Majority Opinion Diffusion in Social Networks: An Adversarial Approach

Ahad N. Zehmakan
Department of Computer Science, ETH Zurich
abdolahad.noori@inf.ethz.ch

Abstract

We introduce and study a novel majority based opinion diffusion model. Consider a graph G, which represents a social network. Assume that initially a subset of nodes, called seed nodes or early adopters, are colored either black or white, which correspond to positive or negative opinion regarding a consumer product or a technological innovation. Then, in each round an uncolored node, which is adjacent to at least one colored node, chooses the most frequent color among its neighbors.

Consider a marketing campaign which advertises a product of poor quality and its ultimate goal is that more than half of the population believe in the quality of the product at the end of the opinion diffusion process. We focus on three types of attackers which can select the seed nodes in a deterministic or random fashion and manipulate almost half of them to adopt a positive opinion toward the product (that is, to choose black color). We say that an attacker succeeds if a majority of nodes are black at the end of the process. Our main purpose is to characterize classes of graphs where an attacker cannot succeed. In particular, we prove that if the maximum degree of the underlying graph is not too large or if it has strong expansion properties, then it is fairly resilient to such attacks.

Furthermore, we prove tight bounds on the stabilization time of the process (that is, the number of rounds it needs to end) in both settings of choosing the seed nodes deterministically and randomly. We also provide several hardness results for some optimization problems regarding stabilization time and choice of seed nodes.

1 Introduction

In real life, we usually have specific perspectives on various topics, such as consumer products, technological innovations, life styles, and political events and by communicating with friends, family, and colleagues, our opinions are influenced. Opinion diffusion and (mis)-information spreading can affect different aspects of our lives such as economy, defense, fashion, even personal affairs. Therefore, there has been a growing interest to understand how opinions form and diffuse because of the existence of social ties among a community's members and how the structure of a social network can influence this process. This would enable us to obtain better predictions of electoral results, control the effect of marketing and political campaigns, and in general advance our knowledge of the cognitive processes behind social influence.

The study of opinion diffusion and the evolution of social dynamics on networks has attracted the attention of researchers from a vast spectrum of disciplines such as economics (Bharathi, Kempe, and Salek 2007), epidemiology (Pastor-Satorras and Vespignani 2001), social psychology (Yin et al. 2019), statistical physics (Gärtner and Zehmakan 2020), and political sciences (N. Zehmakan and Galam 2020). It has also gained significant popularity in theoretical computer science, especially in the quickly growing literature focusing on the interface between social choice and social networks, cf. (Bredereck and Elkind 2017) and (Auletta, Ferraioli, and Greco 2018).

From a theoretical viewpoint, it is natural to introduce and study mathematical models which mimic different opinion dynamics. Of course in the real world, they are too complex to be explained in purely mathematical terms. However, the main idea is to comprehend their general principles and make crude approximations at discovering certain essential aspects of them which are otherwise totally hidden by the complexity of the full phenomenon.

In these models, it is usually assumed that we have a graph G and initially some nodes are colored, say black or white. Then, in each round a group of nodes get colored or update their color based on a predefined rule. Graph G is meant to represent a social network, where each agent is modeled as a node and edges indicate relations between them, e.g., friendship, common interests, advice, or various forms of interactions. Furthermore, black and white stand for the opinion of an agent regarding an innovation or a political party, etc.

In the plethora of opinion diffusion models, threshold-based ones are certainly the best known, cf. (Kempe, Kleinberg, and Tardos 2003), (Apt and Markakis 2014), and (Zehmakan 2020). There, nodes (i.e., agents) adopt a color (i.e., opinion) if it is shared by a certain number or fraction of their connections. Particularly, the majority-based models, where each node chooses the most frequent color among its neighbors, have received a substantial amount of attention, cf. (Chistikov et al. 2020). This imitating behavior can be explained in several ways: an agent that sees a majority agreeing on an opinion might think that her neighbors have access to some information unknown to her and hence they have made the better choice; also agents can directly benefit from adopting the same behavior as their friends (e.g., prices going down).

Nowadays, the identification of factors leading to a successful innovation in a given market is a question of considerable practical importance. Marketing campaigns routinely use online social networks to attempt to sway people's opinions in their favor, for instance by targeting segments of agents with free sample of their products or misleading information. Consequently, the study of control and manipulation of collective decision-making has gained increasing popularity in mechanism design, algorithmic game theory, and computational social choice, cf. (Bredereck and Elkind 2017). Especially, majority-based models have been postulated as one potential explanation for the success or failure of collective action and the diffusion of innovations. For example for different majority-based opinion diffusion models, a substantial amount of attention has been devoted to the study of characterizing graph structures for which the dominant color at the end of the process is the same as the dominant color in the initial coloring, cf. (Auletta et al. 2015).

Our Contribution

We introduce a novel majority-based opinion diffusion model and consider three different types of attackers whose goal is to engineer the output of the diffusion process. Our central problem is to characterize classes of graphs for which an attacker fails to reach its goal.

Our model. Consider a graph G. Assume that initially a subset of nodes, which are called the *seed nodes*, are colored black or white and the rest of nodes are uncolored. Then, in discrete-time rounds each node, which is uncolored and is adjacent to at least one colored node, chooses the most frequent color among its colored neighbors. In case of a tie, it chooses black color with probability (w.p.) $1/2$ independently and white otherwise. (We should mention that the results provided in the present paper actually hold for any choice of the tie-breaking rule.)

The seed nodes correspond to early adopters, who are the first customers to adopt a new product or technology before the rest of the population does and usually constitute 10-20% of the population. They are often called lighthouse customers because they serve as a beacon of light for the rest of the population to follow, which will take the technology or product mainstream. The term early adopters comes from the technology adoption curve, which was popularized by the book Diffusion of Innovations (Rogers 1962).

The seed nodes and their color might be chosen randomly or deterministically. We take an adversarial perspective and focus on the following three types of attackers, where we assume that $\alpha, \epsilon \in (0, 1/2)$ are some arbitrary constants and n is the number of nodes in the underlying graph.

Definition 1.1 (strong attacker) *An (α,ϵ)-strong attacker selects a seed set of size αn and color $(1/2 + \epsilon)$ fraction of its nodes white and the rest black.*

Definition 1.2 (moderate attacker) *An (α,ϵ)-moderate attacker selects a seed set of size αn and color each seed node white, independently, w.p. $(1/2 + \epsilon)$ and black otherwise.*

Definition 1.3 (weak attacker) *An (α,ϵ)-weak attacker selects each node to be a seed node, independently, w.p. α and then color each seed node white, independently, w.p. $(1/2 + \epsilon)$ and black otherwise.*

Assume that the process runs for t rounds. Then, the main goal of an attacker is to maximize the ratio of the number of black nodes to white ones at the end of the t-th round. We say that the attacker *wins* if this ratio is at least one half. Our goal is to bound the probability that an attacker wins.

To better understand the aforementioned attacker models, you might think of an attacker as a marketer which desires to advertise an innovation or a product of poor quality and it can select the set of early adopters (i.e., seed nodes) in a random or deterministic fashion. Then, it manages to convince almost half of them to adopt a positive opinion about the product (i.e., choose black color). More precisely, a strong attacker can choose $(1/2-\epsilon)$ fraction of the seed nodes to be black and a moderate/weak attacker colors each seed node black w.p. $(1/2 - \epsilon)$. The attacker's ultimate goal is that a majority of nodes are colored black after some number of rounds, even though initially white is the dominant color.

As an extreme example, consider a star graph S_n, which includes $n-1$ leaves and an internal node of degree $n-1$. An (α,ϵ)-strong attacker can choose the internal node and $\alpha n-1$ of the leaves to be the seed nodes and color the internal node black. Then after one round of the process, all the remaining $(1 - \alpha)n$ nodes will be colored black. Thus, the attacker wins. Even for an (α,ϵ)-weak attacker (which essentially has no real selection power and only runs a random procedure), the internal node is selected w.p. α and will be colored black w.p. $(1/2 - \epsilon)$. Hence, the attacker wins w.p. $(1/2 - \epsilon)\alpha$.

Our main purpose is to characterize classes of graphs which are *resilient* to such attacks (i.e., an attacker cannot win) where we assume that the attacker has full knowledge of the graph structure. We should emphasize that for a moderate or weak attacker since each seed node is white independently w.p. $(1/2+\epsilon)$, the probability that the attacker wins is at most $1/2$. However, we are interested in graph structures where the probability of winning is extremely small.

Strong attacker. An (α,ϵ)-strong attacker is quite powerful. However, there are graphs, such as a complete graph, where it cannot win. We prove that if the graph is regular and has strong expansion properties, then a strong attacker fails. More precisely we prove that for a regular graph G if $\sigma(G) \leq \epsilon\sqrt{\alpha(1-\alpha)}$, then an ($\alpha$,$\epsilon$)-strong attacker cannot win, where $\sigma(G)$ is the second-largest absolute eigenvalue of the normalized adjacency matrix of G. (Please see Section 2 for a more detailed definition of $\sigma(G)$ and its relation to expansion.) We argue that expansion and regularity are not only sufficient conditions, but also somewhat necessary for a graph to be resilient to a strong attacker.

In a nutshell, since a strong attacker has the power to choose the seed nodes and their color, for a graph to be resilient, the "influencing power" should be distributed uniformly among all nodes. In particular, the number of edges among each two node sets must be proportional to their size and this is what basically regularity and expansion provide.

Moderate attacker. Roughly speaking, if there is a small set of nodes of size s with significant influencing power, then a weak attacker can select them to be in the seed set. These nodes will be colored black w.p. $(1/2 - \epsilon)^s$, which is non-negligible for small s, and this can result in the attacker winning. A natural way to avoid such scenarios is to bound the maximum degree of the underlying graph. We show that if the maximum degree is not "too large", then a moderate attacker fails asymptotically almost surely. (We say an event happens asymptotically almost surely (a.a.s.) if it occurs with a probability tending to 1 while we let n go to infinity.) More precisely, we prove that for a graph G and an (α,ϵ)-moderate attacker, if $\Delta \le (Cn/\log(1/\mu))^{\frac{1}{2t}}$ for some small constant $C_{\epsilon,\alpha} > 0$, the attacker cannot win in t rounds w.p. at least $1 - \mu$, where Δ denotes the maximum degree. Furthermore, we argue the tightness of this statement.

Weak attacker. As we discussed, an (α,ϵ)-weak attacker wins on a star graph S_n w.p. at least $(1/2 - \epsilon)\alpha$. Loosely speaking, the graphs of this type, where there is a small subset of nodes with significant influencing power and a large set of nodes with very limited power, are problematic. Our goal is basically to show that graphs which do not fall under the umbrella of this type of structure are resilient to a weak attacker. We prove that if the maximum degree in a graph G is sufficiently small or if a majority of nodes have rather large degrees, then a weak attacker will fail. More accurately, we show that for a graph G and an (α,ϵ)-weak attacker, if $\Delta \le Cn/((\log n)^{C'} \log(1/\mu))$ for some suitable constants $C_{\alpha,\epsilon}, C'_{\alpha,\epsilon} > 0$ or if half of nodes are of degree at least $C'' \log(1/\mu)$ for a sufficiently large constant $C''_{\alpha,\epsilon}$, then the attacker fails w.p. at least $1-\mu$. Note that the bound on Δ does not depend on t unlike the case of a moderate attacker.

Stabilization time. For a connected graph G, if the seed set is non-empty, then all nodes will be colored eventually. The number of rounds the process needs to color all nodes is called the *stabilization time* of the process. As we discuss in Section 5, it is straightforward to prove that the stabilization time is upper-bounded by the diameter of G. This is, in particular, true for the setting of a strong/moderate attacker. However, what if the seed nodes are chosen at random (for example, in case of a weak attacker)? We prove that if each node is in the seed set independently w.p. α, then the stabilization time is a.a.s. lower-bounded by $\Omega(\log_\Delta \log n^{\frac{1}{\alpha}})$ and upper-bounded by $\mathcal{O}((1/\alpha\delta) \log n)$, where δ denotes the minimum degree in G. Furthermore, we argue that these bounds are both tight, up to a constant factor.

Hardness results. We provide some hardness results for a decision problem on the stabilization time and an optimization problem on the choice of seed set and its coloring. Assume that we are given a graph G and some integers t and s. We prove that the problem of determining whether there is a seed set of size s for which the process takes t rounds to end is NP-complete. Furthermore, suppose we are given a graph G and some integers b, w, and t as the input and our goal is to find the minimum (expected) number of white nodes after t rounds if there are b black and w white nodes initially. We prove the best possible approximation factor for this problem is larger than $n^{1-\zeta}$ for any $\zeta > 0$ unless P=NP.

Related Work

Numerous opinion diffusion models have been introduced to investigate how a group of agents modify their opinions under the influence of other agents. It is usually assumed that for a graph G, which represents a social network, initially some nodes are colored black (positive) or white (negative). Then, nodes update their color based on some predefined rule. Among these models, the threshold and majority model are perhaps the closest to ours. In both of them, initially each node is black or white. In the *majority model* (Peleg 1997), in each round all nodes simultaneously update their color to the most frequent color in their neighborhood (and no update in case of a tie). In the *threshold model* (Kempe, Kleinberg, and Tardos 2003), each node v has a threshold value t_v and it becomes black as soon as it has at least t_v black neighbors.

We should emphasize that different variants of these models have been considered; for example, asynchronous updating rule (Auletta et al. 2015), various tie-breaking rules (Peleg 1998), random threshold values (Kempe, Kleinberg, and Tardos 2003), and with a bias toward one of the colors (Anagnostopoulos et al. 2020). Even more complex models such as the ones considered in (Ferraioli and Ventre 2017) and (Auletta, Fanelli, and Ferraioli 2019), which follow an averaging-based updating rule, or the models in (Brill et al. 2016), (Meir et al. 2010), (Auletta, Ferraioli, and Greco 2020), and (Faliszewski et al. 2018) can be seen as extensions of the majority model.

Consider a marketer which advertises a new product. Assume that it can convince a subset of agents to adopt a positive opinion about its product, e.g., by giving them free samples of the product, and it aims to trigger a large cascade of further adoptions. Which agents should it target? A node set A is a *target set* if black color eventually takes over all (or half) of nodes once A is fully black. (This is sometimes also known as dynamic monopoly or percolating set.) For both the majority and threshold model, the minimum size of a target set has been extensively studied on various classes of graphs such as lattice (Balister et al. 2010), Erdős-Rényi random graph (Schoenebeck and Yu 2018), random regular graphs (Gärtner and Zehmakan 2018), power-law random graphs (Amini and Fountoulakis 2014), and expander graphs (Mossel, Neeman, and Tamuz 2014). Furthermore, (Berger 2001) proved that there exist arbitrarily large graphs which have target sets of constant size under the majority model and (Auletta, Ferraioli, and Greco 2018) showed that every n-node graph has a target set of size at most $n/2$ under the asynchronous variant.

Furthermore, it is known that the problem of finding the minimum size of a target set for a given graph G is NP-hard for different variants of these models. For the majority model, (Mishra, Radhakrishnan, and Sivasubramanian 2002) proved that this problem cannot be approximated within a factor of $\log \Delta \log\log \Delta$, unless P=NP; however, there is a polynomial-time $\log \Delta$-approximation algorithm. For the threshold model, (Chen 2009) proved that the problem is NP-hard even when all thresholds t_v are 1 or 2 and G is a bounded-degree bipartite graph. On the other hand, the problem is traceable for special classes of graphs such as trees. We study a similar optimization problem in our setting and provide inapproximability results.

Many other adversarial scenarios where an attacker aims to engineer the output of an opinion diffusion process have been investigated. For example, assume that an attacker can manipulate almost half of initially colored nodes to choose black color in a random or deterministic manner, similar to our attacker models. Which graph structures are resilient to such attacks, i.e., at the end of the process the white color is still the dominant color? For the majority model, (Mustafa and Pekeč 2001) proved that a graph is resilient if it is a clique or very close to a clique. (Auletta et al. 2015) provided similar results for the asynchronous setting. For our majority-based model, we prove that expansion, regularity, and maximum degree are chiefly responsible for resilience.

Other kinds of attackers with various manipulation powers, such as adding/deleting edges or changing the order of updates, have been considered, cf. (Bredereck and Elkind 2017), (Corò, D'Angelo, and Velaj 2019), and (Wilder and Vorobeychik 2017). Moreover, the complexity of finding an optimal strategy for different types of attackers has been studied, cf. (Grandi, Stewart, and Turrini 2018), (Corò et al. 2019), and (Borodin, Filmus, and Oren 2010).

In the majority model, since the updating rule is deterministic, the process eventually reaches a cycle of colorings. The length of this cycle and the number of rounds the process needs to reach it are called the periodicity and stabilization time of the process, respectively. (Goles and Olivos 1980) proved that periodicity is always one or two. Recently, (Chistikov et al. 2020) and (Zhuang et al. 2020) showed that it is PSPACE-complete to decide whether the periodicity is one or not for a given coloring of a *directed* graph. Regarding the stabilization time, (Fogelman, Goles, and Weisbuch 1983) showed that it is bounded by $\mathcal{O}\left(n^2\right)$. (Frischknecht, Keller, and Wattenhofer 2013) proved that this bound is tight, up to some poly-logarithmic factor. (Kaaser, Mallmann-Trenn, and Natale 2016) proved that the problem of determining whether there exists a coloring for which the process takes at least t rounds is NP-complete. As mentioned, we provide tight bounds and hardness results on the stabilization time of our diffusion process.

Preliminaries

Graph definitions. Let $G = (V, E)$ be an n-node graph. For two nodes $v, u \in V$, we define $d(v, u)$ to be the length of a shortest path between v, u in terms of the number of edges, which is called the *distance* between v and u (for a node v, we define $d(v, v) := 0$). For $t \in \mathbb{N}_0$, we let $N_t(v) := \{u \in V : d(v, u) = t\}$ denote the set of nodes whose distance from v is exactly t. In particular, $N_0(v) = v$ and $N_1(v)$ is the set of v's *neighbors*. Furthermore, we define $\hat{N}_t(v) := \cup_{i=0}^{t} N_i(v)$ to be the *t-neighborhood* of v. Analogously, for a node set $S \subseteq V$ we have $N_t(S) := \{u \in V : d(S, u) = t\}$ and $\hat{N}_t(S) := \cup_{i=0}^{t} N_i(S)$, where $d(S, u) := \min_{v \in S} d(v, u)$. We let $\deg(v) := |N_1(v)|$ denote the *degree* of v and define $\deg_S(v) := |N_1(v) \cap S|$. We also define $\Delta(G)$ and $\delta(G)$ to be respectively the maximum and minimum degree in graph G. For two node sets S and S', we define $e(S, S') := |\{(v, u) \in S \times S' : \{v, u\} \in E\}|$ where $S \times S'$ is the Cartesian product of S and S'.

Model definitions. In our model, we will denote by R_t the set of nodes which are colored in the t-th round for $t \in \mathbb{N}_0$. In other words, R_0 is equal to the seed set and $R_t := N_t(R_0)$ for $t \in \mathbb{N}$. We also define $\hat{R}_t = \cup_{i=0}^{t} R_i$. Moreover, we let B_t (analogously W_t) denote the set of black (resp. white) nodes in R_t. Similarly, $\hat{B}_t$ and $\hat{W}_t$ denote the set of black and white nodes in $\hat{R}_t$, respectively. We also define $r_t := |R_t|$, $\hat{r}_t = |\hat{R}_t|$, $b_t := |B_t|$, $\hat{b}_t = |\hat{B}_t|$, $w_t = |W_t|$, and $\hat{w}_t = |\hat{W}_t|$. Note that these are random variables even when we choose the seed nodes and their color deterministically since we break a tie at random.

Some inequalities. Now, we present some standard probabilistic inequalities, cf. (Dubhashi and Panconesi 2009), which we utilize several times later.

Theorem 1.4 (Chernoff bound) *Suppose $x_1, \cdots, x_k$ are independent Bernoulli random variables and let X denote their sum, then for any $0 \leq \eta \leq 1$*

- $\Pr[(1 + \eta)\, \mathbb{E}[X] \leq X] \leq \exp\left(-\frac{\eta^2 \mathbb{E}[X]}{3}\right)$
- $\Pr[X \leq (1 - \eta)\, \mathbb{E}[X]] \leq \exp\left(-\frac{\eta^2 \mathbb{E}[X]}{2}\right)$.

If we are given k discrete probability spaces $(\Omega_i, \Pr_i)$ for $1 \leq i \leq k$, then their product is defined to be the probability space over the ground set $\Omega := \Omega_1 \times \Omega_2 \times \cdots \times \Omega_k$ with the probability function $\Pr[(\omega_1, \ldots, \omega_k)] := \prod_{i=1}^{k} \Pr_i[\omega_i]$, where $\omega_i \in \Omega_i$. Let $(\Omega, \Pr)$ be the product of k discrete probability spaces, and let $X : \Omega \to \mathbb{R}$ be a random variable over Ω. We say that the *effect* of the i-th coordinate is at most c_i if for all $\omega, \omega' \in \Omega$ which differ only in the i-th coordinate we have $|X(\omega) - X(\omega')| \leq c_i$. Azuma's inequality states that X is sharply concentrated around its expectation if the effect of the individual coordinates is not too big.

Theorem 1.5 (Azuma's inequality) *Let $(\Omega, \Pr)$ be the product of k discrete probability spaces $(\Omega_i, \Pr_i)$ for $1 \leq i \leq k$, and let $X : \Omega \to \mathbb{R}$ be a random variable with the property that the effect of the i-th coordinate is at most c_i. Then,* $\Pr[X \leq \mathbb{E}[X] - a] \leq \exp(-a^2/(2\sum_{i=1}^{k} c_i^2))$.

Moreover, we sometimes use the basic inequalities $1 - z \leq \exp(-z)$ for any z and $4^{-z} \leq 1 - z$ for any $0 < z < 1/2$.

Assumptions. All logarithms are to base e, otherwise we point out explicitly. Furthermore, we let n (the number of nodes) tend to infinity and $\alpha, \epsilon \in (0, 1/2)$ are constants. We assume that the error probability μ is larger than $1/\sqrt{n}$. (Actually, most of our proofs also work when $\mu > 1/f(n)$ for any function $f(n)$ sub-exponential in n.)

2 Strong Attacker

In this section, our goal is to prove that a regular expander graph is resilient to a strong attacker. Roughly speaking, one says a graph has strong expansion properties if it is highly connected. There exist different parameters to measure the expansion of a graph. We consider an algebraic characterization of expansion. However, since the relation between other measures, such as vertex and edge expansion, and ours is well-understood, cf. (Hoory, Linial, and Wigderson 2006), our results can be immediately rephrased.

Let $\sigma(G)$ be the second-largest absolute eigenvalue of the normalized adjacency matrix of graph G. There is a rich literature about the relation between the value of σ and the expansion of graph G. However for our purpose here, all one needs to know is that G has stronger expansion properties when σ is smaller. Now, we present the expander mixing lemma, cf. (Hoory, Linial, and Wigderson 2006), which basically states that the number of edges between any two node sets is almost completely determined by their cardinality if the value of σ is small.

Lemma 2.1 (Expander mixing lemma) *For a d-regular graph $G=(V,E)$ and $S,S'\subset V$*

$$\left|e\left(S,S'\right)-\frac{|S|\,|S'|d}{n}\right|\leq\sigma d\sqrt{|S|\,|S'|}\,.$$

Theorem 2.2 *For a d-regular graph $G=(V,E)$ and an (α,ϵ)-strong attacker, if $\sigma\leq\epsilon\sqrt{\alpha(1-\alpha)}$, then the attacker cannot win for any number of rounds.*

Proof. We know that $w_0=(1/2+\epsilon)\alpha n>\alpha n/2$. We prove that $w_1\geq(1-\alpha)n/2$, which yields $w_0+w_1>n/2$. Thus, for any number of rounds, white is the dominant color.

Let $A:=V\setminus(R_0\cup W_1)$ be the set of non-seed nodes which are black or uncolored after one round. It suffices to prove that $a:=|A|\leq(1-\alpha)n/2$ (since this implies that $w_1\geq(1-\alpha)n/2$). For each node $v\in A$, $\deg_{W_0}(v)\leq\deg_{B_0}(v)$ (because otherwise it would become white after one round). Hence, $e(A,W_0)\leq e(A,B_0)$. By applying Lemma 2.1 to both sides of this inequality, we get

$$\frac{aw_0d}{n}-\sigma d\sqrt{aw_0}\leq\frac{ab_0d}{n}+\sigma d\sqrt{ab_0}.$$

Dividing by $\sqrt{a}$ and re-arranging the terms give us

$$\sqrt{a}(w_0-b_0)\leq\sigma n(\sqrt{w_0}+\sqrt{b_0}).$$

Since $\sqrt{w_0}+\sqrt{b_0}=\sqrt{(1/2+\epsilon)\alpha n}+\sqrt{(1/2-\epsilon)\alpha n}\leq\sqrt{2\alpha n}$, $w_0-b_0=2\alpha\epsilon n$, and $\sigma\leq\epsilon\sqrt{\alpha(1-\alpha)}$, we get

$$\sqrt{a}\leq\frac{\epsilon n\sqrt{\alpha(1-\alpha)}\sqrt{2\alpha n}}{2\alpha\epsilon n}\Rightarrow a\leq(1-\alpha)n/2.\quad\square$$

Random regular graphs. The random d-regular graph $\mathcal{G}_{n,d}$ is the random graph with a uniform distribution over all d-regular graphs on n nodes. It is proven by (Friedman 2003) that $\sigma(\mathcal{G}_{n,d})\leq 2/\sqrt{d}$ for $d\geq 3$ a.a.s. Putting this statement in parallel with Theorem 2.2 implies that if $d\geq 4/(\alpha(1-\alpha)\epsilon^2)$, then $\mathcal{G}_{n,d}$ is resilient to an (α,ϵ)-attacker a.a.s.

Irregular graphs. So far, we limited ourselves to regular graphs. However, our result can be generalized to capture irregular graphs by applying basically the same proof ideas. All we need to do is to apply a more general variant of Lemma 2.1, cf. (Hoory, Linial, and Wigderson 2006), and replace d with δ or Δ, according to the case. Then, we can conclude that for a graph $G=(V,E)$ and an (α,ϵ)-strong attacker, if $\sigma\leq((1+\gamma)\epsilon/2-(1-\gamma)/4)\sqrt{\alpha(1-\alpha)}$, then the attacker cannot win, where $\gamma:=\delta/\Delta$. (This is equivalent to the statement of Theorem 2.2 for $\gamma=1$.)

Erdős-Rényi random graph. In the Erdős-Rényi random graph $\mathcal{G}_{n,p}$ each edge is added independently w.p. p on a node set of size n. Let K be a sufficiently large constant. It is proven by (Le, Levina, and Vershynin 2017) that if $p\geq K\log n/n$, then $\sigma(\mathcal{G}_{n,p})=\mathcal{O}(1/\sqrt{np})$ a.a.s. (recall that $\log n/n$ is the connectivity threshold). Furthermore, it is well known, cf. (Dubhashi and Panconesi 2009), that $\gamma\geq(1-\epsilon)$ for any $\epsilon>0$ a.a.s. if $p\geq K\log n/n$. Combining the last two statements and our proposition about irregular graphs implies that $\mathcal{G}_{n,p}$ for $p\geq K\log n/n$ is a.a.s. resilient to an (α,ϵ)-strong attacker.

Tightness. We believe that regularity and expansion are not only sufficient conditions for a graph to be resilient to a strong attacker, but also somehow necessary. (See the full version of the paper for more details.)

3 Moderate Attacker

In this section, our main purpose is to prove Theorem 3.2. To do so, let us first provide Lemma 3.1.

Lemma 3.1 *For a graph $G=(V,E)$ and an (α,ϵ)-moderate attacker, $(1/2+\epsilon)\alpha n-\alpha\epsilon n/2\leq w_0\leq(1/2+\epsilon)\alpha n+\alpha\epsilon n/2$ w.p. $1-\exp(-\Theta(\epsilon^2\alpha n))$.*

Proof. Label the seed nodes arbitrarily from v_1 to $v_{\alpha n}$. We define Bernoulli random variable x_i to be 1 if and only if v_i is colored white. Then, $w_0=\sum_{i=1}^{\alpha n}x_i$ and $\mathbb{E}[w_0]=(1/2+\epsilon)\alpha n$. Since x_is are independent, applying the Chernoff bound (Theorem 1.4) for $\eta=\epsilon/2$ yields our claim. $\square$

Theorem 3.2 *For a graph $G=(V,E)$ and an (α,ϵ)-moderate attacker, if $\Delta\leq(Cn/\log(1/\mu))^{\frac{1}{2t}}$ for some sufficiently small constant $C_{\alpha,\epsilon}>0$, then the attacker cannot win in t rounds w.p. at least $1-\mu$.*

Proof. Let $r=\hat{r}_t-r_0$ be the number of nodes which are colored during rounds 1 to t. We define $X:=\hat{w}_t-w_0$ to be the number of nodes which are colored white among these nodes. We prove that $X>r/2-\alpha\epsilon n/2$ w.p. at least $1-\mu/2$. Furthermore based on Lemma 3.1, $w_0\geq r_0/2+\alpha\epsilon n/2$ w.p. $1-\exp(-\Theta(\epsilon^2\alpha n))\geq 1-1/(2\sqrt{n})\geq 1-\mu/2$ (where we used our assumptions that α,ϵ are constant and $\mu\geq 1/\sqrt{n}$). Therefore, we have $\hat{w}_t=X+w_0>(r/2-\alpha\epsilon n/2)+(r_0/2+\alpha\epsilon n/2)=(r+r_0)/2=\hat{r}_t/2$ w.p. at least $1-\mu$, which implies that the attacker does not win in t rounds.

It remains to prove that $X>r/2-\alpha\epsilon n/2$ w.p. at least $1-\mu/2$. Let us label all nodes in R_0 arbitrarily from v_1 to $v_{\alpha n}$. Corresponding to each node v_i for $1\leq i\leq\alpha n$, we define the probability space $(\Omega_i,\Pr_i)$. Then, $X:\Omega\rightarrow\mathbb{R}$ is defined over Ω, where $(\Omega,\Pr)$ is the product of discrete probability spaces $(\Omega_i,\Pr_i)$. Recall that we say the effect of the i-th coordinate is at most c_i if for all $\omega,\omega'\in\Omega$ which differ only in the i-th coordinate we have $|X(\omega)-X(\omega')|\leq c_i$. Obviously, if we change the color of node v_i, then the color of at most $|\hat{N}_t(v_i)|\leq 2\Delta^t$ of the nodes, which are colored in the first t rounds, will be affected. Therefore, $c_i\leq 2\Delta^t$, which implies that $\sum_{i=1}^{\alpha n}c_i^2\leq 4\alpha n\Delta^{2t}$. Furthermore, $\mathbb{E}[X]\geq r/2$ since each node in $V\setminus R_0$ is white w.p. at least $1/2$ (actually, one can get the stronger lower bound of $(1/2+\epsilon)$ with a coupling argument, which is not needed here). Now, applying Azuma's inequality (Theorem 1.5) yields

$$\Pr[X \le \frac{r}{2} - \frac{\alpha\epsilon n}{2}] \le \exp(-\frac{\alpha^2\epsilon^2 n^2}{8\sum_{i=1}^{\alpha n} c_i^2}) \le \exp(-\frac{\alpha\epsilon^2 n}{32\Delta^{2t}}).$$

For a suitable choice of $C_{\alpha,\epsilon}$, we have $\Delta \le (\frac{\alpha\epsilon^2 n}{32\log(2/\mu)})^{\frac{1}{2t}}$. Thus, the above probability is at most $\mu/2$. □

Tightness. Now, we provide Proposition 3.3 which asserts that if we could replace the exponent $1/2t$ with $1/t$ in Theorem 3.2, then our bound would be tight, up to some constant factor. We believe that this is actually doable with some case distinctions and more careful calculations; however, we only prove such statement for $t = 1$, in Proposition 3.4. Please see the full version of the paper for the proof of Propositions 3.3 and 3.4.

Proposition 3.3 *For any $\epsilon, \alpha, \mu > 0$ and $t \in \mathbb{N}$, there exists a graph G with $\Delta = (C'n/\log(1/\mu))^{\frac{1}{t}}$ for some constant $C'_{\alpha,\epsilon}$ such that an (α,ϵ)-moderate attacker can win in t rounds w.p. larger than μ.*

Proposition 3.4 *For a graph $G = (V, E)$ and an (α,ϵ)-moderate attacker, if $\Delta \le C''n/\log^2(1/\mu)$ for some sufficiently small constant $C''_{\alpha,\epsilon} > 0$, then the attacker cannot win in one round w.p. at least $1 - \mu$.*

4 Weak Attacker

Theorem 4.1 states that if the maximum degree of a graph is sufficiently smaller than n or if the degree of at least half of the nodes is pretty large, then it is resilient to a weak attacker. In other words, a weak attacker can be successful only on graphs with some nodes of very high degree and a large set of nodes with small degree such as a star graph.

Theorem 4.1 *For a graph $G = (V, E)$ and an (α,ϵ)- weak attacker, if $\Delta \le Cn/((\log n)^{C'} \log(4/\mu))$ for some suitable constants $C_{\alpha,\epsilon}, C'_{\alpha,\epsilon} > 0$ or if half of the nodes are of degree at least $d_1^* := (6/\alpha\epsilon^2)\log(12/\alpha\epsilon\mu)$, then for any $t \in \mathbb{N}$ the attacker cannot win in t rounds w.p. at least $1 - \mu$.*

We observe that the bound on Δ in Theorem 4.1 does not depend on the number of rounds t, unlike the case of a moderate attacker. Furthermore for a moderate attacker, we cannot prove that if the degree of half of the nodes is larger than a certain degree threshold, then the attacker fails, except if the threshold is very large. Assume that graph G is the union of a clique of size $\alpha n - 1$ and a clique of size $(1-\alpha)n + 1$. Suppose that an (α,ϵ)-moderate attacker chooses all $\alpha n - 1$ nodes in the first clique and a node v in the second one to be the seed nodes. If v is colored black, then the attacker wins and this happens w.p. $(1/2 - \epsilon)$. Note that all nodes in G are of degree at least $\alpha n - 2$.

Proof sketch. The complete proof is fairly long and is given in the full version of the paper. Here, we deliver some key ideas which the proof is built on.

Since each of a node's neighbors is initially black independently w.p. $(1/2 - \epsilon)\alpha$ and white w.p. $(1/2 + \epsilon)\alpha$, the dominant color in the neighborhood of a high-degree node is very likely to be white. This lets us prove that if half of the nodes are of degree at least d_1^*, then after one round more than $n/2$ nodes are white w.p. at least $1-\mu$. In that case, the attacker cannot win after any number of rounds.

For the case of $\Delta \le Cn/((\log n)^{C'} \log(4/\mu))$, we first show that $w_0 \ge b_0 + \alpha\epsilon n$ w.p. at least $1 - \mu/4$, similar to the proof of Lemma 3.1. Furthermore, one can prove that at most $\alpha\epsilon n/2$ nodes are colored after the t_1^*-th round w.p. at least $1 - \mu/4$, for $t_1^* := (2/\alpha)\log(4/\alpha\epsilon)$. In the worst case scenario, all these nodes are colored black, but this is not an issue since we already have $\alpha\epsilon n$ extra white nodes in R_0. Moreover building on the argument from the previous paragraph, we can prove that all non-seed nodes whose degree is larger than $d_2^* := 8\log n/(\alpha\epsilon^2)$ are colored white after one round w.p. at least $1 - \mu/4$. Therefore, it only remains to show that at least half of the nodes whose degree is smaller than d_2^* and are in $\cup_{i=1}^{t_1^*} R_i$ will be colored white w.p. at least $1 - \mu/4$. For this, we rely on Azuma's inequality, similar to the proof of Theorem 3.2. However, we can prove a much tighter bound on $\sum_{i=1}^{r_0} c_i^2$ by a smarter counting argument and using the fact that we only need to focus on "low-degree" nodes that are colored up to the t_1^*-th round. □

Tightness. The bound $(6/\alpha\epsilon^2)\log(12/\alpha\epsilon\mu)$ in Theorem 4.1 is tight in terms of μ. However, we believe the dependency on α and ϵ is not best possible. The upper bound on Δ is also tight in μ and the dependency on n is optimal, up to the poly-logarithmic term. Please see the full version of the paper for more details including the proof of these claims.

5 Stabilization Time

In this section, we prove tight bounds on the stabilization time of our opinion diffusion process. Let G be a connected graph. (Otherwise, we just need to consider the maximum stabilization time among all the connected components.) Note that stabilization time is only a function of the choice of seed nodes, not their color. After t rounds, all nodes whose distance from the seed set is at most t will be colored. Therefore, the stabilization time is upper-bounded by the diameter $D(G)$, which is the greatest distance between any pair of nodes in G. This bound is obviously tight. Consider two nodes v, u such that $d(v, u) = D$. If v is the only seed node, then it takes D rounds until node u is colored.

What if the seed set is chosen randomly (as in case of a weak attacker)? We prove that if each node is chosen to be a seed node independently w.p. α, then a.a.s. the stabilization time is between $\Omega(\log_\Delta \log n^{\frac{1}{\alpha}})$ (see Theorem 5.4) and $\mathcal{O}((1/\alpha\delta)\log n)$ (see Theorem 5.1).

Theorem 5.1 *For a graph $G = (V, E)$, if each node is a seed node independently w.p. α, then the stabilization time is in $\mathcal{O}((1/\alpha\delta)\log n)$ a.a.s.*

To prove Theorem 5.1, we first provide Lemmas 5.2 and 5.3.

Lemma 5.2 *For a node v in a graph $G = (V, E)$ and $t \in \mathbb{N}$, if $N_t(v) \neq \emptyset$, then $|\hat{N}_t(v)| \ge (t-1)\delta/3$.*

Proof. Consider a triple $(N_{3i}(v), N_{3i+1}(v), N_{3i+2}(v))$ for some $0 \le i \le \lfloor (t+1)/3 \rfloor - 1$. Since $N_{3i+1}(v)$ is non-empty, it must include a node u. All the neighbors of u are in this triple, i.e., $N(u) \subset N_{3i}(v) \cup N_{3i+1}(v) \cup N_{3i+2}(v)$. Since $\deg(u) = |N(v)| \geq \delta$, the triple must include at least $\delta + 1$ nodes. Thus, we get $|\hat{N}_t(v)| \geq \lfloor \frac{t+1}{3} \rfloor (\delta+1) \geq (t-1)\delta/3$. □

Lemma 5.3 *Assume each node in a graph $G = (V, E)$ is a seed node independently w.p. α. Then, a.a.s. there is no node v such that $|\hat{N}_t(v)| \geq (2/\alpha)\log n$ and $\hat{N}_t(v) \cap R_0 = \emptyset$.*

Proof. The probability that $\hat{N}_t(v) \cap R_0 = \emptyset$ is equal to

$$(1-\alpha)^{|\hat{N}_t(v)|} \leq (1-\alpha)^{\frac{2\log n}{\alpha}} \leq \exp(-2\log n) = \frac{1}{n^2}.$$

Thus, a union bound implies that w.p. $1 - 1/n$ there is no node v such that $|\hat{N}_t(v)| \geq \frac{2}{\alpha}\log n$ and $\hat{N}_t(v) \cap R_0 = \emptyset$. □

Proof of Theorem 5.1. Let $t_2^* := (6/\alpha\delta)\log n + 1$. If $N_{t_2^*}(v) = \emptyset$ for a node $v \in V$, it will be colored in less than t_2^* rounds or it will never be colored. On the other hand, if $N_{t_2^*}(v) \neq \emptyset$, by Lemma 5.2 we have $|\hat{N}_{t_2^*}(v)| \geq (t_2^* - 1)\delta/3 = (2/\alpha)\log n$. By Lemma 5.3, we know that a.a.s. every node whose t_2^*-neighborhood is of size at least $(2/\alpha)\log n$ has a seed node in its t_2^*-neighborhood, and thus is colored in at most t_2^* rounds. Hence, a.a.s. after $t_2^* = (6/\sigma\alpha)\log n + 1 = \mathcal{O}((1/\alpha\delta)\log n)$ rounds every node is colored or it will never be colored. □

Tightness. Let C_n^δ be the δ-th power of a cycle C_n, which is a 2δ-regular graph. We can prove that if each node is a seed node independently w.p. α, then a.a.s. there is a node u such that $\hat{N}_{t_3^*}(u) \cap R_0 = \emptyset$, for $t_3^* := (1/16\alpha\delta)\log_2 n - 1$, which implies that it will be colored after the t_3^*-th round. Thus, the process takes at least $t_3^* = \Omega((1/\alpha\delta)\log n)$ rounds to end. A complete proof is given in the full version.

Theorem 5.4 *For a connected graph $G = (V, E)$, if each node is a seed node independently w.p. α, then the stabilization time is in $\Omega(\log_\Delta \log n^{\frac{1}{\alpha}})$ a.a.s.*

Tightness. We can prove that if each node in a complete Δ-ary tree is a seed node independently w.p. α, then a.a.s. after $\mathcal{O}(\log_\Delta \log n^{\frac{1}{\alpha}})$ rounds all nodes are colored. This implies that the lower bound in Theorem 5.4 is tight. The proof of this statement and Theorem 5.4 are provided in the full version of the paper.

6 Hardness Results

Assume you are given a graph $G = (V, E)$ and some $t \in \mathbb{N}$. The goal is to determine whether there is a choice of the seed set for which the process needs exactly t rounds to end. This problem is polynomial-time solvable. If $t > D$, where D is the diameter of G, then the answer is No since the stabilization time is upper-bounded by D. If $t \leq D$, then the answer is Yes. Assume that $d(v, u) = D$ for some nodes $v, u \in V$. Then, if the seed set is equal to $\hat{N}_{D-t}(v)$, the process takes exactly t rounds. However, if in the above problem we require that the seed set must be of a given size s, then the problem becomes NP-hard. It is because the *t-hop dominating set problem*, which is proven to be NP-hard by (Basuchowdhuri and Majumder 2014), can be reduced to this problem. Assume we are given a graph G and integers t, s as input. In the t-hop dominating set problem, the task is to decide whether there is a node set $S \subseteq V$ of size s such that $d(S, v) \leq t$ for any $v \in V$. We observe that the answer to this problem is Yes if and only if there is a seed set of size s for which the process takes at most t rounds. Hence, this problem can be reduced to our stabilization time problem.

Now, assume that you are given an n-node graph G and some integers b, w, and t. Your task is to choose a seed set of b black nodes and w white nodes such that $\hat{w}_t$ is as small as possible. (This is similar to what a strong attacker aims to do.) We call this the Minimum Influence (MI) problem.

Theorem 6.1 *Let $\mathcal{A}$ be a polynomial-time β-approximation algorithm for the MI problem. Then, $\beta > n^{1-\zeta}$ for any $\zeta > 0$ unless P=NP.*

Proof sketch. In the *clique problem*, which is NP-hard, a graph G' and an integer k are given and the task is to determine whether G' has a clique of size k. Let G' and k be the input of the clique problem. Then, we construct an n-node graph G such that the solution of the MI problem for $b = k$, $w = \binom{k}{2}$ on G is $\binom{k}{2}$ if G' has a clique of size k and it is larger than $n^{1-\zeta}\binom{k}{2}$ if G' does not. This yields our claim.

To construct G, we essentially replace each edge and node in a copy of G' with some gadget, which includes a large clique. Our construction is tailored in a way that if G' has a clique of size k, then we can place $w = \binom{k}{2}$ white nodes on the gadgets corresponding to the edges of this clique such that no new white node will be generated. However, when there is no clique of size k, then a large group of nodes will become white after two rounds, no matter how we place the w white nodes. A complete proof is given in the full version of the paper. □

7 Conclusion

It would be interesting to study other variants of our model or attackers. Our results apply to some other settings with some minor changes. For example, our propositions regarding the attackers should still hold if in each round only a random subset of nodes update their color or if each node does not follow the majority rule with some noise probability smaller than ϵ. However, our proof techniques cannot immediately be applied to characterize graph structures resilient to an attacker which can choose the order of updates or add/delete edges. Hence, this is a potential avenue for further research.

The maximum degree in most of real-world social graphs is not too large, which makes them fairly resilient to a moderate and weak attacker. However, they are usually far from being regular or expander, and thus are vulnerable to a strong attacker. A counter-measure might be to add an Erdős-Rényi random graph on top of the underlying graph structure. In the real world, this is essentially the same as asking agents to choose a fraction of their connections at random. Our preliminary experiments on data from SNAP (Leskovec and Krevl 2014) suggest that such modification enhances the resilience significantly. However, a detailed and formal investigation of this subject is left for future research.

Acknowledgments

The author is thankful to Václav Rozhoň and Christoph Grunau for several stimulating conversations, especially on Theorem 6.1.

References

Amini, H.; and Fountoulakis, N. 2014. Bootstrap percolation in power-law random graphs. *Journal of Statistical Physics* 155(1): 72–92.

Anagnostopoulos, A.; Becchetti, L.; Cruciani, E.; Pasquale, F.; and Rizzo, S. 2020. Biased Opinion Dynamics: When the Devil Is in the Details. *arXiv preprint arXiv:2008.13589* .

Apt, K. R.; and Markakis, E. 2014. Social networks with competing products. *Fundamenta Informaticae* 129(3): 225–250.

Auletta, V.; Caragiannis, I.; Ferraioli, D.; Galdi, C.; and Persiano, G. 2015. Minority becomes majority in social networks. In *WINE-15*, 74–88.

Auletta, V.; Fanelli, A.; and Ferraioli, D. 2019. Consensus in opinion formation processes in fully evolving environments. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 6022–6029.

Auletta, V.; Ferraioli, D.; and Greco, G. 2018. Reasoning about Consensus when Opinions Diffuse through Majority Dynamics. In *IJCAI-18*, 49–55.

Auletta, V.; Ferraioli, D.; and Greco, G. 2020. On the Effectiveness of Social Proof Recommendations in Markets with Multiple Products. In *ECAI-20*.

Balister, P.; Bollobás, B.; Johnson, J. R.; and Walters, M. 2010. Random majority percolation. *Random Structures & Algorithms* 36(3): 315–340.

Basuchowdhuri, P.; and Majumder, S. 2014. Finding influential nodes in social networks using minimum k-hop dominating set. In *International Conference on Applied Algorithms*, 137–151. Springer.

Berger, E. 2001. Dynamic monopolies of constant size. *Journal of Combinatorial Theory, Series B* 83(2): 191–200.

Bharathi, S.; Kempe, D.; and Salek, M. 2007. Competitive influence maximization in social networks. In *International workshop on web and internet economics*, 306–311. Springer.

Borodin, A.; Filmus, Y.; and Oren, J. 2010. Threshold models for competitive influence in social networks. In *International workshop on internet and network economics*, 539–550. Springer.

Bredereck, R.; and Elkind, E. 2017. Manipulating Opinion Diffusion in Social Networks. In *IJCAI-17*, 894–900.

Brill, M.; Elkind, E.; Endriss, U.; Grandi, U.; et al. 2016. Pairwise Diffusion of Preference Rankings in Social Networks. In *IJCAI-16*, 130–136.

Chen, N. 2009. On the approximability of influence in social networks. *SIAM Journal on Discrete Mathematics* 23(3): 1400–1415.

Chistikov, D.; Lisowski, G.; Paterson, M.; and Turrini, P. 2020. Convergence of Opinion Diffusion is PSPACE-Complete. In *AAAI*, 7103–7110.

Corò, F.; Cruciani, E.; D'Angelo, G.; and Ponziani, S. 2019. Exploiting social influence to control elections based on scoring rules. In *IJCAI-19*, 201–207.

Corò, F.; D'Angelo, G.; and Velaj, Y. 2019. Recommending Links to Control Elections via Social Influence. *Algorithms* 12(10): 207.

Dubhashi, D. P.; and Panconesi, A. 2009. *Concentration of measure for the analysis of randomized algorithms*. Cambridge University Press.

Faliszewski, P.; Gonen, R.; Koutecký, M.; and Talmon, N. 2018. Opinion Diffusion and Campaigning on Society Graphs. In *IJCAI-18*, 219–225.

Ferraioli, D.; and Ventre, C. 2017. Social pressure in opinion games. In *IJCAI-17*.

Fogelman, F.; Goles, E.; and Weisbuch, G. 1983. Transient length in sequential iteration of threshold functions. *Discrete Applied Mathematics* 6(1): 95–98.

Friedman, J. 2003. A proof of Alon's second eigenvalue conjecture. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, 720–724.

Frischknecht, S.; Keller, B.; and Wattenhofer, R. 2013. Convergence in (social) influence networks. In *DISC-13*, 433–446.

Gärtner, B.; and Zehmakan, A. N. 2018. Majority model on random regular graphs. In *Latin American Symposium on Theoretical Informatics*, 572–583. Springer.

Gärtner, B.; and Zehmakan, A. N. 2020. Threshold behavior of democratic opinion dynamics. *Journal of Statistical Physics* 1–25.

Goles, E.; and Olivos, J. 1980. Periodic behaviour of generalized threshold functions. *Discrete mathematics* 30(2): 187–189.

Grandi, U.; Stewart, J.; and Turrini, P. 2018. The Complexity of Bribery in Network-Based Rating Systems. In *AAAI*, 1047–1054.

Hoory, S.; Linial, N.; and Wigderson, A. 2006. Expander graphs and their applications. *Bulletin of the American Mathematical Society* 43(4): 439–561.

Kaaser, D.; Mallmann-Trenn, F.; and Natale, E. 2016. On the voting time of the deterministic majority process. In *MFCS-16*.

Kempe, D.; Kleinberg, J.; and Tardos, É. 2003. Maximizing the spread of influence through a social network. In *KDD-03*, 137–146.

Le, C. M.; Levina, E.; and Vershynin, R. 2017. Concentration and regularization of random graphs. *Random Structures & Algorithms* 51(3): 538–561.

Leskovec, J.; and Krevl, A. 2014. SNAP Datasets: Stanford Large Network Dataset Collection. http://snap.stanford.edu/data.

Meir, R.; Polukarov, M.; Rosenschein, J.; and Jennings, N. 2010. Convergence to equilibria in plurality voting. In *AAAI-10*, volume 24, 823–828.

Mishra, S.; Radhakrishnan, J.; and Sivasubramanian, S. 2002. On the hardness of approximating minimum monopoly problems. In *FSTTCS-02*, 277–288.

Mossel, E.; Neeman, J.; and Tamuz, O. 2014. Majority dynamics and aggregation of information in social networks. *Autonomous Agents and Multi-Agent Systems* 28(3): 408–429.

Mustafa, N. H.; and Pekeč, A. 2001. Majority consensus and the local majority rule. In *ICALP-01*, 530–542.

N. Zehmakan, A.; and Galam, S. 2020. Rumor spreading: A trigger for proliferation or fading away. *Chaos: An Interdisciplinary Journal of Nonlinear Science* 30(7): 073122.

Pastor-Satorras, R.; and Vespignani, A. 2001. Epidemic spreading in scale-free networks. *Physical review letters* 86(14): 3200.

Peleg, D. 1997. Local majority voting, small coalitions and controlling monopolies in graphs: A review. In *SIROCCO-97*, 152–169.

Peleg, D. 1998. Size bounds for dynamic monopolies. *Discrete Applied Mathematics* 86(2-3): 263–273.

Rogers, E. M. 1962. *Diffusion of innovations*. Simon and Schuster.

Schoenebeck, G.; and Yu, F.-Y. 2018. Consensus of interacting particle systems on erdös-rényi graphs. In *SODA-18*, 1945–1964.

Wilder, B.; and Vorobeychik, Y. 2017. Controlling elections through social influence. In *AAMAS-17*.

Yin, X.; Wang, H.; Yin, P.; and Zhu, H. 2019. Agent-based opinion formation modeling in social network: A perspective of social psychology. *Physica A: Statistical Mechanics and its Applications* 532: 121786.

Zehmakan, A. N. 2020. Opinion forming in Erdős–Rényi random graph and expanders. *Discrete Applied Mathematics* 277: 280–290.

Zhuang, Z.; Wang21, K.; Wang, J.; Zhang, H.; Wang, Z.; and Gong, Z. 2020. Lifting Majority to Unanimity in Opinion Diffusion. In *ECAI-20*.