{a vertex at level } i \text{ having clashes}\} \leq \frac{3^{H-i}-1}{n-1} \approx \frac{3^{H-i}}{n} = \varepsilon_i$

Sprinkling model: analysis

- p_i : probability of a vertex at level i being blue, $p_0 = \frac{1}{2} - \delta$. $\varepsilon_i = 3^{H-i}/d = O(\log n)/d$.
- Let $T_1 = O\left(\log_2(\delta^{-1})\right)$, $T_2 = \log_2 \log_2 n$, and $T = T_1 + T_2 + const$. Then,

$$\begin{aligned} T_1 &: p_{T_1} < 1/2 - 1/(2\sqrt{3}) ; \\ T-1 &: p_{T_1+T_2+2} < \varepsilon_{T_1+T_2+2} = O\left((\log n)/d\right) ; \text{and} \\ T &: p_{T_1+T_2+3} < \varepsilon^2_{T_1+T_2+3} = O\left((\log n^2)/d^2\right) . \end{aligned}$$

Sprinkling model: analysis

- p_i : probability of a vertex at level i being blue, $p_0 = \frac{1}{2} - \delta$. $\varepsilon_i = 3^{H-i}/d = O(\log n)/d$.
- Let $T_1 = O\left(\log_2(\delta^{-1})\right)$, $T_2 = \log_2 \log_2 n$, and $T = T_1 + T_2 + const$. Then,

$$\begin{aligned} T_1 &: p_{T_1} < 1/2 - 1/(2\sqrt{3}); \\ T-1 &: p_{T_1+T_2+2} < \varepsilon_{T_1+T_2+2} = O\left((\log n)/d\right) \text{; and} \\ T &: p_{T_1+T_2+3} < \varepsilon^2_{T_1+T_2+3} = O\left((\log n^2)/d^2\right) . \end{aligned}$$

Sprinkling model: upper bounds

- $T = T_1 + T_2 + 3 = log_2 \log_2 n + \frac{16}{5} \log_2(\delta^{-1})$.

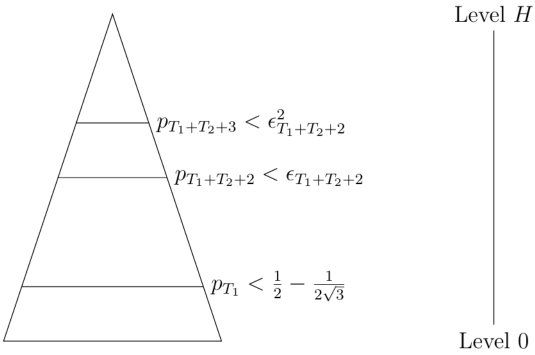

Figure: Upper bounds of p_i given by the Sprikling model.

- Deal with the top part from level T to H?

Duplicating model

- Consider the top part of the pseudo-tree from level T to H.
- Transform the top part into a ternary sub-tree.
- $B_T = \#$ blues at level T of the pseudo-tree after Sprinkling, $B'_T = \#$blues at level T of the ternary sub-tree after Duplicating.

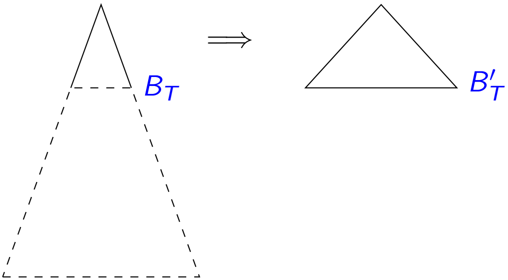

Duplicating model: idea

Level T to H: bottom to top, left to right. (add copy + subtree)

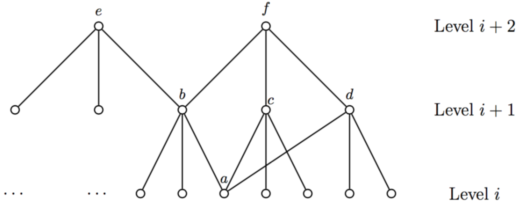

(a) Vertices with clashes.

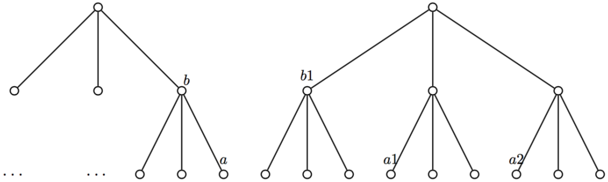

(b) Duplicating the vertex and its subtree.

Duplicating model: an upper bound on blue vertices

- Recall that $B_T = \#$ blue vertices at level T of the pseudo-tree after Sprinkling, and $B'_T = \#$ blues at level T of the ternary sub-tree after Duplicating.

$$\mathbb{P}(\text{root is blue}) \leq \mathbb{P}\left(B'_T \geq 2^{H-T}\right) .$$

- Let $K = \#$ levels containing clashes from level T to H, then an upper bound is: $B'_T \leq B_T \cdot 2^K$.
- $\mathbb{P}(\text{root is blue}) \leq \mathbb{P}\left(B'_T \geq 2^{H-T}\right) \leq \mathbb{P}\left(B_T \geq 2^{H-T-K}\right)$.

Duplicating model: an upper bound on blue vertices

- Recall that $B_T = \#$ blue vertices at level T of the pseudo-tree after Sprinkling, and $B'_T = \#$ blues at level T of the ternary sub-tree after Duplicating.

$$\mathbb{P}(\text{root is blue}) \leq \mathbb{P}\left(B'_T \geq 2^{H-T}\right) .$$

- Let $K = \#$ levels containing clashes from level T to H, then an upper bound is: $B'_T \leq B_T \cdot 2^K$.
- $\mathbb{P}(\text{root is blue}) \leq \mathbb{P}\left(B'_T \geq 2^{H-T}\right) \leq \mathbb{P}\left(B_T \geq 2^{H-T-K}\right)$.

Duplicating model: an upper bound on blue vertices

- Recall that $B_T = \#$ blue vertices at level T of the pseudo-tree after Sprinkling, and $B'_T = \#$ blues at level T of the ternary sub-tree after Duplicating.

$$\mathbb{P}(\text{root is blue}) \leq \mathbb{P}\left(B'_T \geq 2^{H-T}\right) .$$

- Let $K = \#$ levels containing clashes from level T to H, then an upper bound is: $B'_T \leq B_T \cdot 2^K$.
- $\mathbb{P}(\text{root is blue}) \leq \mathbb{P}\left(B'_T \geq 2^{H-T}\right) \leq \mathbb{P}\left(B_T \geq 2^{H-T-K}\right)$.

Duplicating model: analysis

- $\mathbb{P}(\text{root is blue}) \le \mathbb{P}\left(B'_T \ge 2^{H-T}\right) \le \mathbb{P}\left(B_T \ge 2^{H-T-K}\right)$.
- $K \prec Bin\left(H-T, 9^{H-T}/d\right)$,
  $B_T \preceq Bin(3^{H-T}, 3^{H-T}/d)$.
- To ensure $\mathbb{P}(\text{root is blue}) \le \mathbb{P}\left(B_T \cdot 2^K \ge 2^{H-T}\right) = O(1/n^2)$, it is required that

$$d = n^{\Omega(1/\log_2 \log_2 n)} .$$

  Hence, $\mathbb{P}(\text{blue wins}) = O(1/n)$ by union bound .

Duplicating model: analysis

- $\mathbb{P}(\text{root is blue}) \leq \mathbb{P}\left(B'_T \geq 2^{H-T}\right) \leq \mathbb{P}\left(B_T \geq 2^{H-T-K}\right)$.
- $K \prec Bin\left(H-T, 9^{H-T}/d\right)$,
  $B_T \preceq Bin(3^{H-T}, 3^{H-T}/d)$.
- To ensure $\mathbb{P}(\text{root is blue}) \leq \mathbb{P}\left(B_T \cdot 2^K \geq 2^{H-T}\right) = O(1/n^2)$, it is required that

$$d = n^{\Omega(1/\log_2 \log_2 n)} .$$

Hence, $\mathbb{P}(\text{blue wins}) = O(1/n)$ by union bound .

Duplicating model: analysis

- $\mathbb{P}(\text{root is blue}) \leq \mathbb{P}\left(B'_T \geq 2^{H-T}\right) \leq \mathbb{P}\left(B_T \geq 2^{H-T-K}\right)$.
- $K \prec Bin\left(H - T, 9^{H-T}/d\right)$,
  $B_T \preceq Bin(3^{H-T}, 3^{H-T}/d)$.
- To ensure $\mathbb{P}(\text{root is blue}) \leq \mathbb{P}\left(B_T \cdot 2^K \geq 2^{H-T}\right) = O(1/n^2)$, it is required that

$$d = n^{\Omega(1/\log_2 \log_2 n)} .$$

Hence, $\mathbb{P}(\text{blue wins}) = O(1/n)$ by union bound .

Theorem in graphs with min degree $d = n^\alpha$

Theorem:
Let G be a graph of n vertices, such that each vertex of G is initially assigned a colour **R** with probability $\frac{1}{2} + \delta$, and **B** with probability $\frac{1}{2} - \delta$, where $\delta \in (0, \frac{1}{2})$. Under the Best-of-Three protocol:

- If G has minimum degree $d = n^{\Omega(1/\log_2 \log_2 n)}$, and $\delta = \log_2 n^{-O(1)}$, then with probability $1 - O(1/n)$, every vertex of G has colour **R** after $O(\log_2 \log_2 n) + O\left(\log_2\left(\delta^{-1}\right)\right)$ timesteps.
- (does not depend on the structure of G)

Theorem in graphs with min degree $d = n^\alpha$

Theorem:
Let G be a graph of n vertices, such that each vertex of G is initially assigned a colour **R** with probability $\frac{1}{2} + \delta$, and **B** with probability $\frac{1}{2} - \delta$, where $\delta \in (0, \frac{1}{2})$. Under the Best-of-Three protocol:

- If G has minimum degree $d = n^{\Omega(1/\log_2 \log_2 n)}$, and $\delta = \log_2 n^{-O(1)}$, then with probability $1 - O(1/n)$, every vertex of G has colour **R** after $O(\log_2 \log_2 n) + O\left(\log_2 \left(\delta^{-1}\right)\right)$ timesteps.
- (does not depend on the structure of G)

Theorem in graphs with min degree $d = n^{\alpha}$

Theorem:
Let G be a graph of n vertices, such that each vertex of G is initially assigned a colour **R** with probability $\frac{1}{2} + \delta$, and **B** with probability $\frac{1}{2} - \delta$, where $\delta \in (0, \frac{1}{2})$. Under the Best-of-Three protocol:

- If G has minimum degree $d = n^{\Omega(1/\log_2 \log_2 n)}$, and $\delta = \log_2 n^{-O(1)}$, then with probability $1 - O(1/n)$, every vertex of G has colour **R** after $O(\log_2 \log_2 n) + O\left(\log_2 \left(\delta^{-1}\right)\right)$ timesteps.
- (does not depend on the structure of G)

Future Work

Further research may concern the following special classes of graphs:

- High dimensional grids
- Hypercubes
- $G_{n,p}$ with small p
- Regular expanders with low degree
- (more than two opinions)