Princeton COS 521: Advanced Algorithm Design Fall 2018

Expansion and Consensus for Social Networks

Names: Natalie Collina, Albert Zuo

1 Introduction

Our paper builds upon results from Feldman et. al.'s paper, Consensus via non-Bayesian Asynchronous Learning in Social Networks, which determines certain classes of social network graphs which convergence to a consensus under an asynchronous updating model. This paper contributes to the question of social network convergence in two ways. In part 1, we establish that the decoupled analysis approach performed in [1] is insufficient to derive a result that holds for all λ-expanders. In part 2, we prove that all d-regular λ-expanders converge to consensus with high probability, even if d is on the scale of n.

2 Model and Definitions

We start with a graph $G(V, E)$, with $|V| = n$, and some "ground truth" for the entire graph, say *red*. Initialize the graph as follows: give each vertex a private signal $X(v) \in \{red, blue\}$ randomly, with bias toward the ground truth. That is, with probability $\frac{1}{2} + \delta$, $X(v) = red$, and with probability $\frac{1}{2} - \delta$, $X(v) = blue$. Note that the value of $X(v)$ is known only to v.

Let $C^t(v)$ be the public color of vertex v (all vertices can observe $C^t(v)$), where $C^t(v) \in \{red, blue, uncolored\}$. At the start, $C^0(v) = uncolored\ \forall v$. Let $N_B^t(v)$ be the number of blue neighbors vertex v has at time t, and let $N_R^t(v)$ be the number of red neighbors vertex v has at time t. At each time t, pick a vertex v uniformly at random, and update $C^t(v)$ in the following way: if $N_R^t(v) > N_B^t(v)$, then $C^t(v) = red$. If $N_R^t(v) < N_B^t(v)$, then $C^t(v) = blue$. If $N_R^t(v) = N_B^t(v)$, then $C^t(v) = X(v)$. Stop when $C^t(v) = C^k(v)$, $\forall v$, $\forall k \geq t+1$ (i.e. no vertex will ever change their color).

Definition 1: Weighted Adjacency Matrix. Let $d(v)$ be the degree of vertex v. For a graph G, the weighted adjacency matrix $M(G)$ is an $n \times n$ matrix defined by

$$M(i,j) = \begin{cases} \frac{1}{\sqrt{d(i)d(j)}} & \text{if } i \text{ and } j \text{ are adjacent in } G \\ 0 & \text{otherwise} \end{cases}$$

Definition 2: λ-expander. A graph G is a λ-expander if all but the first eigenvalues of $M(G)$ lie in $[-\lambda, \lambda]$.

3 Related Work

Feldman et. al's main result shows that with probability at least $1 - O(\frac{1}{(\delta \ln \ln n)^2})$, applying the above process to max-degree d λ-expanders with $\lambda \leq \frac{\delta}{6}$ will result in a red (correct) consensus.
Feldman's is composed of two steps: they first show that max-degree d λ-expanders will reach a red supermajority (in terms of volume) with high probability. In the second step, they show that if a max-degree d λ-expander contains a red supermajority, it will converge to a red consensus with high probability. In the second step, the authors are able to derive a result without making any assumptions about the possible types of graphs that could be formed from step 1 [1].
They conjecture that assumptions on the sparsity of the graph are not necessary to show that expansive graphs with a supermajority of either color will converge to a consensus of that color with high probability.

4 Our Results

In Section 5 we show that is not possible to derive a similar result for general λ-expanders with no assumption on the maximum degree, using the same analysis structure as [1]. Our counterexample consists of a λ-expander graph that has a blue supermajority but converges to a red consensus with at least constant probability. The specific graph arrangement of our counterexample occurs with low probability in the construction of the graph, so this proof does not necessarily imply that λ-expanders initialized via Feldman et. al's algorithm converge against their supermajority with constant probability. However, it does demonstrate that such a proof would only be possible if it also reasoned about the types of graphs that could be reasonably constructed. We hope that this guide future researchers towards a correct proof. Section 6 introduces a new class of graphs that converge to a consensus of their supermajority color with high probability: d-regular λ-expanders, with no assumptions on d.

5 Counter-Convergence

We construct an example that is a λ-expander for $\lambda \leq \frac{\delta}{6}$ with a blue supermajority, and show that with at least constant probability, this example converges to a red consensus. One of the open questions in Feldman et. al is whether all λ-expander graphs for appropriate λ converge with high probability given a supermajority of that color. Our counterexample proves that this is not true for all λ-expanders. It certainly may still be true that λ-expander graphs initialized according to the asynchronous algorithm stated in section 2 do converge to a correct consensus with high probability, as our counterexample would be constructed in this process with low probability. However, this counterexample demonstrates that any proof about general λ-expansive graphs must include reasoning about how the algorithm could have colored the graph.

[1] Feldman et. al posits that the second step of their proof relies only on the expansiveness of the graph, not the sparseness. However, in speaking to one of the paper's authors and analyzing the probability bounds provided, we are confident that the second step of their proof does in fact need the max degree to be d.

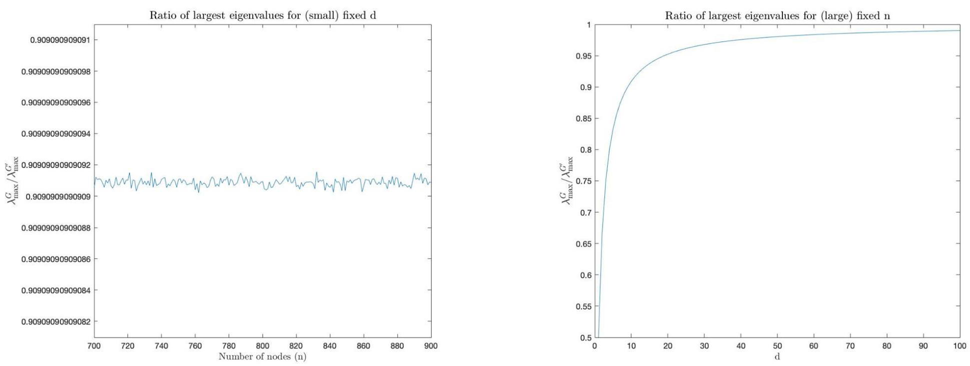

Figure 1

5.1 Counterexample Construction

Let G' be a $(d-1)$-regular random λ-expander for $\lambda \leq \frac{\delta}{6}$. Construct G by adding a node v to G' that is connected to all other nodes. Assume that all nodes in G have declared a color, that $Vol(B) = Vol(R) + \frac{\delta}{2} Vol(G)$, and that v has declared blue. G therefore has a blue super-majority. Our main result is as follows:

Counter-Convergence Theorem: With at least constant probability, G will converge to a red consensus.

At a high level, the proof consists in first showing that if a majority of nodes are red when v is selected, then the graph will converge to a red consensus with high probability, then showing that with at least constant probability, a majority of nodes will be red when v is selected.

5.2 λ-expansiveness of counterexample graph

Conjecture: If G' is a λ-expander, then G is a λ-expander (demonstrated empirically)

The d-regular graph is a λ-expander for a small λ. Thus in order to prove that G is also a λ-expander for small λ, it would be sufficient to demonstrate that G is always more expansive than G'. We show this empirically in Figure 1.

From these graphs we can observe a relationship between the expansiveness of G', $\lambda_{G'}$, and the expansiveness of G, λ_G. As n increases and d is held constant, at least for $d = o(n)$, the relationship between $\lambda_{G'}$ and λ_G seems to remain fairly consistent, with $_G < \lambda_{G'}$. As d increases and n is held constant, λ_G approaches $\lambda_{G'}$ and seems to always be less than or equal to it.

In order to draw a conclusion from these figures, we must assume that the patterns of the relationship between the graphs are consistent. If this is true, then our G is at least as expansive as G' for any d, n. To prove this, let's assume that we have some d, n for which we want to reason about the expansiveness of G. Now we can choose some d', n' that we have empirically demonstrates has $\frac{\lambda_G}{\lambda_{G'}} \leq 1$. If $d' < d$, then we can hold n constant and decrease d' to d. We conjecture based upon our empirical data that this will only make λ_G smaller relative to $\lambda_{G'}$. If $d' > d$, we can hold n' constant and increase d' to d. We conjecture that $\frac{\lambda_G}{\lambda_{G'}}$ value will approach, but not surpass, 1. Now holding d constant we can increase or decrease n' to n. We conjecture that this change to n will not impact $\frac{\lambda_G}{\lambda_{G'}}$ in any significant manner. Thus $\frac{\lambda_G}{\lambda_{G'}} \leq 1$ at our new n and d as well, so it is true at any n and d that G is a λ-expander for $\lambda \leq \frac{\delta}{6}$.

5.3 Conditions for Convergence

Convergence Lemma: If, when v is first selected, $|R| > \frac{n}{2}$, then with probability at least $1 - 4n \cdot e^{-\delta n/48d^2}$, G will reach a red consensus.

Since v is connected to all other nodes and a majority of all other nodes are red, v will declare red. Removing v' to form G' results in a λ-expander for $\lambda \leq \frac{\delta}{6}$ with max degree d. By [1], G' converges to a red consensus with probability at least $1 - 4n \cdot e^{-\delta n/48d^2}$. Since v' declared red and is connected to all other nodes, the probability that G converges to a red consensus is also lower bounded by $1 - 4n \cdot e^{-\delta n/48d^2}$. We continue by proving that with at least constant probability, v's first selection occurs when $|R| > \frac{n}{2}$.

5.4 Arrival Times

Arrival Theorem: With at least constant probability, when v is first selected, a majority of nodes are red.

The proof is broken down into two parts. First, we lower bound the probability that v is selected within n updates. We then lower bound the probability that a majority of nodes are red after n updates, conditioned on not selecting v, which proves the theorem.

We start by calculating the number of blue versus red nodes.

Lemma: For G as constructed above, $|R| \geq |B| + cn$, for constant c that satisfies $d \leq \frac{4-\delta}{\delta+4c}$.

In G, $Vol(B) = Vol(R) + \frac{\delta vol(G)}{2}$.

$$\sum_{\forall b \in B} deg(b) = \sum_{\forall r \in R} deg(r) + \frac{\delta \sum_{\forall x \in N} deg(x)}{2}$$

We know that all nodes except for v have degree d, and that v (which is blue) has degree $n - 1$. Therefore,

$$d(|B| - 1) + n - 1 = d|R| + \frac{\delta(d+1)(n-1)}{2}$$
$$d|R| = d(|B| - 1) + n - 1 - \frac{\delta(d+1)(n-1)}{2}$$
$$d|R| = d(|B| - 1) + (n-1)(1 - \frac{\delta(d+1)}{2})$$

Dividing by d, we get

$$|R| = |B| - 1 + (n-1)\frac{(1 - \frac{\delta(d+1)}{2})}{d}$$

We would like to lower bound the expression $(n-1)\frac{(1-\frac{\delta(d+1)}{2})}{d}$ by cn for some constant c.

$$\begin{aligned}(n-1)\frac{(1-\frac{\delta(d+1)}{2})}{d} - 1 &\geq cn \\ 1 - \frac{(d+1)\delta}{2} &\geq \frac{d(cn+1)}{n-1} \\ \frac{(d+1)\delta}{2} &\leq 1 - \frac{d(cn+1)}{n-1} \\ (d+1)\delta &\leq 2 - 2\frac{d(cn+1)}{n-1} \\ &\leq 4 - 4dc \\ d &\leq \frac{4-\delta}{\delta+4c}\end{aligned}$$

If $c << 1 - \frac{\delta}{2}$, then this bound on d is not too restrictive. δ is very small, and therefore we can find some constant $0 < c < 1$ such that this is true.
Thus,

$$|R| \geq |B| + cn$$

5.5 Arrival Time of Selection of v

Let S be the time that v is first selected.
Arrival Lemma: $\mathbb{P}[S \leq n] \geq 1 - \frac{1}{e}$. S is the first arrival of a Bernoulli process with parameter $p = \frac{1}{n}$. We can conclude this directly from the CDF of the distribution.

$$\begin{aligned}Pr[S \leq n] &\geq 1 - (1-p)^{n-1} \\ &= 1 - \left(1 - \frac{1}{n}\right)^{n-1} \\ &= 1 - \frac{1}{e} \qquad \text{(For large } n\text{)}\end{aligned}$$

5.6 Hitting Time of Half Blue, Half Red

Hitting Time Theorem: After n steps, with probability at least $1 - e^{-O(\log n)}$, there will still be a majority of red nodes.
The outline of the proof is as follows: we model the number of nodes that are red versus blue as a biased, lazy random walk on the interval 0 to n, where position i signifies there are i red nodes, with starting position $\frac{n}{2} + cn$. We bound the probabilities of moving left, right, or staying stationary, then compute the probability that the hitting time for the boundary of $\frac{n}{2}$ is at least n.

5.6.1 Setup

At time t, let $P_{l,t}$ be the probability that we take a step to the left, $P_{r,t}$ be the probability that we take a step to the right, and $P_{s,t}$ be the probability that we stay stationary. If we took a step to the left, then this occurred because a red node was selected that declared blue. If we took a step to the right, then this occurred because a blue node was selected that declared red. Let t' be the hitting time for the boundary of $\frac{n}{2}$, i.e. the time it takes for at least half of the nodes to have declared blue.

Lemma: $\forall t < t'$, the probability that a red node was selected but declared blue is at most $\left(\frac{1}{2}+\frac{1}{2^d}\right)\left(\frac{1+c}{2}\right)$
The proof of this lemma can be broken into two parts: first, we compute the probability that a randomly selected node is red. Then, we compute the probability that a randomly selected node declares blue. We note that these two events are roughly independent for large enough n, which allows us to conclude that the probability that both occur is their product.

Since $|R| \geq |B| + cn$ and $|R| + |B| = n$, the number of red nodes is $n\left(\frac{1+c}{2}\right)$. The probability that a red node $u \neq v$ was selected is then at most $\frac{1+c}{2}$.

Next, we consider when a randomly selected node u will declare blue, and we do so by first calculating the probability that u will declare red. Node u will declare red if a majority of its neighbors are red. Since the graph was random regular graph, the color of its neighbors (excluding the butterfly) are independent (for large enough n). The probability that a given neighbor $j \neq v$ is red is lower-bounded by $\frac{1}{2}$, $\forall t < t'$, and in order for a majority of u's neighbors to be red, there must be at least $\frac{d}{2}+1$ red neighbors, since v is blue. Then, the probability that u will declare red is at least

$$\sum_{i=\frac{d}{2}+1}^{d-1} \binom{d}{i} \left(\frac{1}{2}\right)^d = \frac{1}{2} - \frac{1}{2^d}$$

Since a node must either declare red or blue, the probability that u declares blue is at most $\frac{1}{2}+\frac{1}{2^d}$ Multiplying this probability with the probability that a red node was selected gives us $\left(\frac{1}{2}+\frac{1}{2^d}\right)\left(\frac{1+c}{2}\right)$. □

Corollary: The probability that a blue node $u \neq v$ was selected and declared red is at least $\left(\frac{1}{2}+\frac{1}{2^d}\right)\left(\frac{1+c}{2}\right)$
This follows using similar logic from the above lemma.

5.6.2 Analysis of the Random Walk

We have shown that in the random walk, $\forall t < t'$, $P_{l,t} \leq \left(\frac{1}{2}+\frac{1}{2^d}\right)\left(\frac{1+c}{2}\right)$ and $P_{r,t} \geq \left(\frac{1}{2}-\frac{1}{2^d}\right)\left(\frac{1-c}{2}\right)$. We can simplify our analysis by considering a random walk where the probabilities of moving left or right are static, with $P_{l,t} = \left(\frac{1}{2}+\frac{1}{2^d}\right)\left(\frac{1+c}{2}\right)$ and $P_{r,t} = \left(\frac{1}{2}-\frac{1}{2^d}\right)\left(\frac{1-c}{2}\right)$. Note that if we can lower bound the probability that the equivalent hitting time is at least n in the static random walk, then this is also a lower bound on $\mathbb{P}[t' > n]$ (the hitting time for the non-static random walk), since the bias toward the left is always at least as large in the static random walk. Henceforth we will use all of the notation related to the random walk in reference to the static random walk.

In the static random walk, $P_{l,t} = \left(\frac{1}{2}+\frac{1}{2^d}\right)\left(\frac{1+c}{2}\right)$, $P_{r,t} = \left(\frac{1}{2}-\frac{1}{2^d}\right)\left(\frac{1-c}{2}\right)$, and $P_{s,t} = 1 - P_{l,t} - P_{r,t} = \frac{1}{2} - \frac{c}{2^d}$. We analyze this random walk by using a 3-level decision tree for deciding how to move, i.e. 3 different biased coin flips. The first coin has probability $\frac{1}{2}-\frac{c}{2^d}$ of coming up heads. The second coin has probability c of coming up heads. The third coin has probability $\frac{1}{2}$ of coming up heads (the third coin is unbiased). The random variables representing our position in the original, static random walk, and in our random walk using this decision tree are equivalent, so it is sufficient to analyze the decision tree.

The decision tree is as follows: If the outcome of the first coin flip is heads, we do nothing. If the outcome is tails, we flip the second coin. If the second coin comes up heads, then we move to the left. If the second coin comes up tails, then we flip the third coin. If the third coin comes up heads, we move to the right. If the third coin comes up tails, we move to the left.

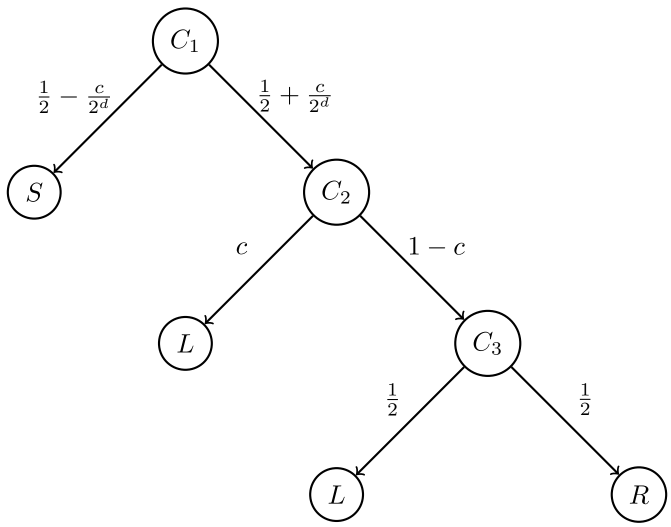

First, we lower bound the number of times that the first coin comes up heads, i.e. a lower bound on the number of times that we stay stationary. Let S_i be the number of times that the first coin comes up heads in i flips. Then, by the Chernoff bound,

$$\mathbb{P}\left[S_n < (1-\alpha)\,\mathbb{E}[S_n]\right] \leq e^{\frac{-\alpha^2\,\mathbb{E}[S_n]}{2}}$$

$$\mathbb{P}\left[S_n < (1-\alpha)n\left(\frac{1}{2}-\frac{c}{2^d}\right)\right] \leq e^{\frac{-\alpha^2 n\left(\frac{1}{2}-\frac{c}{2^d}\right)}{2}}$$

For $\alpha = \frac{\sqrt{n\log n}}{n\left(\frac{1}{2}-\frac{c}{2^d}\right)}$, we have

$$\begin{aligned}
\mathbb{P}\left[S_n < n\left(\frac{1}{2}-\frac{c}{2^d}\right)-\sqrt{n\log n}\right] &\leq e^{\frac{-\log n\left(2^{d-1}-c\right)}{2^{d+2}\left(\frac{1}{2}-\frac{c}{2^d}\right)}} \\
&= e^{-O(\log n)} \qquad \text{(For appropriate } d, c\text{)} \\
\mathbb{P}\left[S_n > n\left(\frac{1}{2}-\frac{c}{2^d}\right)-\sqrt{n\log n}\right] &\geq 1-e^{-O(\log n)}
\end{aligned}$$

Next, we upper bound the number of times that the second coin comes up heads, i.e. an upper bound on the number of times that we move left due to the drift, conditioned on $S_n >$

$n\left(\frac{1}{2}-\frac{c}{2^d}\right)-\sqrt{n\log n}$. Let $q = n - S_n$ be the total number of times we flip the second coin.

$$q \le n - \left(n\left(\frac{1}{2}-\frac{c}{2^d}\right)-\sqrt{n\log n}\right)$$
$$= n\left(\frac{1}{2}+\frac{c}{2^d}\right)+\sqrt{n\log n}$$

Let D_i be the number of times that the second coin comes up heads in i flips. By the Chernoff bound,

$$\mathbb{P}\left[D_q > (1+\epsilon)\,\mathbb{E}[D_q]\right] \le e^{\frac{-\epsilon^2\,\mathbb{E}[D_q]}{3}}$$
$$\mathbb{P}\left[D_q > (1+\epsilon)cq\right] \le e^{\frac{-\epsilon^2 cq}{3}}$$

For $\epsilon = \frac{\sqrt{n\log n}}{cq}$, we have

$$\mathbb{P}\left[D_q > cq - \sqrt{n\log n}\right] \le e^{\frac{-n\log n}{3cq}}$$
$$\mathbb{P}\left[D_q > cn\left(\frac{1}{2}+\frac{c}{2^d}\right)-(1-c)\sqrt{n\log n}\right] \le e^{\frac{-n\log n}{3cq}}$$
$$= e^{-O(\log n)}$$
$$\mathbb{P}\left[D_q \le cn\left(\frac{1}{2}+\frac{c}{2^d}\right)-(1-c)\sqrt{n\log n}\right] \ge 1-e^{-O(\log n)}$$

Finally, we lower bound the probability that an unbiased random walk stays within k steps of its starting position, i.e. a lower bound on the probability that we do not ever get too many more heads than tails in the third coin flip, conditioned on the previous two events occurring. Let Z_i be our position in the unbiased random walk after i steps. Let $w = n - S_n - D_q$ be the total number of steps we flip the third coin. Let h' be the hitting time for the boundary l. By the reflection principle,

$$\mathbb{P}[Z_w \ge l | h' \le w] = \frac{1}{2}$$
$$\frac{\mathbb{P}[Z_w \ge l, h' \le w]}{\mathbb{P}[h' \le w]} = \frac{1}{2} \qquad \text{(By Bayes' theorem)}$$

Note that the event $Z_w \ge l \cap h' \le w$ is equivalent to the event $Z_w \ge l$.

$$\frac{\mathbb{P}[Z_w \ge l]}{\mathbb{P}[h' \le w]} = \frac{1}{2}$$
$$\mathbb{P}[h' > w] = 1 - 2\,\mathbb{P}[Z_w \ge l]$$

By the Chernoff bound,

$$\mathbb{P}[Z_w \ge l] \le e^{\frac{-l^2}{2w}}$$
$$\mathbb{P}[h' > w] \ge 1 - e^{\frac{-l^2}{2w}}$$

For $l = \sqrt{w \log n} = \sqrt{((1-c)n\left(\frac{1}{2} + \frac{c}{2^d}\right) + (2-c)\sqrt{n \log n})(\sqrt{n \log n})}$, we have

$$\begin{aligned}\mathbb{P}[h' > w] &\geq 1 - e^{\frac{-\log n}{2}} \\ &= 1 - e^{-O(\log n)}\end{aligned}$$

We can now lower bound the probability that in the static random walk, we have walked less than cn to the left after n steps. The distance walked to the left is at most $D_q + l$. With probability at least $1 - e^{-O(\log n)}$, this distance is at most

$$\begin{aligned}D_q + l &\leq cn\left(\frac{1}{2} + \frac{c}{2^d}\right) - (1-c)\sqrt{n \log n} + \sqrt{(1-c)n\left(\frac{1}{2} + \frac{c}{2^d}\right) + (2-c)\sqrt{n \log n}} \\ &= \frac{cn}{2} + o(n)\end{aligned}$$

This concludes the proof of the Hitting Time Theorem. Combining the Hitting Time Theorem with the Arrival Lemma, with probability at least $(1 - \frac{1}{e})(1 - e^{-O(\log n)})$, when v is first selected, a majority of nodes are red. This concludes the proof of the Arrival Theorem. Combining the Arrival Theorem with the Convergence Lemma, with at least constant probability, this proves the Counter-Convergence Lemma.

6 Convergence for d-regular graphs for unbounded d

Here, we extend Feldman et. al's result of the high probability convergence of super-majorities to d-regular graphs. Previously this had been demonstrated for small d, but here we provide a proof that holds for all d. .
Lemma: The probability that a λ-expansive, d-regular graph with a supermajority will converge to the color of that supermajority is $\geq 1 - 16ne^{-\delta n(\frac{1}{128})}$.

We will assume without loss of generality that the supermajority is red. Thus $vol(R) \geq vol(B) + \frac{\delta}{2}vol(G)$. Since G is d-regular, $vol(R) = d|R|$ and $vol(B) = d|B|$. Thus

$$\begin{aligned}d|R| &\geq d|B| + \frac{\delta}{2}nd \\ |R| &\geq |B| + \frac{\delta}{2}n\end{aligned}$$

As $|R| + |B| = n$, we can solve for $|B|$:

$$\begin{aligned}n - |B| &\geq |B| + \frac{\delta}{2}n \\ \frac{n(1-\frac{\delta}{2})}{2} &\geq |B|\end{aligned}$$

For any given graph, switching a blue node to a red node will never increase the probability of converging to red. Therefore we can upper-bound $|B|$ to $\frac{n(1-\frac{\delta}{2})}{2}$.
We can view changes to G as a random walk on the number of red nodes from 0 to n, where hitting n is a red consensus. Let us define Z_t as our position at time t. Thus $Z_0 = \frac{n}{2} + \frac{\delta}{2}n$.

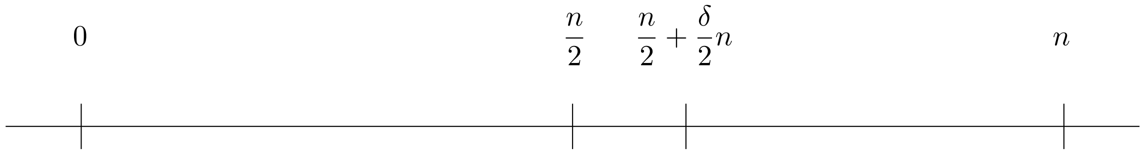

At each step, we are selecting a node at random and changing its color to the majority of its neighbors. The total number of red nodes changes by at most 1 at each step, and therefore the number of red nodes at time t is represented by a 1-bounded lazy random walk.
Let us consider the interval between $\frac{n}{2}+\frac{\delta}{4}n$ and $\frac{n}{2}+\frac{3\delta}{4}n$. $E[Z_{t+1}-Z_t]$ for any Z_t on this interval is lower bounded by $E[Z_{t+1}-Z_t]$ for $Z_t=\frac{n}{2}+\frac{\delta}{4}n$. At this position, $|B|=\frac{n}{2}-\frac{\delta}{8}n$ (come clean this up and make sure the fractions are correct). Corollary 2 in Feldman et. al implies that $|B'|\leq\frac{|B|}{2}$, where B' is the set of nodes that would declare blue if they were selected, and B is the set of nodes that are currently blue. Thus

$$|B'|\leq\frac{n}{4}-\frac{\delta}{8}n$$

At each step a node is selected at random. The probability that this node is blue after being selected is at most $\frac{n}{4}-\frac{\delta}{8}n$. Thus the probability that we select a node and it declares red is at least $1-\frac{n}{4}-\frac{\delta}{8}n$. By definition this is also the probability that we select a node and the majority of nodes connected to is are red. Let us define RB as the set of nodes that are red but would declare blue if selected, and RR, BB and BR correspondingly. From our bound on the number of nodes that would switch to red, we get

$$|B'|=|BB|+|RB|=\frac{n}{4}-\frac{\delta}{8}n$$
$$|BB|=\frac{n}{4}-\frac{\delta}{8}n-|RB|$$

From our bound on the number of initial blue nodes, we get

$$|B|=|BB|+|BR|=\frac{n}{2}-\frac{\delta}{4}n$$
$$|BB|=\frac{n}{2}-\frac{\delta}{4}n-|BR|$$

Combining these, we can get a bound on the relationship between $|BR|$ and $|RB|$.

$$\frac{n}{2}-\frac{\delta}{4}n-|BR|=|BB|=\frac{n}{4}-\frac{\delta}{8}n-|RB|$$
$$\frac{n}{4}-\frac{\delta}{8}n-|BR|=-|RB|$$
$$\frac{n}{4}-\frac{\delta}{8}n=|BR|-|RB|$$

Now we can reason about $E[Z_{t+1}|Z_{t+1} \neq Z_t]$, or the expected movement on any step that has any movement at all.

$$\begin{aligned} E[Z_{t+1}|Z_{t+1} \neq Z_t] &= Z_t + 1 * P(select\ node\ in\ BR) - 1 * P(select\ node\ in\ RB) \\ &= \frac{|BR|}{n} - \frac{|RB|}{n} \\ &= \frac{|BR| - |RB|}{n} \\ &= \frac{1}{4} - \frac{\delta}{8} \end{aligned}$$

We would like to reason about the probability that this random walk hits one threshold before another threshold, and therefore it is equivalent to reason about a random walk ignoring the steps without any movement. In this case, $E[Z_{t+1}] \geq \frac{1}{4} - \frac{\delta}{8} + Z_t$ for all $\frac{n}{2} + \frac{\delta}{4}n \leq Z_{t-1} \leq \frac{n}{2} + \frac{3\delta}{4}n$. We have constructed a 1-bounded, $\frac{1}{4} - \frac{\delta}{8}$-biased random walk.
Now we have fulfilled all the conditions to utilize lemma 3 provided in Feldman et. al, which gives us that the probability that the walk hits $\frac{n}{2} + \frac{1\delta}{4}n$ before hitting $\frac{n}{2} + \frac{3\delta}{4}n$ is at most

$$\frac{2\frac{\delta n}{4}}{\frac{1}{4} - \frac{\delta}{8}} e^{-\frac{(\frac{1}{4} - \frac{\delta}{8})(\frac{\delta n}{4})}{4}} = \frac{2\delta n}{1 - \frac{\delta}{2}} e^{-\delta n(\frac{1}{64} - \frac{\delta}{128})}$$

The graph will hit a red consensus after $\frac{n}{\frac{\delta n}{4}} = \frac{4}{\delta}$ successful steps. We can take the union bound of this to determine an upper bound on the probability that we will ever hit a left boundary before a right boundary between our staring point and n. As we move towards more reds in our graph we can continue to lower bound p by $\frac{1}{4} - \frac{\delta}{8}$.

$$\begin{aligned} P(failure) &\leq \frac{4}{\delta} \frac{2\delta n}{1 - \frac{\delta}{2}} e^{-\delta n(\frac{1}{64} - \frac{\delta}{128})} \\ &= \frac{8n}{1 - \frac{\delta}{2}} e^{-\delta n(\frac{1}{64} - \frac{\delta}{128})} \leq 16ne^{-\delta n(\frac{1}{128})} \end{aligned}$$

This probability of failure approaches 0 for large n, so this concludes the proof. Therefore d-regular graphs for unbounded d with a supermajority of one color converge to that color with high probability.

7 Conclusion and Future Work

We have expanded the class of graphs that, given a red supermajority, converges to a red consensus with high probability. However, we constructed an example that shows that this decoupled analysis can not be used to answer the generalized question, which asks whether λ-expanders with no additional assumptions on the sparsity can start from a supermajority of a given color and converge to a consensus on that color with high probability. We hope that this proof can help guide future researchers in approaching a proof for λ-expanders.

References

[1] Michal Feldman, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. Reaching Consensus via non-Bayesian Asynchronous Learning in Social Networks. *CoRR*, abs/1408.5192, 2014.