

Asynchronous Opinion Dynamics in Social Networks

Petra Berenbrink
Universität Hamburg
Hamburg, Germany
petra.berenbrink@uni-hamburg.de

Martin Hoefler
Goethe University Frankfurt
Frankfurt, Germany
mhoefler@cs.uni-frankfurt.de

Dominik Kaaser
Universität Hamburg
Hamburg, Germany
dominik.kaaser@uni-hamburg.de

Pascal Lenzner
Hasso Plattner Institute
Potsdam, Germany
pascal.lenzner@hpi.de

Malin Rau
Universität Hamburg
Hamburg, Germany
malin.rau@uni-hamburg.de

Daniel Schmand
Universität Bremen
Bremen, Germany
schmand@uni-bremen.de

ABSTRACT

Opinion spreading in a society decides the fate of elections, the success of products, and the impact of political or social movements. The model by Hegselmann and Krause is a well-known theoretical model to study such opinion formation processes in social networks. In contrast to many other theoretical models, it does not converge towards a situation where all agents agree on the same opinion. Instead, it assumes that people find an opinion reasonable if and only if it is close to their own. The system converges towards a stable situation where agents sharing the same opinion form a cluster, and agents in different clusters do not influence each other.

We focus on the social variant of the Hegselmann-Krause model where agents are connected by a social network and their opinions evolve in an iterative process. When activated, an agent adopts the average of the opinions of its neighbors having a similar opinion. By this, the set of influencing neighbors of an agent may change over time. To the best of our knowledge, social Hegselmann-Krause systems with asynchronous opinion updates have only been studied with the complete graph as social network. We show that such opinion dynamics with random agent activation are guaranteed to converge for any social network. We provide an upper bound of $O(n|E|^2(\epsilon/\delta)^2)$ on the expected number of opinion updates until convergence, where $|E|$ is the number of edges of the social network. For the complete social network we show a bound of $O(n^3(n^2 + (\epsilon/\delta)^2))$ that represents a major improvement over the previously best upper bound of $O(n^9(\epsilon/\delta)^2)$. Our bounds are complemented by simulations that indicate asymptotically matching lower bounds.

CCS CONCEPTS

• **Theory of computation** → **Social networks**.

KEYWORDS

Hegselmann-Krause Systems; Opinion Dynamics; Social Influence; Asynchronous Updates; Social Networks; Convergence Time

ACM Reference Format:

Petra Berenbrink, Martin Hoefler, Dominik Kaaser, Pascal Lenzner, Malin Rau, and Daniel Schmand. 2022. Asynchronous Opinion Dynamics in Social Networks. In *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), Online, May 9–13, 2022, IFAAMAS*, 9 pages.

Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), P. Faliszewski, V. Mascardi, C. Pelachaud, M.E. Taylor (eds.), May 9–13, 2022, Online. © 2022 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

1 INTRODUCTION

Our opinions are not static. On the contrary, opinions are susceptible to dynamic changes and this is heavily exploited by (social) media, influencers, politicians, and professionals for public relations campaigns and advertising. The way we form our opinions is not a solitary act that simply combines our personal experiences with information from the media. Instead, it is largely driven by interactions with our peers in our social network. We care about the opinions of our peers and relatives, and their opinions significantly influence our own opinion in an asynchronous dynamic process over time. Such opinion dynamics are pervasive in many real-world settings, ranging from small scale townhall meetings, community referendum campaigns, parliamentary committees, and boards of enterprises to large scale settings like political campaigns in democratic societies or peer interactions via online social networks.

The aim for understanding how opinions are formed and how they evolve in multi-agent systems is the driving force behind an interdisciplinary research effort in diverse areas such as sociology, economics, political science, mathematics, physics, and computer science. Initial work on these issues dates back to Downs [32] and early agent-based opinion formation models as proposed by Abelson and Bernstein [1].

In this paper we study an agent-based model for opinion formation on a social network where the opinion of an agent depends both on its own intrinsic opinion and on the opinions of its network neighbors. One of the earliest influential models in this direction was defined by DeGroot [31]. In this model the opinion of an agent is iteratively updated to the weighted average of the opinions of its neighbors. Later, Friedkin and Johnsen [38] extended this by incorporating private opinions. Every agent has a private opinion which does not change and an expressed opinion that changes over time. The expressed opinion of an agent is determined as a function of the expressed opinions of its neighbors and its private opinion.

The main focus of our paper is the very influential model by Hegselmann and Krause [43] that adds an important feature: the set of neighbors that influence a given agent is no longer fixed and the agents' opinions and their respective sets of influencing neighbors co-evolve over time. At any point in time the set of influencing neighbors of an agent are all the neighbors in a given static social network with an opinion close to their own opinion. Hence, agents only adapt their opinions to neighboring agents having an opinion that is not too far away from their own opinion. Note that this adaptation, in turn, might lead to a new set of influencing neighbors.

In sociology this wide-spread behavior is known as homophily [45], which, for example, governs the formation of social networks and explains residential segregation. Co-evolutionary opinion formation helps to analyze and explain current phenomena like filter bubbles in the Internet [48] and social media echo chambers [22] that inhibit opinion exchange and amplify extreme views. The co-evolution of opinions and the sets of influencing neighbors is the key feature of a Hegselmann-Krause system (HKS). However, it is also the main reason why the analysis of the dynamic behavior of a HKS is highly non-trivial and challenging.

Typical questions studied are the convergence properties of the opinion dynamics: Is convergence to stable states guaranteed and if yes, what are upper and lower bounds on the convergence time? Guaranteed convergence is essential since otherwise the predictive power of the model is severely limited. Moreover, studying the convergence *time* of opinion dynamics is crucially important. In general, the analysis of stable states is significantly more meaningful if these states are likely to be reached in a reasonable amount of time, i.e., if quick convergence towards such states is guaranteed. If systems do not stabilize in a reasonable time, stable states lack justification as a prediction of the system’s behavior.

Researchers have investigated the convergence to stable states and the corresponding convergence speed in many variants of the Hegselmann-Krause model. The existing work can be categorized along two dimensions: complete or arbitrary social network and *synchronous* or *asynchronous* updates of the opinions. Synchronous opinion updates means that *all* agents update their opinion at the same time. In systems with asynchronous updates a single agent is selected uniformly at random and only this agent updates its opinion. The main body of recent work focuses on HKSs assuming the complete graph as social network and the synchronous update rule. Interestingly, convergence guarantees and convergence times for the case with asynchronous updates on an arbitrary social network are, to the best of our knowledge, absent from the literature so far. This case is arguably the most realistic setting as social networks are typically sparse, i.e., non-complete, and social interactions and thereby opinion exchange usually happens in an uncoordinated asynchronous fashion.

In this paper we study the following *Hegselmann-Krause system (HKS)*. We have n agents and their opinions are modeled by points in d -dimensional Euclidean space \mathbb{R}^d , for some $d \geq 1$. The agents are connected by a *social network* which does not change over time. At any point of time the set of *influencing neighbors* of an agent is the subset of its neighbors (in the social network) with an opinion of distance at most $\epsilon > 0$ from its own opinion. We assume that in each step a random agent is activated and its opinion is updated to the average of its current opinion and the opinion of all current influencing neighbors. Note in such an asynchronous HKS stable states in the sense that no agent will change its opinion might never be reached. This can be seen by a simple example with two nodes and one edge. Hence, we adopt a natural stability criterion defined by Bhattacharyya and Shiragur [14]. A HKS is in a δ -stable state if and only if each edge in the influence network has length at most δ . For this scenario we prove that the convergence of the opinion dynamics is guaranteed and we give an upper bound on

the expected convergence time of

$$O(n|E|^2(\epsilon/\delta)^2) \leq O(n^5(\epsilon/\delta)^2),$$

where $|E|$ is the cardinality of the edge set of the given social network. We demonstrate the tightness of our derived upper bound by providing empirical agent-based simulations for several topologies of the underlying social network topologies. Note that for complete graphs as social network our bound of $O(n^3(n^2 + (\epsilon/\delta)^2))$ improves the best previously known upper bound of $O(n^9(\epsilon/\delta)^2)$ [34].

1.1 Related Work

We focus our discussion on recent research on Hegselmann-Krause systems and other opinion formation models.

Synchronous HKSs on Complete Networks. Most recent research focused on synchronous opinion updates in complete social networks. For this setting it is known that the process always converges to a state where no agent changes its opinion anymore [21]. We denote such states as *perfectly stable states*. Touri and Nedic [49] prove that any one-dimensional HKS converges in $O(n^4)$ synchronous update rounds to a perfectly stable state. Bhattacharyya et al. [13] improve this upper bound to $O(n^3)$. For d dimensions they show a convergence time of $O(n^{10}d^2)$. For arbitrary d Etesami and Bařar [34] establish a bound of $O(n^6)$ rounds, which is independent of the dimension d . Finally, Martinsson [44] shows that any synchronous d -dimensional HKS converges within $O(n^4)$ update rounds to a perfectly stable state.

Regarding lower bounds, Bhattacharyya et al. [13] construct two-dimensional instances that need at least $\Omega(n^2)$ update rounds before a perfectly stable state is reached. Later, Wedin and Hegarty [50] show that this lower bound holds even in one-dimensional systems.

Synchronous HKSs on Arbitrary Social Networks. In [47], the authors use the probabilistic method to prove that the expected convergence time to a perfectly stable state is infinite for general networks. This also holds for a slightly weaker stability concept than perfect stability: in all future steps an agent’s opinion will not move further than by a given distance δ . To show their result the authors construct a HKSs with infinitely many oscillating states. Their stability notion is also different to the one considered in this paper. We analyze the time to reach a δ -stable state which is defined as a state where any edge in the influence network has length at most δ (see Section 1.2). For δ -stability Bhattacharyya and Shiragur [14] prove that a synchronous HKS with an arbitrary social network reaches a δ -stable state in $O(n^5(\epsilon/\delta)^2)$ synchronous rounds.

Asynchronous HKSs. Compared to the synchronous case, the existing results for asynchronous HKSs are rather limited. To the best of our knowledge, they were first studied by Etesami and Bařar [34] where the authors consider δ -equilibria in contrast to δ -stable states. They define a δ -equilibrium as a state where each connected component of the influence network has an Euclidean diameter of at most δ and prove that the expected number of update steps to reach such a state is bounded by $O(n^9(\epsilon/\delta)^2)$ for the complete social network. In general, δ -equilibria are a proper subset of the set of δ -stable states. However, in Section 2 we discuss the equivalence of both stability notions on complete social networks.

Other Opinion Formation Models. In the seminal models by Friedkin and Johnsen [38] (extending earlier work by DeGroot [31]) each agent has an innate opinion and strategically selects an expressed opinion that is a compromise of its innate opinion and the opinions of its neighbors. Recently, co-evolutionary and game-theoretic variants were studied [15–17, 33, 37], and the results focus on equilibrium existence and social quality, measured by the price of anarchy. In the AI and multi-agent systems community, opinion formation is studied intensively. In [4] a co-evolutionary model is investigated, where also the innate opinion may change over time. There is also substantial work on understanding opinion diffusion, i.e., the process of how opinions spread in a social network [2, 18–20, 30, 36]. Moreover, in [24, 25] a framework and a simulator for agent-based opinion formation models is presented. Opinion dynamics and in particular the emergence of echo chambers is modeled with tools from statistical physics in [35, 39]

Another line of related research on opinion dynamics has its roots in randomized rumor spreading and distributed consensus processes (see [6] for a rather recent survey). Communication in these models is typically restricted to constantly many neighbors. A simple and natural protocol in this context is the VOTER process [11, 26, 42, 46], where every agent adopts in each round the opinion of a single, randomly chosen neighbor. Similar processes are the TWOCHOICES process [27–29], the 3MAJORITY dynamics [8, 9, 40], and the Undecided State Dynamics [3, 5, 7, 10, 23, 41].

1.2 Model and Notation

A *Hegselmann-Krause system (HKS)* $(G = (V, E), \varepsilon, x)$ in d dimensions is defined as follows. We are given a *social network* $G = (V, E)$ and a *confidence bound* $\varepsilon \in \mathbb{R}_+$. The n nodes of the social network correspond to the agents, and each agent $v \in V$ has an *initial opinion* $x_0(v) \in \mathbb{R}^d$. We will use the terms agents and nodes interchangeably. As the opinion of agent v is represented by a point in the d -dimensional Euclidean space, we sometimes call it *the position of v* . In step $t \in \mathbb{Z}_{\geq 0}$ the opinion of agent $v \in V$ is denoted as $x_t(v) \in \mathbb{R}^d$. For some constant confidence bound $\varepsilon \in \mathbb{R}_+$ we define the *influencing neighborhood* of agent $v \in V$ at time t as

$$\mathcal{N}_t(v) = \{u \in V \mid \{u, v\} \in E, \|x_t(u) - x_t(v)\|_2 \leq \varepsilon\} \cup \{v\}.$$

In each step t one agent $v \in V$ is chosen uniformly at random and updates its position according to the rule

$$x_{t+1}(v) = \frac{\sum_{u \in \mathcal{N}_t(v)} x_t(u)}{|\mathcal{N}_t(v)|}.$$

If $x_t(v) \neq x_{t+1}(v)$, then we say that (the opinion of) agent v has *moved*. Also, in an update of agent v 's position in step t , all other agents do not change their positions, i.e., $x_{t+1}(u) = x_t(u)$ for $u \neq v$.

Given a social network $G = (V, E)$, we define for any edge $e = \{u, v\} \in E$ at time t the length of e as $\|x_t(e)\|_2 = \|x_t(u) - x_t(v)\|_2$. We define the *movement* $m_t(v)$ of agent $v \in V$ at time t as the d -dimensional vector

$$m_t(v) = \frac{\sum_{u \in \mathcal{N}_t(v)} (x_t(u) - x_t(v))}{|\mathcal{N}_t(v)|}.$$

Note that $m_t(v) = x_{t+1}(v) - x_t(v)$ if v is chosen in step t , and hence $\|m_t(v)\|_2$ denotes the distance the agent moves when activated in step t . The *influence network* I_t in step t is given by the social

network G restricted to edges that have length at most ε . More formally, it is defined as $I_t = (V, \mathcal{E}_t)$, where $e = \{u, v\} \in \mathcal{E}_t$ if and only if $u \in \mathcal{N}_t(v)$, i.e., $\|x_t(e)\|_2 \leq \varepsilon$. We define the *state* of a HKS $(G = (V, E), \varepsilon, x)$ at time t as $S_t = (G = (V, E), \varepsilon, x_t)$ and it refers to the positions of the agents at that specific time. If clear from the context, we omit the parameter t . For a fixed state S , the term $\mathcal{N}(v)$ denotes the influencing neighborhood in this state.

We are interested in the expected number of steps that are required until the HKS reaches a δ -stable state, which is a natural stability criterion defined by Bhattacharyya and Shiragur [14]. A HKS is in a δ -stable state, if and only if each edge in the influence network has length at most δ . We call the number of steps to reach a δ -stable state the *convergence time* of the system.

1.3 Our Contribution

We study the convergence time to a δ -stable state in Hegselmann-Krause systems with an arbitrary initial state and an arbitrary given social network, where we update one uniformly at random chosen agent in each step. To the best of our knowledge, this is the first analysis of the variant of HKSs that feature asynchronous opinion updates on a given arbitrary social network. For these systems, we prove the following:

THEOREM 1. *For a d -dimensional HKS $S_0 = (G = (V, E), \varepsilon, x)$, the expected convergence time to a δ -stable state under uniform random asynchronous updates is $O(\Phi(S_0)n|E|/\delta^2) \leq O(n|E|^2 (\varepsilon/\delta)^2)$.*

For graphs with $|E| = O(n)$, for example graphs with constant maximum node degree, the theorem immediately shows an expected convergence time of $O(n^3 (\varepsilon/\delta)^2)$. Interestingly, our upper bound on the expected convergence time in the asynchronous process on arbitrary social networks is of the same order as the best known upper bound of $O(n^5 (\varepsilon/\delta)^2)$ for the synchronous process [14] where *all* agents are activated in parallel.

Furthermore, we show that the convergence time stated in Theorem 1 also transfers to the model of Etesami and Başar [34]. They showed that a HKS with asynchronous opinion updates on a complete social network converges to a δ -equilibrium in $O(n^9 (\varepsilon/\delta)^2)$ steps, thus it is a major improvement over their analysis. However, since on arbitrary social networks δ -stability does not imply a δ -equilibrium, it is open if the bound given in Theorem 1 also holds for the convergence time to δ -equilibria.

Moreover, for the special case of a complete social network with asynchronous opinion updates, i.e., the case considered by Etesami and Başar [34], we show the following even stronger result that holds for arbitrary δ :

THEOREM 2. *Let $(G = (V, E), \varepsilon, x)$ be any instance of a d -dimensional HKS and let $G = K_n$ be the complete social network. Using uniform random asynchronous update steps, the expected convergence time to a δ -stable state is at most $O(n^3 (n^2 + (\varepsilon/\delta)^2))$.*

To prove these results we extend the potential function used in [34]. The main ingredient for strongly improving the upper bound derived in [34] is to significantly tighten and generalize the proof by Etesami and Başar [34]. To do so we develop a projection argument (see Lemma 3) and a new analysis of the expected movement of a randomly chosen agent. This allows us to improve the bound on the expected drop of the potential function (see Lemma 7).

To complement our upper bound results we demonstrate that our analysis method is tight in the sense that using this potential function and studying step by step drop, one cannot improve the results. We present a family of examples where the expected potential drop is exactly of the same order as our upper bound (see Theorem 8). Moreover, we present a family of one-dimensional HKSs where in expectation $\Omega(n^2)$ steps are needed to reach a δ -stable state (see Theorem 11). Last but not least, in Section 5 we provide some simulation results for two specific social network topologies. Our empirically derived lower bounds asymptotically match our theoretically proven upper bound from Theorem 1.

2 SOCIAL HEGSELMANN-KRAUSE SYSTEMS

In this section, we prove Theorem 1 using three intermediate steps. First, we show that there is a projection of any state of a d -dimensional HKS to one dimension while preserving the main properties of the HKS. In the next step, we prove for any 1-dimensional HKS that the term $\sum_{v \in V} (|\mathcal{N}_t(v)| \cdot \|m_t(v)\|_2)$ can be lower bounded by the twice the length of the longest edge in the system (see Corollary 5). Finally, we prove that the drop in the potential when activating an agent v can be lower bounded by a function of its movement $m_t(v)$ (see Lemma 6). Combining these three properties on HKSs enables us to prove the theorem.

Let $S = (G = (V, E), \varepsilon, x)$ be a state of some d -dimensional HKS with influential network $I = (V, \mathcal{E})$. For some arbitrary edge $e = \{u, w\} \in \mathcal{E}$, we will project the state S to a state \bar{S}_e of some 1-dimensional HKS. We define the projected state \bar{S}_e along edge $e = \{u, w\}$ with the help of the projection vector

$$p = \frac{(x(u) - x(w))}{\|x(u) - x(w)\|_2},$$

where the order of u and w is chosen arbitrarily. We define

$$\bar{S}_e = (\bar{G} = (V, \bar{E}), \varepsilon, \bar{x}),$$

as follows. We project the position of each agent $v \in V$ to

$$\bar{x}(v) = x(v)^T p \in \mathbb{R}.$$

Furthermore, in the graph $\bar{G} = (V, \bar{E})$ of the projected system, we restrict the set of edges \bar{E} to the ones, which are edges of the influence network in the original state, i.e., $\bar{E} = \mathcal{E}$. For an agent $v \in V$, we denote by $\bar{\mathcal{N}}(v)$ its influencing neighborhood, and by $\bar{m}(v)$ its movement in \bar{S}_e .

In the following lemma, we prove that the projected system behaves similarly to the original system in the sense that the length of the edge e stays the same and the influence network does not change. Furthermore, the agents in the original HKS move at least as much as the agents in the projected state, when activated.

LEMMA 3. *Let $S = (G = (V, E), \varepsilon, x)$ be a state of a d -dimensional HKS with influence network $I = (V, \mathcal{E})$ and $e = \{u, w\} \in \mathcal{E}$. Then it holds for the projected state \bar{S}_e defined as above that*

- (1) $\|x(u) - x(w)\|_2 = |\bar{x}(u) - \bar{x}(w)|$,
- (2) $\mathcal{N}(v) = \bar{\mathcal{N}}(v)$,
- (3) $\sum_{v \in V} (|\mathcal{N}(v)| \cdot \|m(v)\|_2) \geq \sum_{v \in V} (|\bar{\mathcal{N}}(v)| \cdot |\bar{m}(v)|)$.

PROOF. Let p be the projection vector used to generate \bar{S}_e . To see statement (1) note that

$$|\bar{x}(u) - \bar{x}(w)| = |x(u)^T p - x(w)^T p| = |(x(u) - x(w))^T p|$$

$$\begin{aligned} &= \left| \frac{(x(u) - x(w))^T (x(u) - x(w))}{\|(x(u) - x(w))\|_2} \right| \\ &= \|(x(u) - x(w))\|_2. \end{aligned}$$

To prove statement (2), we show that for each pair $v, v' \in V$ it holds that $\|x(v) - x(v')\|_2 \geq |\bar{x}(v) - \bar{x}(v')|$.

$$\begin{aligned} |\bar{x}(v) - \bar{x}(v')| &= |x(v)^T p - x(v')^T p| = |(x(v) - x(v'))^T p| \\ &= \left| \frac{(x(v) - x(v'))^T (x(u) - x(w))}{\|(x(u) - x(w))\|_2} \right| \end{aligned}$$

$$\stackrel{\text{Cauchy-Schwarz}}{\leq} \|(x(v) - x(v'))\|_2.$$

Since the difference between projected positions of agents is at most as large as the difference between their original positions and since \bar{E} contains only the edges of the influence network in the original state, it holds that $\mathcal{N}(v) = \bar{\mathcal{N}}(v)$. Finally, it holds that

$$\begin{aligned} \|m(v)\|_2 &= \left\| \frac{\sum_{u \in \mathcal{N}(v)} (x(u) - x(v))}{|\mathcal{N}(v)|} \right\|_2 \\ &= \left\| \frac{\sum_{u \in \mathcal{N}(v)} (x(u) - x(v))}{|\mathcal{N}(v)|} \right\|_2 \\ &\stackrel{\text{Cauchy-Schwarz}}{\geq} \frac{\left(\sum_{u \in \mathcal{N}(v)} (x(u) - x(v)) \right)^T (x(u) - x(w))}{|\mathcal{N}(v)| \|x(u) - x(w)\|_2} \\ &= \left| \frac{\left(\sum_{u \in \mathcal{N}(v)} (x(u)^T p - x(v)^T p) \right)}{|\mathcal{N}(v)|} \right| \\ &= \left| \frac{\sum_{j \in \bar{\mathcal{N}}(v)} (\bar{x}(u) - \bar{x}(v))}{|\bar{\mathcal{N}}(v)|} \right| = |\bar{m}(v)|, \end{aligned}$$

and hence

$$\sum_{v \in V} |\mathcal{N}(v)| \|m(v)\|_2 \geq \sum_{v \in V} |\bar{\mathcal{N}}(v)| |\bar{m}(v)|. \quad \square$$

We now prove a lower bound on the total movement of agents.

LEMMA 4. *Let $S = (G = (V, E), \varepsilon, x)$ be a state of a 1-dimensional HKS, let $c \in \mathbb{R}$ and $V_\ell = \{v \in V \mid x(v) \leq c\}$ and $V_r = V \setminus V_\ell$. Define $E_{\ell,r} = \{\{u, w\} \in E \mid u \in V_\ell, w \in V_r\}$. Then it holds that*

$$\sum_{v \in V} |\mathcal{N}(v)| \|m(v)\|_2 \geq 2 \sum_{e \in E_{\ell,r}} \|x(e)\|_2$$

PROOF. We observe

$$\begin{aligned} \sum_{v \in V_\ell} |\mathcal{N}(v)| \|m(v)\|_2 &\geq \sum_{v \in V_\ell} |\mathcal{N}(v)| |m(v)| \\ &= \sum_{v \in V_\ell} |\mathcal{N}(v)| \sum_{u \in \mathcal{N}(v)} \frac{x(u) - x(v)}{|\mathcal{N}(v)|} \\ &= \sum_{v \in V_\ell} \sum_{u \in V_r} (x(u) - x(v)) = \sum_{e \in E_{\ell,r}} \|x(e)\|_2. \end{aligned}$$

Similarly, it holds that

$$\sum_{v \in V_r} |\mathcal{N}(v)| \|m(v)\|_2 \geq \left| \sum_{v \in V_r} |\mathcal{N}(v)| |m(v)| \right| = \sum_{e \in E_{\ell,r}} \|x(e)\|_2.$$

The lemma follows by combining the two results as $V_r = V \setminus V_\ell$. \square

COROLLARY 5. Let $S = (G = (V, E), \varepsilon, x)$ be a state of a d -dimensional HK system. Let λ be the length of a longest edge in the influence network. Then

$$\sum_{v \in V} |\mathcal{N}_t(v)| \cdot \|m_t(v)\|_2 \geq 2\lambda.$$

PROOF. Let edge $e = \{v_\ell, v_r\} \in \mathcal{E}$ be a longest edge in the influence network and $\|x(e)\|_2 = \lambda$. Let $\bar{S}_e = (\bar{G} = (V, \bar{E}), \varepsilon, \bar{x})$ be the state projected to one dimension along the edge e . By Lemma 3, we know that

$$\sum_{v \in V} |\mathcal{N}(v)| \cdot \|m(v)\|_2 \geq \sum_{v \in V} |\bar{\mathcal{N}}(v)| \cdot |\bar{m}(v)|.$$

Furthermore, we know that the influence network in both systems has the same set of edges, and that the length of the longest edge in the influence network of \bar{I}_t is equal to the length of the longest edge in I_t . Hence $e = \{u, w\} \in E$ is a longest edge in the influence network \bar{I} with $\|x(e)\|_2 = \lambda$.

Analogously to Lemma 4, we partition V into two sets V_ℓ and V_r at $c = (x(u) + x(w))/2$ and define $E_{\ell,r} = \{\{v, v'\} \mid v \in V_\ell, v' \in V_r\}$. Note that $e \in E_{\ell,r}$ and hence

$$\sum_{v \in V} |\bar{m}(v)| |\bar{\mathcal{N}}(v)| \geq 2 \sum_{e \in E_{\ell,r}} \|x(e)\|_2 \geq 2\lambda. \quad \square$$

For any state $S = (G = (V, E), \varepsilon, x)$ of a d -dimensional HKS ($G = (V, E), \varepsilon, x$), we define the following potential function.

$$\Phi(S) = \sum_{\{u,v\} \in E} \min\{\|x(u) - x(v)\|_2^2, \varepsilon^2\}.$$

This potential is upper-bounded by $|E|\varepsilon^2$. In the next step, we will prove a lower bound on the drop in the potential when updating any agent $v \in V$.

LEMMA 6. Let $S_t = (G = (V, E), \varepsilon, x_t)$ be the state of some d -dimensional HKS ($G = (V, E), \varepsilon, x$). Suppose we update the position of agent v and v moves by $m_t(v)$. Let

$$S_{t+1} = (G = (V, E), \varepsilon, x_{t+1})$$

be the new state. The potential decreases by at least

$$\Phi(S_t) - \Phi(S_{t+1}) \geq (|\mathcal{N}_t(v)| + 1) \cdot \|m_t(v)\|_2^2.$$

If the influence network does not change from step t to $t+1$, we obtain equality.

The proof is inspired by the work of Etesami and Başar [34] and can be found in the full version of our paper [12].

We now have the tools to prove a lower bound on the expected potential drop in a single step.

LEMMA 7. For any state $S_t = (G = (V, E), \varepsilon, x_t)$ of some HKS ($G = (V, E), \varepsilon, x$) in step t , when updating an agent chosen uniformly at random resulting in state $S_{t+1} = (G = (V, E), \varepsilon, x_{t+1})$, the expected potential drop is at least

$$\mathbb{E}[\Phi(S_t) - \Phi(S_{t+1})] \geq \frac{2(\lambda_t)^2}{n|\mathcal{E}_t|},$$

where λ_t is the length of the longest edge in the influence network I_t in step t .

PROOF. From Lemma 6 we know that the potential never increases: if we choose agent v to be updated, the potential decreases by at least

$$\Phi(S_t) - \Phi(S_{t+1}) \geq (|\mathcal{N}_t(v)| + 1) \cdot \|m_t(v)\|_2^2.$$

Let e_t be a longest edge in the corresponding influence network of S_t . By Corollary 5, we know that

$$\sum_{v \in V} |\mathcal{N}_t(v)| \cdot \|m_t(v)\|_2 \geq 2\|e_t\|_2.$$

Using Cauchy-Schwarz ($\sum_{v \in V} a_v b_v$)² $\leq \sum_{v \in V} a_v^2 \cdot \sum_{v \in V} b_v^2$ with $a_v = \sqrt{|\mathcal{N}_t(v)|} \cdot \|m_t(v)\|_2$ and $b_v = \sqrt{|\mathcal{N}_t(v)|}$, we conclude that the expected potential drop in each step with an edge with length at least λ_t is at least

$$\begin{aligned} \mathbb{E}[\Phi(S_t) - \Phi(S_{t+1})] &= \sum_{v \in V} \frac{1}{n} \mathbb{E}[\Phi(S_t) - \Phi(S_{t+1}) | v \text{ is updated}] \\ &\geq \frac{1}{n} \sum_{v \in V} (|\mathcal{N}_t(v)| + 1) \|m_t(v)\|_2^2 \\ &\geq \frac{1}{n} \sum_{v \in V} (\sqrt{|\mathcal{N}_t(v)|} \cdot \|m_t(v)\|_2)^2 \\ &\geq \frac{1}{n} \frac{(\sum_{v \in V} |\mathcal{N}_t(v)| \cdot \|m_t(v)\|_2)^2}{\sum_{v \in V} |\mathcal{N}_t(v)|^2} \\ &\geq \frac{1}{n} \cdot \frac{4(\lambda_t)^2}{2|\mathcal{E}_t|}. \quad \square \end{aligned}$$

The proof of Theorem 1 is a direct consequence of Lemma 7.

THEOREM 1. For a d -dimensional HKS $S_0 = (G = (V, E), \varepsilon, x)$, the expected convergence time to a δ -stable state under uniform random asynchronous updates is $O(\Phi(S_0)n|E|/\delta^2) \leq O(n|E|^2(\varepsilon/\delta)^2)$.

PROOF. Note that by definition of the potential function, we have $\Phi(S) \leq \varepsilon^2|E|$ for all states S . We know by Lemma 7 that the expected potential drop at any step t is at least

$$\mathbb{E}[\Phi(S_t) - \Phi(S_{t+1})] \geq \frac{2\delta^2}{n|\mathcal{E}_t|} \geq \frac{2\delta^2}{n|E|}$$

as long as there is an edge with length at least δ . Thus, the expected number of steps to reach a δ -stable state is upper bounded by

$$\frac{|E|\varepsilon^2}{\frac{2\delta^2}{n|E|}} = \frac{n|E|^2}{2} \left(\frac{\varepsilon}{\delta}\right)^2. \quad \square$$

The results from Theorem 1 directly improve the results from Etesami and Başar [34] even though they use a slightly different convergence criterium. In their paper convergence is reached if the diameter of each connected component is bounded by δ and they call this state a δ -equilibrium. They bound the expected number of update steps to reach a δ -equilibrium in the complete social network by $O(n^9(\varepsilon/\delta)^2)$.

Note that if the social network is the complete graph, each connected component in the influence network where each edge has a length of at most $\varepsilon/2$ also must be a complete graph and hence the diameter of this connected component is also bounded by $\varepsilon/2$. Hence, if $\delta \leq \varepsilon/2$, a δ -stable state must be in δ -equilibrium as well. On the other hand, if $\delta > \varepsilon/2$, the expected number of steps to




Figure 1: A state S of a HKS with $\Phi(S) = \Theta(n^2\varepsilon)$ and an expected potential drop of $\Theta(\varepsilon^2/n^3)$. Only edges in \mathcal{E}_0 are presented, and $\hat{m} = \varepsilon/(n^2/16+5n/4-1)$ represents the equal movement of all nodes. Note that the state S is a one-dimensional instance and the position of all nodes of the cliques C_ℓ and C_r have the same position, respectively. We use the second dimension only for a better illustration of the influencing network. We call the state S with its social network reduced to the edges in \mathcal{E}_0 a *Dumbbell* instance.

reach a $\varepsilon/2$ -stable state and hence a δ -equilibrium is bounded by $O(n^5)$ by Theorem 1.

The next theorem shows that our bound on the potential drop per step is tight. Consequently, if we would like to improve the theorem, we have to choose a different potential function and/or consider multiple activations at once.

THEOREM 8. *There is a family of examples with $|E| = \Theta(n^2)$, a potential of $\Theta(n^2\varepsilon^2)$, where the expected potential drop is $\Theta(\varepsilon^2/n^3)$ for the first activation.*

PROOF. Consider the following family of 1-dimensional HKSs $HK_n = (G = (V, E), \varepsilon, x_0)$ such that $|V| = 4n$ for any $n \in \mathbb{N}$, see Fig. 1 for the example for $n = 4$. The set of nodes V is partitioned into sets $C_\ell, C_r, P, \{\ell, r\} \subseteq V$, such that $|C_\ell| = |C_r| = n$ and $|P| = 2n - 2$. The set of edges E is given such that C_ℓ, C_r , and P are cliques while nodes ℓ and r are connected to all nodes.

To define the opinions of the agents that correspond to the nodes V at state S_0 , define $\hat{m} = \varepsilon/(n^2 + 5n - 1)$ and choose

- $x_0(v) = 0$ for each $v \in C_\ell$,
- $x_0(\ell) = \hat{m} \cdot (n + 1)$,
- for each $j \in \{1, \dots, n\}$ there exists a node $v_j \in P$ with

$$x_0(v_j) = x_0(v_{j-1}) + \varepsilon - 3(n - j)\hat{m}$$

where we define $v_0 = \ell$.

- for each $j \in \{n + 1, \dots, 2n - 2\}$ there exists a node $v_j \in P$ with

$$x_0(v_j) = x_0(v_{j-1}) + \varepsilon - 3(j - n)\hat{m}$$

- $x_0(r) = x_0(v) + \varepsilon - 3(n - 1)\hat{m}$
- $x_0(v) = x_0(r) + (n + 1)\hat{m}$ for each $v \in C_r$.

Note that all the edges inside the cliques $C_\ell \cup \{\ell\}$ and $C_r \cup \{r\}$ are in the influence network I_0 , as well as each edge between v_j and v_{j+1} for $j \in \{0, \dots, 2n - 2\}$, where $v_0 = \ell$ and $v_{2n-1} = r$. Also,

$$\begin{aligned} |x_0(v_i) - x_0(v_j)| &\geq x_0(v_2) - x_0(v_0) \\ &= \varepsilon - 3(n - 2)\hat{m} + \varepsilon - 3(n - 1)\hat{m} \\ &= 2\varepsilon - 3(2n - 3)\varepsilon/(n^2 + 5n - 1) \\ &> \varepsilon \end{aligned}$$

for all $0 \leq i, j \leq 2n$ with $|i - j| \geq 2$ and therefore the above mentioned edges are the only ones in I_0 .

We proceed by verifying that for each $v \in V$ it holds that $|m_0(v)| = \hat{m}$. We calculate the movement for ℓ . Let $v \in C_\ell$. Since all n agents in C_ℓ have the same initial position $x_0(v) = 0$, it holds that

$$\begin{aligned} m_0(\ell) &= (n \cdot x_0(v) + x_0(v) - (n + 1)x_0(\ell))/(n + 2) \\ &= (x_0(\ell) + \varepsilon - 3(n - 1)\hat{m} - (n + 1)x_0(\ell))/(n + 2) \\ &= (\varepsilon - 3(n - 1)\hat{m} - n \cdot x_0(\ell))/(n + 2) \\ &= (\varepsilon - 3(n - 1)\hat{m} - n \cdot \hat{m} \cdot (n + 1))/(n + 2) \\ &= (\varepsilon - (n^2 + 4n - 3)\hat{m})/(n + 2) \\ &= (\varepsilon - (n^2 + 4n - 3)\varepsilon/(n^2 + 5n - 1))/(n + 2) \\ &= \varepsilon((n^2 + 5n - 1) - (n^2 + 4n - 3))/((n^2 + 5n - 1)(n + 2)) \\ &= \varepsilon((n + 2)/((n^2 + 5n - 1)(n + 2))) \\ &= \hat{m}. \end{aligned}$$

The calculation of the movement of the other agents is analogous.

Note that the nodes on the path from ℓ to r are the only nodes that have edges not included in I_0 . However, independently of the chosen agent to be updated, no new edge will be activated, since the distance between the corresponding nodes always stays larger than ε . Hence, by Lemma 6, the expected potential drop is given by

$$\begin{aligned} \mathbb{E}[\Phi(S_0) - \Phi(S_1)] &= \frac{1}{n} \sum_{v \in V} (|N_0(v)| + 1) |m_0(v)|^2 \\ &= \frac{1}{n} \left(\frac{n}{2} \left(\frac{n}{4} + 2 \right) + 2 \left(\frac{n}{4} + 3 \right) + \left(\frac{n}{2} - 2 \right) \cdot 4 \right) \hat{m}^2 \\ &= (n/8 + 7/2 - 2/n) \hat{m}^2 \\ &= (n/8 + 7/2 - 2/n) (\varepsilon/(n^2/16 + 5n/4 - 1))^2 \\ &= \Theta(\varepsilon^2/n^3). \end{aligned}$$

On the other hand, there exist $n/2(n/2 - 2)$ edges with length longer than ε and hence $\Phi(S_0) = \Theta(\varepsilon^2 n^2)$. \square

3 IMPROVED RESULTS FOR SPECIFIC SOCIAL NETWORK TOPOLOGIES

In this section we will prove two improved upper bounds, each for a more restricted set of graph classes. The first result holds for the case that the social network is a complete graph, while the second is for the case that in each step of the HKS the influence network is the same as the social network.

THEOREM 2. *Let $(G = (V, E), \varepsilon, x)$ be any instance of a d -dimensional HKS and let $G = K_n$ be the complete social network. Using uniform random asynchronous update steps, the expected convergence time to a δ -stable state is at most $O(n^3 (n^2 + (\varepsilon/\delta)^2))$.*

PROOF. We split this proof into two steps. First, we count the number of possible steps where the influence network has an edge of length at least $\varepsilon/2$. Secondly, we upper-bound the number of steps where the longest edge of the influence network is in $[\delta, \varepsilon/2]$.

Assume in step t there is an edge in the influence network with length at least $\varepsilon/2$. Let S_t and S_{t-1} denote the states of the HKS in steps t and $t - 1$, respectively. In this case, by Lemma 7, we have

$$\mathbb{E}[\Phi(S_t) - \Phi(S_{t-1})] \geq \frac{\varepsilon^2}{n|\mathcal{E}_t|}.$$

As a consequence, the expected number of such steps is bounded by $|E|^2 n = O(n^5)$.

CLAIM 9. *Let I_t be the current influence network. The following property holds in I_t . If all edges have length at most $\varepsilon/2$, each connected component $C_i = (V_i, E_i)$ for $V_i \subseteq V$ is a complete graph.*

PROOF OF THE CLAIM. Assume that $v, u \in V_i$, but $\{v, u\} \notin E_i$. Then there exists a shortest path $P = (v, w_1, \dots, w_k, u)$ of length at least 2 from v to u where each edge has a length of at most $\varepsilon/2$. As a consequence, the distance between v and w_2 has to be smaller than ε and the edge between v and w_2 has to exist in the influence network. Hence P is not the shortest path contradicting the assumption. \square

For the rest of the proof assume that all edges in the influence network are shorter than $\varepsilon/2$ and there exists one edge with length at least δ . We project the HKS to one dimension along the longest edge. By Lemma 3 we know that in the projected graph no edge increases its length and there exists an edge with length at least δ .

For each connected component $C_i = (V_i, E_i)$ define by $\lambda_i(t)$ the length of the longest edge in the connected component. We bound the total movement in this component from below using Lemma 4.

Let $e_i = \{u, w\}$ be the longest edge of the connected component $C_i = (V_i, E_i)$. We partition V_i into $V_{i,\ell}$ and $V_{i,r}$ at $c = (x_t(u) + x_t(w))/2$ and we define the set $E_{\ell,r}$ as in Lemma 4. Since each node from $V_{i,\ell}$ is connected to w while each node from $V_{i,r}$ is connected to u , the set $E_{\ell,r}$ contains at least $(|V_i| - 1)$ edges of length at least $\lambda_i/2$ and one of them has length λ_i . As a consequence, $\sum_{e \in E_{\ell,r}} \|e\| \geq |V_i| \lambda_i/2$ and hence, by Lemma 4,

$$|V_i| \sum_{v \in V_i} |m_t(v)| = \sum_{v \in V_i} |\mathcal{N}_t(v)| |m_t(v)| \geq 2 \sum_{e \in E_{\ell,r}} \|e\| \geq |V_i| \lambda_i(t)$$

and therefore

$$\sum_{v \in V_i} |m_t(v)| \geq \lambda_i(t).$$

As a consequence, it holds that

$$\begin{aligned} \mathbb{E}[\Phi(S_t) - \Phi(S_{t-1})] &\geq \frac{1}{n} \sum_{v \in V} (|\mathcal{N}_t(v)| + 1) \|m_t(v)\|_2^2 \\ &\geq \frac{1}{n} \sum_{i=1}^k (|V_i| + 1) \sum_{v \in V_i} \|m_t(v)\|_2^2 \\ &\geq \frac{1}{n} \sum_{i=1}^k (|V_i| + 1) \left(\sum_{v \in V_i} \|m_t(v)\|_2 \right)^2 / |V_i| \\ &> \frac{1}{n} \cdot \sum_{i=1}^k (\lambda_i(t))^2. \end{aligned}$$

Since one of the edges $\lambda_i(t)$ has length at least δ , the expected potential drop is at least δ^2/n . Therefore, in expectation there are at most $O(|E|n(\varepsilon/\delta)^2)$ steps where the length of the longest edge is in $[\delta, \varepsilon/2]$. Combining the two results finishes the proof. \square

We say a HKSs is *socially stable* if independently of the update steps the influence network is always equal to the social network. For these systems, we can prove a better upper bound on the expected number of steps needed to reach a δ -stable state. Examples for such graphs are the path, where all the nodes are positioned with equal

distance of at most ε and the graph from Theorem 8, if the social network for the latter is reduced to the set of edges in E_0 .

THEOREM 10. *Let $(G = (V, E), \varepsilon, x)$ be a HKSs where the social network and the influence network are equal in each step. Using uniform asynchronous update steps, the expected convergence time to a δ -stable state is bounded by $O(n|E|^2 \log(\varepsilon/\delta))$.*

PROOF. Note that at any step it holds that $\Phi(S_t) \leq |E|(\lambda_t)^2$, where λ_t is the length of the longest edge at time t . By Lemma 7, the expected drop of the potential in each step is bounded by $2(\lambda_t)^2/(n|E|)$. As a consequence, for each $i \in \mathbb{N}$ the expected number of steps with $\lambda_t \in [\varepsilon/2^{i+1}, \varepsilon/2^i]$ is bounded by $O(n|E|^2)$. Since for $\lambda_t \in [\delta, \varepsilon]$ there are at most $\log(\varepsilon/\delta)$ such intervals, the expected number of update steps is bounded by $O(n|E|^2 \log(\varepsilon/\delta))$. \square

4 LOWER BOUND

In this section we complement our upper bounds on the expected convergence time with a lower bound. To the best of our knowledge no lower bound for asynchronous updates is known so far.

THEOREM 11. *There exists a family of social 1-dimensional HKSs where, in expectation, at least $\Omega(n^2)$ steps are needed to reach a δ -stable state.*

PROOF. We consider the following family of 1-dimensional HKSs. Let the social network $G = P_n$ be the path with n nodes such that v_i is the i 'th node on the path, and the position of agent v_i is $x_0(v_i) = i \cdot \varepsilon$. It follows that all edges have length exactly $\varepsilon > \delta$, which implies that at least one of the two agents $v_{\lfloor n/2 \rfloor}, v_{\lfloor n/2 \rfloor + 1}$ has to move at some step to reach a δ -stable state.

In the following we prove by induction that agent v_i moves the first time after $\Omega(\min\{n(i-1), n(n-i)\})$ steps in expectation. This certainly holds for agents v_1 and v_n , since they move for the first time upon their first activation.

Note that agent v_i can only move for the first time, after agent v_{i-1} or agent v_{i+1} has moved for the first time because otherwise both incident edges $\{v_{i-1}, v_i\}, \{v_i, v_{i+1}\}$ have the same length ε and hence agent v_i cannot move. By induction hypothesis one of these agents moves for the first time after

$$\min\{\Omega(\min\{n(i-2), n(n-(i-1))\}), \Omega(\min\{ni, n(n-(i+1))\})\} \\ = \Omega(\min\{n(i-2), n(n-i-1)\})$$

steps. Since agents are activated uniformly at random, after activating one of the neighboring agents of v_i for the first time, additional $\Omega(n)$ activations in expectation are needed so that agent v_i can finally move for the first time. Hence in expectation agent v_i moves for the first time after $\Omega(\min\{n(i-1), n(n-i)\})$ steps.

As a consequence, agent $v_{\lfloor n/2 \rfloor}$ or agent $v_{\lfloor n/2 \rfloor + 1}$ moves for the first time after $\Omega(n^2)$ steps in expectation. \square

5 SIMULATION RESULTS

For corroborating our theoretical findings, we performed agent-based simulations of asynchronous Hegselmann-Krause opinion dynamics in one dimension on two types of initial HKS states called *Path* and *Dumbbell*. They are defined as follows:

- *Path*: The given social network is a path graph. Initially, the agents' opinions are uniformly distributed in one dimension




Figure 2: The plot shows the normalized convergence time: the number of agent activations until a δ -stable state has been reached, divided by n^3 . The data indicate that the convergence time on *Path* instances with equal distances scales as n^3 and on *Dumbbell* instances it scales as n^4 .

with equal distance of ε so that the influence network forms a path graph with uniform edge length of ε .

- *Dumbbell*: This is the state constructed in the proof of Theorem 8 using the dumbbell graph, except that the social network contains only the edges that are in \mathcal{E}_0

In our simulations we fixed $\varepsilon = 100$ and $\delta = 1$. For each initial HKS state on social networks with varying numbers of agents n , we simulated 100 independent runs of random activations needed to reach a δ -stable state.

We present our simulation results in Fig. 2. There, the obtained number of activations divided by n^3 is plotted via a box plots that summarize the results for each configuration. Since for *Path* instances the number of activations appear to be constant, we observe that we need $\Theta(n^3)$ activations for *Path* instances. On the other hand, the number of activations seem to grow linearly in n for *Dumbbell* instances. This hints at $\Theta(n^4)$ activations until *Dumbbell* instances reach a δ -stable for constant ε and δ .

Note that by construction, in the first step the potential function of both instance types is bounded by $\Phi(S_0) = \Theta(n\varepsilon^2)$. Applying Theorem 1 yields an upper bound of $\Phi(S_0)/(2\delta^2/(n|E|)) = O(n\varepsilon^2/(2\delta^2/(n|E|)))$, which yields an upper bound of $O(n^3(\varepsilon/\delta)^2)$ for *Path* instances and $O(n^4(\varepsilon/\delta)^2)$ for *Dumbbell* instances. Thus, if the empirically observed lower bounds on the expected number of steps until convergence are in fact true, our theoretical analysis is tight for these two graph classes with respect to the dependence on the number of agents.

6 CONCLUSION

In this paper we present the first analysis of the convergence time of asynchronous Hegselmann-Krause opinion dynamics on arbitrary social networks. As our main result, we derive an upper bound of $O(n|E|^2(\varepsilon/\delta)^2)$ expected random activations until a δ -stable state is reached. This bound significantly improves over the state-of-the-art upper bound for the special case with a given complete social

network. Moreover, our simulation results on one dimensional instances with a path graph or a dumbbell graph as social network indicate that our theoretical upper bound is tight for these instances. Our theoretical lower bound on the expected convergence time is the first proven non-trivial lower bound for asynchronous opinion updates. A challenging open problem is to improve this lower bound so that it matches our proven upper bound. As the experimental results suggest, this might be possible. However, proving lower bounds for the asynchronous setting seems to be much more involved compared to the analysis of synchronous opinion dynamics as the specific order of agent activations determines which of the possibly many δ -stable states with possibly very different potential function values is reached.

It might be possible to prove better bounds for specific social network topologies. Regarding this, it would be interesting to consider social networks that have similar features as real-world social networks. Moreover, another direction for future work is to consider social networks with directed and possibly weighted edges. This would more closely mimic the structure of real-world neighborhood influences and it would allow to study asymmetric influence settings found in online social networks like Twitter. Another promising extension would be to incorporate the influence of external factors like publicity campaigns.

Acknowledgements

We thank the anonymous reviewers for suggesting additional real-world applications for our model, for pointing us to additional related work, and for sharing interesting ideas on future work.

This work was supported by DFG Research Group ADYN under grant DFG 411362735.

REFERENCES

- [1] Robert Abelson and Alex Bernstein. 1963. A computer simulation model of community referendum controversies. *Public Opinion Quarterly* 27, 1 (1963), 93–122.
- [2] Aris Anagnostopoulos, Luca Becchetti, Emilio Cruciani, Francesco Pasquale, and Sara Rizzo. 2020. Biased Opinion Dynamics: When the Devil is in the Details. In *Proc. 29th Int. Joint Conf. Artif. Intell. (IJCAI)*. 53–59.
- [3] Dana Angluin, James Aspnes, and David Eisenstat. 2008. A simple population protocol for fast robust approximate majority. *Distrib. Comput.* 21, 2 (2008), 87–102.
- [4] Vincenzo Auletta, Angelo Fanelli, and Diodato Ferrarioli. 2019. Consensus in Opinion Formation Processes in Fully Evolving Environments. In *Proc. 33rd Conf. Artif. Intell. (AAAI)*. 6022–6029.
- [5] Gregor Bankhamer, Petra Berenbrink, Felix Biermeier, Robert Elsässer, Hamed Hosseinpour, Dominik Kaaser, and Peter Kling. 2022. Fast Consensus via the Unconstrained Undecided State Dynamics. In *Proc. Symp. Discret. Algorithms (SODA)*. 3417–3429.
- [6] Luca Becchetti, Andrea E. F. Clementi, and Emanuele Natale. 2020. Consensus Dynamics: An Overview. *SIGACT News* 51, 1 (2020), 58–104.
- [7] Luca Becchetti, Andrea E. F. Clementi, Emanuele Natale, Francesco Pasquale, and Riccardo Silvestri. 2015. Plurality Consensus in the Gossip Model. In *Proc. 26th Symp. Discret. Algorithms (SODA)*. 371–390.
- [8] Luca Becchetti, Andrea E. F. Clementi, Emanuele Natale, Francesco Pasquale, Riccardo Silvestri, and Luca Trevisan. 2017. Simple dynamics for plurality consensus. *Distrib. Comput.* 30, 4 (2017), 293–306.
- [9] Petra Berenbrink, Andrea E. F. Clementi, Robert Elsässer, Peter Kling, Frederik Mallmann-Trenn, and Emanuele Natale. 2017. Ignore or Comply?: On Breaking Symmetry in Consensus. In *Proc. Symp. Principles of Distrib. Comput. (PODC)*. 335–344.
- [10] Petra Berenbrink, Tom Friedetzky, George Giakkoupis, and Peter Kling. 2016. Efficient Plurality Consensus, Or: the Benefits of Cleaning up from Time to Time. In *Proc. 43rd Int. Colloq. Autom. Lang. Programming (ICALP)*. 136:1–136:14.
- [11] Petra Berenbrink, George Giakkoupis, Anne-Marie Kermmarrec, and Frederik Mallmann-Trenn. 2016. Bounds on the Voter Model in Dynamic Networks. In

- Proc. 43rd Int. Colloq. Autom. Lang. Programming (ICALP)*. 146:1–146:15.
- [12] Petra Berenbrink, Martin Hoefer, Dominik Kaaser, Pascal Lenzner, Malin Rau, and Daniel Schmand. 2022. Asynchronous Opinion Dynamics in Social Networks. arXiv:2201.12923 [cs.DS]
- [13] Arnab Bhattacharyya, Mark Braverman, Bernard Chazelle, and Huy L. Nguyen. 2013. On the convergence of the Hegselmann-Krause system. In *Proc. Symp. Innov. Theoret. Comput. Sci. (ITCS)*, Robert D. Kleinberg (Ed.). 61–66.
- [14] Arnab Bhattacharyya and Kirankumar Shiragur. 2015. How friends and non-determinism affect opinion dynamics. In *Proc. 54th IEEE Conf. Decision and Control (CDC)*. IEEE, 6466–6471.
- [15] Kshipra Bhawalkar, Sreenivas Gollapudi, and Kamesh Munagala. 2013. Coevolutionary opinion formation games. In *Proc. Symp. Theory Comput. (STOC)*. 41–50.
- [16] Vittorio Bilò, Angelo Fanelli, and Luca Moscardelli. 2018. Opinion formation games with dynamic social influences. *Theor. Comput. Sci.* 746 (2018), 73–87.
- [17] David Bindel, Jon M. Kleinberg, and Sigal Oren. 2015. How bad is forming your own opinion? *Games Econ. Behav.* 92 (2015), 248–265.
- [18] Sirin Botan, Umberto Grandi, and Laurent Perrussel. 2019. Multi-Issue Opinion Diffusion under Constraints. In *Proc. 18th Conf. Auton. Agents and Multi-Agent Syst. (AAMAS)*. 828–836.
- [19] Robert Brederick and Edith Elkind. 2017. Manipulating Opinion Diffusion in Social Networks. In *Proc. 26th Int. Joint Conf. Artif. Intell. (IJCAI)*, Carles Sierra (Ed.). 894–900.
- [20] Robert Brederick, Lilian Jacobs, and Leon Kellerhals. 2020. Maximizing the Spread of an Opinion in Few Steps: Opinion Diffusion in Non-Binary Networks. In *Proc. 29th Int. Joint Conf. Artif. Intell. (IJCAI)*. 1622–1628.
- [21] Bernard Chazelle. 2011. The Total s-Energy of a Multiagent System. *SIAM J. Control. Optim.* 49, 4 (2011), 1680–1706.
- [22] Matteo Cinelli, Gianmarco De Francisci Morales, Alessandro Galeazzi, Walter Quattrociocchi, and Michele Starnini. 2021. The echo chamber effect on social media. *Proceedings of the National Academy of Sciences* 118, 9 (2021).
- [23] Andrea E. F. Clementi, Mohsen Ghaffari, Luciano Gualà, Emanuele Natale, Francesco Pasquale, and Giacomo Scornavacca. 2018. A Tight Analysis of the Parallel Undecided-State Dynamics with Two Colors. In *Proc. 43rd Symp. Math. Found. Comput. Sci. (MFCS)*. 28:1–28:15.
- [24] Adam Coates, Liangxiu Han, and Anthony Kleerekoper. 2018. A Unified Framework for Opinion Dynamics. In *Proc. 17th Conf. Auton. Agents and Multi-Agent Syst. (AAMAS)*. 1079–1086.
- [25] Adam Coates, Liangxiu Han, and Anthony Kleerekoper. 2018. A Unified Opinion Framework Simulator. In *Proc. 17th Conf. Auton. Agents and Multi-Agent Syst. (AAMAS)*. 1803–1805.
- [26] Colin Cooper, Robert Elsässer, Hirotaka Ono, and Tomasz Radzik. 2012. Coalescing random walks and voting on graphs. In *Proc. Symp. Principles of Distrib. Comput. (PODC)*. 47–56.
- [27] Colin Cooper, Robert Elsässer, and Tomasz Radzik. 2014. The Power of Two Choices in Distributed Voting. In *Proc. 41st Int. Colloq. Autom. Lang. Programming (ICALP)*. 435–446.
- [28] Colin Cooper, Robert Elsässer, Tomasz Radzik, Nicolas Rivera, and Takeharu Shiraga. 2015. Fast Consensus for Voting on General Expander Graphs. In *Proc. 29th Symp. Distrib. Comput. (DISC)*. 248–262.
- [29] Colin Cooper, Tomasz Radzik, Nicolas Rivera, and Takeharu Shiraga. 2017. Fast Plurality Consensus in Regular Expanders. In *Proc. 31st Symp. Distrib. Comput. (DISC)*. 13:1–13:16.
- [30] Abir De, Sourangshu Bhattacharya, and Niloy Ganguly. 2018. Shaping Opinion Dynamics in Social Networks. In *Proc. 17th Conf. Auton. Agents and Multi-Agent Syst. (AAMAS)*. 1336–1344.
- [31] Morris H DeGroot. 1974. Reaching a consensus. *J. Amer. Stat. Assoc.* 69, 345 (1974), 118–121.
- [32] Anthony Downs. 1957. *An economic theory of democracy*. Harper & Row New York.
- [33] Markos Epitropou, Dimitris Fotakis, Martin Hoefer, and Stratis Skoulakis. 2019. Opinion Formation Games with Aggregation and Negative Influence. *Theory Comput. Syst.* 63, 7 (2019), 1531–1553.
- [34] Seyed Rasoul Etesami and Tamer Başar. 2015. Game-Theoretic Analysis of the Hegselmann-Krause Model for Opinion Dynamics in Finite Dimensions. *IEEE Trans. Autom. Control.* 60, 7 (2015), 1886–1897.
- [35] Tucker Evans and Feng Fu. 2018. Opinion formation on dynamic networks: identifying conditions for the emergence of partisan echo chambers. *Royal Society open science* 5, 181122 (2018).
- [36] Piotr Faliszewski, Rica Gonen, Martin Koutecký, and Nimrod Talmon. 2018. Opinion Diffusion and Campaigning on Society Graphs. In *Proc. 27th Int. Joint Conf. Artif. Intell. (IJCAI)*. 219–225.
- [37] Dimitris Fotakis, Dimitris Palyvos-Giannas, and Stratis Skoulakis. 2016. Opinion Dynamics with Local Interactions. In *Proc. 25th Int. Joint Conf. Artif. Intell. (IJCAI)*. 279–285.
- [38] Noah E Friedkin and Eugene C Johnsen. 1990. Social influence and opinions. *J. Math. Soc.* 15, 3-4 (1990), 193–206.
- [39] Feng Fu and Long Wang. 2008. Coevolutionary dynamics of opinions and networks: From diversity to uniformity. *Phys. Rev. E* 78, 1 (2008), 016104.
- [40] Mohsen Ghaffari and Johannes Lengler. 2018. Nearly-Tight Analysis for 2-Choice and 3-Majority Consensus Dynamics. In *Proc. Symp. Principles of Distrib. Comput. (PODC)*. 305–313.
- [41] Mohsen Ghaffari and Merav Parter. 2016. A Polylogarithmic Gossip Algorithm for Plurality Consensus. In *Proc. Symp. Principles of Distrib. Comput. (PODC)*. 117–126.
- [42] Yehuda Hassin and David Peleg. 2001. Distributed Probabilistic Polling and Applications to Proportionate Agreement. *Inf. Comput.* 171, 2 (2001), 248–268.
- [43] Rainer Hegselmann and Ulrich Krause. 2002. Opinion dynamics and bounded confidence models, analysis, and simulation. *J. Artif. Societ. Soc. Simul.* 5, 3 (2002).
- [44] Anders Martinsson. 2016. An improved energy argument for the Hegselmann-Krause model. *J. Diff. Equat. Appl.* 22, 4 (2016), 513–518.
- [45] Miller McPherson, Lynn Smith-Lovin, and James M Cook. 2001. Birds of a feather: Homophily in social networks. *Annual Review of Sociology* 27, 1 (2001), 415–444.
- [46] Toshio Nakata, Hiroshi Imahayashi, and Masafumi Yamashita. 2000. A probabilistic local majority polling game on weighted directed graphs with an application to the distributed agreement problem. *Networks* 35, 4 (2000), 266–273.
- [47] Rohit Parasnian, Massimo Franceschetti, and Behrouz Touri. 2019. On the Convergence Properties of Social Hegselmann-Krause Dynamics. arXiv:1909.03485 [math.OA]
- [48] Eli Pariser. 2011. *The filter bubble: What the Internet is hiding from you*. Penguin UK.
- [49] Behrouz Touri and Angelia Nedic. 2011. Discrete-time opinion dynamics. In *Conference Record of the 45th Asilomar Conference on Signals, Systems and Computers (ACSSC)*, Michael B. Matthews (Ed.). 1172–1176.
- [50] Edwin Wedin and Peter Hegarty. 2015. A Quadratic Lower Bound for the Convergence Rate in the One-Dimensional Hegselmann-Krause Bounded Confidence Dynamics. *Discret. Comput. Geom.* 53, 2 (2015), 478–486.