



Asynchronous opinion dynamics in social networks

Petra Berenbrink¹ · Martin Hoefler² · Dominik Kaaser¹ · Pascal Lenzner³ · Malin Rau¹ · Daniel Schmand⁴

Received: 21 November 2022 / Accepted: 28 March 2024
© The Author(s) 2024

Abstract

Opinion spreading in a society decides the fate of elections, the success of products, and the impact of political or social movements. A prominent model to study opinion formation processes is due to Hegselmann and Krause. It has the distinguishing feature that stable states do not necessarily show consensus, i.e., the population of agents might not agree on the same opinion. We focus on the social variant of the Hegselmann–Krause model. There are n agents, which are connected by a social network. Their opinions evolve in an iterative, asynchronous process, in which agents are activated one after another at random. When activated, an agent adopts the average of the opinions of its neighbors having a similar opinion (where similarity of opinions is defined using a parameter ε). Thus, the set of influencing neighbors of an agent may change over time. We show that such opinion dynamics are guaranteed to converge for any social network. We provide an upper bound of $O(n|E|^2(\varepsilon/\delta)^2)$ on the expected number of opinion updates until convergence to a stable state, where $|E|$ is the number of edges of the social network, and δ is a parameter of the stability concept. For the complete social network we show a bound of $O(n^3(n^2 + (\varepsilon/\delta)^2))$ that represents a major improvement over the previously best upper bound of $O(n^9(\varepsilon/\delta)^2)$.

Keywords Hegselmann–Krause Systems · Opinion Formation · Asynchronous Dynamics · Social Networks · Convergence Time

✉ Martin Hoefler
mhoefler@em.uni-frankfurt.de

Petra Berenbrink
petra.berenbrink@uni-hamburg.de

Dominik Kaaser
dominik.kaaser@uni-hamburg.de

Pascal Lenzner
pascal.lenzner@hpi.de

Malin Rau
malin.rau@uni-hamburg.de

Daniel Schmand
schmand@uni-bremen.de

¹ Department of Informatics, Universität Hamburg, Vogt-Kölln-Straße 30, 22527 Hamburg, Germany

² Department of Computer Science, Goethe University Frankfurt, Robert-Mayer-Straße 11-15, 60325 Frankfurt, Germany

³ Hasso Plattner Institute, Prof.-Dr.-Helmert-Str. 2-3, 14482 Potsdam, Germany

⁴ Center for Industrial Mathematics, University of Bremen, Bibliothekstraße 5, 28359 Bremen, Germany

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 [31] 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 [30]. 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 adaption, in turn, might lead to a new set of influencing neighbors. In sociology this wide-spread behavior is known as homophily [46], 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 [49] and social media echo chambers [21] 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). 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. While the main body of recent work focuses on HKSs assuming the complete graph as social network and the synchronous update rule, empirical simulations have also been performed with asynchronous updates on arbitrary social networks. Interestingly, convergence guarantees and convergence times for the latter case 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 $\varepsilon > 0$ 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 [13]. 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. We give an upper bound on the expected convergence time of

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

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 analytical lower bounds as well as empirical simulations for several topologies of the underlying social network topologies. Note that for complete graphs as social network our bound of $O(n^3(n^2 + (\varepsilon/\delta)^2))$ improves the best previously known upper bound of $O(n^9(\varepsilon/\delta)^2)$ [33].

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 com-

plete social networks. For this setting it is known that the process always converges to a state where no agent changes its opinion anymore [20]. We denote such states as *perfectly stable states*. Touri and Nedic [50] prove that any one-dimensional HKS converges in $O(n^4)$ synchronous update rounds to a perfectly stable state. Bhattacharyya et al. [12] 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 Başar [33] establish a bound of $O(n^6)$ rounds, which is independent of the dimension d . Finally, Martinsson [45] 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. [12] construct two-dimensional instances that need at least $\Omega(n^2)$ update rounds before a perfectly stable state is reached. Later, Wedin and Hegarty [51] show that this lower bound holds even in one-dimensional systems.

Synchronous HKSs on Arbitrary Social Networks In [48], 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 HKS 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 Sect. 1.2). For δ -stability Bhattacharyya and Shiragur [13] prove that a synchronous HKS with an arbitrary social network reaches a δ -stable state in $O(n^5(\varepsilon/\delta)^2)$ synchronous rounds.

Asynchronous HKSs Compared to the synchronous case, the existing results for asynchronous HKSs are rather limited. On the empirical side, Fortunato [36] investigated the consensus threshold with uniformly chosen initial opinions in asynchronous dynamics on non-complete social networks like grids, Erdős-Rényi graphs, or scale-free random graphs. To the best of our knowledge, convergence guarantees and convergence times on non-complete networks were first studied by Etesami and Başar [33] 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 $O(n^9(\varepsilon/\delta)^2)$ for the complete social network. In general, δ -equilibria are a proper subset of the set of δ -stable states. However, in Sect. 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 [30]) 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 [14–16, 32, 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, 17–19, 29, 35]. Moreover, in [23, 24] 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 [34, 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, 25, 42, 47], where every agent adopts in each round the opinion of a single, randomly chosen neighbor. Similar processes are the TWO-CHOICES process [26–28], the 3MAJORITY dynamics [8, 9, 40], and the Undecided State Dynamics [3, 5, 7, 10, 22, 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* $x(v) \in \mathbb{R}^d$. 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 $x_t(v) \in \mathbb{R}^d$, where $x_0(v) = x(v)$. For some constant confidence bound $\varepsilon \in \mathbb{R}_+$ we define the *influencing neighborhood* of agent $v \in V$ at time t as

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

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 *available movement* $m_t(v)$ of agent $v \in V$ at time t as the d -dimensional vector

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

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 [13]. A HKS is in a δ -stable state if and only if each edge in the influence network has length at most δ . Intuitively, in such a state each agent has a small incentive to further revise the opinion. Hence, it is reasonable to assume that such states represent a stable configuration of the system. Strictly speaking, however, in a δ -stable state the HKS might not be stabilized *entirely* in the sense that agents are unable to achieve further improvement by a deviation *at all*. If agents would continue to revise their opinions, the HKS might subsequently be able to leave the δ -stable state. Put differently, not all such states are attractive. We note, however, that this is a condition shared by the vast majority of approximate stability or equilibrium concepts defined in the literature.

We call the number of steps to reach a δ -stable state the *convergence time* of the system. To track the progress towards convergence, we define the following potential function for any state $S = (G = (V, E), \varepsilon, x)$ of a d -dimensional HKS $(G = (V, E), \varepsilon, x)$:

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

This potential is upper-bounded by $\Phi(S) \leq |E|\varepsilon^2$.

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 [13] 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 [33]. 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 [33], 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 [33]. The main ingredient for strongly improving the upper bound derived in [33] is to significantly tighten and generalize their proof. To do so, we develop a projection argument (see Lemma 4) and a new analysis of the expected available movement of a randomly chosen agent. This allows us to improve the bound on the expected drop of the potential function (see Lemma 8).

To complement our upper bound results, we demonstrate that our analysis method is tight in the sense that by using this potential function and studying the step-by-step drop, one cannot improve the results. We present a family of instances and initial states where the expected potential drop is exactly of the same order as our upper bound (see Theorem 9). Moreover, we present a family of one-dimensional HKSs and initial states where $\Omega(\Phi(S_0)n|E|/\varepsilon^2)$ steps are needed to

reach a δ -stable state (see Theorem 3), thereby matching the upper bound shown in Theorem 1 in terms of the order of n .

Theorem 3 For $\varepsilon = 1$ and $\delta < 1/2$, there exists a family of social 1-dimensional HKSs $(S_{0,4n})_{n \in \mathbb{N}}$ where any given update sequence needs at least $\Omega(\Phi(S_{0,4n})n|E|/\varepsilon^2) = \Omega(n^4)$ steps to reach a δ -stable state.

Notably, this lower bound applies for arbitrary update sequences, while our upper bound applies when the updating agent is chosen uniformly at random. Thus, even when resorting to a smarter choice of the updating agent, one cannot drastically reduce the convergence time in the worst case.

Last but not least, in Sect. 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 in three steps. Recall that for a HKS in d dimensions the opinion are represented by points in the d -dimensional Euclidean space. First, in Lemma 4 we show that for each HKS in d dimensions there exists a mapping to a suitable 1-dimensional HKS, such that the length of all edges does not increase, and the influence network (consisting of the active edges) as well as the length of the longest edge λ is preserved. We use this projection in the second step (see Lemma 6) where we only consider HKS in one dimension: We prove that $\sum_{v \in V} (|\mathcal{N}_t(v)| \cdot \|m_t(v)\|_2) \geq 2\lambda$, where $\mathcal{N}_t(v)$ is the set of neighbors in the influence network and $m_t(v)$ is the available movement of the node. In the third step, we prove that the potential drop due to activating an agent v can be lower bounded by $(|\mathcal{N}_t(v)| + 1) \cdot \|m_t(v)\|_2^2$ (see Lemma 7). Finally, in Lemma 8 we combine these three insights to bound the potential drop.

Let $S = (G = (V, E), \varepsilon, x)$ be a state of some d -dimensional HKS with influence 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

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

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 available 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 4 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 for any $v, v' \in V$ and the projected state \bar{S}_e defined as above it holds that

$$\|x(u) - x(w)\|_2 = |\bar{x}(u) - \bar{x}(w)|, \tag{1}$$

$$\|x(v) - x(v')\|_2 \geq |\bar{x}(v) - \bar{x}(v')|, \tag{2}$$

$$\mathcal{N}(v) = \bar{\mathcal{N}}(v), \text{ and} \tag{3}$$

$$\sum_{v \in V} (|\mathcal{N}(v)| \cdot \|m(v)\|_2) \geq \sum_{v \in V} (|\bar{\mathcal{N}}(v)| \|\bar{m}(v)\|). \tag{4}$$

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

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

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

$$\begin{aligned} &|\bar{x}(v) - \bar{x}(v')| \\ &= |x(v)^\top p - x(v')^\top p| = |(x(v) - x(v'))^\top p| \\ &= \left| \frac{(x(v) - x(v'))^\top (x(u) - x(w))}{\|x(u) - x(w)\|_2} \right| \\ &\stackrel{\text{C.-S.}}{\leq} \|x(v) - x(v')\|_2. \end{aligned}$$

The last inequality uses the Cauchy-Schwarz inequality (C.-S.).

To see (3), note that (because of (2)) the difference between projected positions of agents is at most as large as the difference between their original positions. 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

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

and hence

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

□

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

Lemma 5 *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

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

Note that for each edge $e = \{v, u\}$ with $v, u \in V_\ell$ the second sum contains $x(u) - x(v)$ as well as $x(v) - x(u)$. As such,

$$\sum_{v \in V_\ell} \sum_{u \in V_\ell \cap \mathcal{N}(v)} (x(u) - x(v)) = 0$$

Furthermore, for each edge $e = \{v, u\}$ with $v \in V_\ell$ and $u \in V_r$ it holds that $x(u) - x(v) > 0$, since $x(u) > 0$ and $x(v) < 0$. As a consequence,

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

Similarly, it holds that

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

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

Corollary 6 *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 4 Eq. (4), we know that

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

Furthermore, by Lemma 4, we know that the influence network in both systems has the same set of edges (Eq. (3)), the longest edge e preserves its length in the projection (1), and all other edges do not increase their length (Eq. 2). Therefore, 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 5, 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 $\sum_{v \in V} |\bar{m}(v)| |\tilde{\mathcal{N}}(v)| \geq 2 \sum_{e \in E_{\ell,r}} \|x(e)\|_2 \geq 2\lambda$. □

In the next step, we will prove a lower bound on the drop in the potential when updating any agent $v \in V$.

Lemma 7 *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.

Proof As we activate v , the position of agents $u \neq v$ does not change, but the set of active edges can change and is updated from \mathcal{E}_t to \mathcal{E}_{t+1} . To bound the potential change we consider edges $\mathcal{E}_t \cap \mathcal{E}_{t+1}$, $\mathcal{E}_t \setminus \mathcal{E}_{t+1}$ and $\mathcal{E}_{t+1} \setminus \mathcal{E}_t$. In the set $\mathcal{E}_t \cap \mathcal{E}_{t+1}$ the length of the edges increases above ε while in the set $\mathcal{E}_{t+1} \setminus \mathcal{E}_t$ the length decreases below ε .

By the definition of Φ , we have

$$\begin{aligned} & \Phi(S_t) - \Phi(S_{t+1}) \\ &= \sum_{\{u,v\} \in E} (\min\{\|x_t(v) - x_t(u)\|_2^2, \varepsilon^2\} \\ & \quad - \min\{\|x_{t+1}(v) - x_{t+1}(u)\|_2^2, \varepsilon^2\}) \\ &= \sum_{\substack{\{u,v\} \in \\ \mathcal{E}_t \cap \mathcal{E}_{t+1}}} \left(\|x_t(v) - x_t(u)\|_2^2 - \|x_{t+1}(v) - x_{t+1}(u)\|_2^2 \right) \\ & \quad + \sum_{\substack{\{u,v\} \in \\ \mathcal{E}_t \setminus \mathcal{E}_{t+1}}} \left(\|x_t(v) - x_t(u)\|_2^2 - \varepsilon^2 \right) \\ &= \sum_{\substack{\{u,v\} \in \\ \mathcal{E}_t \cap \mathcal{E}_{t+1}}} \left(\|x_t(v) - x_t(u)\|_2^2 - \|x_{t+1}(v) - x_{t+1}(u)\|_2^2 \right) \\ & \quad + \sum_{\substack{\{u,v\} \in \\ \mathcal{E}_{t+1} \setminus \mathcal{E}_t}} \left(\varepsilon^2 - \|x_{t+1}(v) - x_{t+1}(u)\|_2^2 \right) \\ &\geq \sum_{\substack{\{u,v\} \in \\ \mathcal{E}_t \cap \mathcal{E}_{t+1}}} \left(\|x_t(v) - x_t(u)\|_2^2 - \|x_{t+1}(v) - x_{t+1}(u)\|_2^2 \right) \\ & \quad + \sum_{\substack{\{u,v\} \in \\ \mathcal{E}_t \setminus \mathcal{E}_{t+1}}} \left(\|x_t(v) - x_t(u)\|_2^2 \right) \\ & \quad - \|x_{t+1}(v) - x_{t+1}(u)\|_2^2. \end{aligned}$$

Note that in this step, we have equality if $\mathcal{E}_t = \mathcal{E}_{t+1}$. We conclude

$$\begin{aligned} & \Phi(S_t) - \Phi(S_{t+1}) \\ &\geq \sum_{\{u,v\} \in \mathcal{E}_t} \left(\|x_t(v) - x_t(u)\|_2^2 - \|x_{t+1}(v) - x_{t+1}(u)\|_2^2 \right) \end{aligned}$$

$$\begin{aligned} &= \sum_{u \in \mathcal{N}_t(v)} \left(\|x_t(v) - x_t(u)\|_2^2 - \|x_{t+1}(v) - x_{t+1}(u)\|_2^2 \right) \\ &= \|x_{t+1}(v) - x_t(v)\|_2^2 \\ & \quad + \sum_{u \in \mathcal{N}_t(v)} \left(\|x_t(v) - x_t(u)\|_2^2 \right. \\ & \quad \left. - \|x_{t+1}(v) - x_t(u)\|_2^2 \right) \\ &= \|m_t(v)\|_2^2 \\ & \quad + \sum_{u \in \mathcal{N}_t(v)} \left(\|x_t(v) - x_t(u)\|_2^2 \right. \\ & \quad \left. - \|x_t(v) + m_t(v) - x_t(u)\|_2^2 \right), \end{aligned}$$

Using the definition of $\|\cdot\|$, we obtain

$$\begin{aligned} &= \|m_t(v)\|_2^2 \\ & \quad + \sum_{u \in \mathcal{N}_t(v)} (x_t(v)^\top x_t(v) - 2x_t(v)^\top x_t(u) \\ & \quad + x_t(u)^\top x_t(u) \\ & \quad - (x_t(v) + m_t(v))^\top (x_t(v) + m_t(v)) \\ & \quad + 2(x_t(v) + m_t(v))^\top x_t(u) - x_t(u)^\top x_t(u)) \\ &= \|m_t(v)\|_2^2 \\ & \quad + \sum_{u \in \mathcal{N}_t(v)} (-2m_t(v)^\top x_t(v) - m_t(v)^\top m_t(v) \\ & \quad + 2m_t(v)^\top x_t(u)) \\ &= \|m_t(v)\|_2^2 - |\mathcal{N}_t(v)| \|m_t(v)\|_2^2 \\ & \quad + 2m_t(v)^\top \left(\sum_{u \in \mathcal{N}_t(v)} (x_t(u) - x_t(v)) \right) \\ &= \|m_t(v)\|_2^2 - |\mathcal{N}_t(v)| \|m_t(v)\|_2^2 \\ & \quad + 2|\mathcal{N}_t(v)| m_t(v)^\top m_t(v) \\ &= (|\mathcal{N}_t(v)| + 1) \|m_t(v)\|_2^2, \end{aligned}$$

which concludes the proof. \square

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

Lemma 8 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 7 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 Lemma 6, 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

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

□

The proof of Theorem 1 is a direct consequence of Lemma 8.

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 8 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 .

Applying the classic additive drift theorem (see, e.g., [44, Theorem 2.3.1] and the historic references therein) we directly observe that the expected number of steps to reach a δ -stable state is upper bounded by

$$\frac{\Phi(S)}{\frac{2\delta^2}{n|E|}} \leq \frac{|E|\varepsilon^2}{\frac{2\delta^2}{n|E|}} = \frac{n|E|^2}{2} \left(\frac{\varepsilon}{\delta}\right)^2,$$




Fig. 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 available 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

resulting in the bound from the theorem. □

Our results in Theorem 1 directly improve the results from Etesami and Başar [33] 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 $O(n^9(\varepsilon/\delta)^2)$.

Our result transfers to their notion of convergence as follows. Assume $\delta \leq \varepsilon/2$ and that the length of the longest edge is at most δ . If the social network is the complete graph, each connected component in the influence network must be a complete sub-graph, see Lemma 10. Hence, the diameter of this connected component is also bounded by δ . 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 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 9 There is a family of instances and initial states with $|E| = \Theta(n^2)$ and 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 $n \in \mathbb{N}_{>1}$, 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 available 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

$$\begin{aligned} |m_0(\ell)| &= \frac{n \cdot x_0(v) + x_0(v_1) - (n + 1)x_0(\ell)}{n + 2} \\ &= \frac{x_0(\ell) + \varepsilon - 3(n - 1)\hat{m} - (n + 1)x_0(\ell)}{n + 2} \\ &= \frac{\varepsilon - 3(n - 1)\hat{m} - n \cdot x_0(\ell)}{n + 2} \\ &= \frac{\varepsilon - 3(n - 1)\hat{m} - n \cdot \hat{m} \cdot (n + 1)}{n + 2} \\ &= \frac{\varepsilon - (n^2 + 4n - 3)\hat{m}}{n + 2} \\ &= \frac{\varepsilon - (n^2 + 4n - 3) \cdot \varepsilon / (n^2 + 5n - 1)}{n + 2} \\ &= \varepsilon \cdot \frac{(n^2 + 5n - 1) - (n^2 + 4n - 3)}{(n^2 + 5n - 1)(n + 2)} \\ &= \varepsilon \cdot \frac{n + 2}{(n^2 + 5n - 1)(n + 2)} \\ &= \hat{m} . \end{aligned}$$

The calculation of the available movement of the other agents is analogous. Let $v \in C_\ell$, then it holds that

$$\begin{aligned} |m_0(v)| &= \frac{x_0(\ell) - x_0(v)}{n + 1} + \sum_{v' \in C_\ell} \frac{x_0(v') - x_0(v)}{n + 1} \\ &= \frac{x_0(\ell)}{n + 1} = \frac{\hat{m} \cdot (n + 1)}{n + 1} = \hat{m} . \end{aligned}$$

Let $v \in C_r$, then we have

$$\begin{aligned} |m_0(v)| &= \left| \frac{x_0(r) - x_0(v)}{n + 1} + \sum_{v' \in C_r} \frac{x_0(v') - x_0(v)}{n + 1} \right| \\ &= \left| \frac{x_0(r) - (x_0(r) + (n + 1)\hat{m})}{n + 1} \right| \\ &= \frac{\hat{m} \cdot (n + 1)}{n + 1} = \hat{m} . \end{aligned}$$

Let $v_i \in P$ with $i \leq n - 1$, then we get that

$$\begin{aligned} |m_0(v_i)| &= \frac{|x_0(v_{i-1}) + x_0(v_{i+1}) - 2x_0(v_i)|}{3} \\ &= \frac{|x_0(v_{i-1}) + \varepsilon - 3(n - (i + 1))\hat{m} - x_0(v_i)|}{3} \\ &= \frac{|\varepsilon - 3(n - i)\hat{m} - (\varepsilon - 3(n - (i + 1))\hat{m})|}{3} \\ &= \frac{3\hat{m}}{3} = \hat{m} . \end{aligned}$$

Let $v_n \in P$, then it holds that

$$\begin{aligned} |m_0(v_n)| &= \frac{|x_0(v_{n-1}) + x_0(v_{n+1}) - 2x_0(v_n)|}{3} \\ &= \frac{|x_0(v_{n-1}) + \varepsilon - 3(n + 1 - n)\hat{m} - x_0(v_n)|}{3} \\ &= \frac{|\varepsilon - 3(n + 1 - n)\hat{m} - (\varepsilon - 3(n - n)\hat{m})|}{3} \\ &= \frac{3\hat{m}}{3} = \hat{m} . \end{aligned}$$

Let $v_i \in P$ with $i > n$, then we have

$$\begin{aligned} |m_0(v_i)| &= \frac{|x_0(v_{i-1}) + x_0(v_{i+1}) - 2x_0(v_i)|}{3} \\ &= \frac{|x_0(v_{i-1}) + \varepsilon - 3(i + 1 - n)\hat{m} - x_0(v_i)|}{3} \\ &= \frac{|\varepsilon - 3(i + 1 - n)\hat{m} - (\varepsilon - 3(i - n)\hat{m})|}{3} \\ &= \frac{3\hat{m}}{3} = \hat{m} . \end{aligned}$$

Finally, we get that

$$\begin{aligned}
 |m_0(r)| &= \left| \sum_{v \in C_r \cup v_{2n-2}} \frac{x_0(v) - x_0(r)}{n+2} \right| \\
 &= \frac{(n+1)x_0(r)}{n+2} \\
 &\quad - \frac{n \cdot (x_0(r) + (n+1)\hat{m})}{n+2} \\
 &\quad - \frac{(x_0(r) - \varepsilon + 3(n-1)\hat{m})}{n+2} \\
 &= \frac{\varepsilon - 3(n-1)\hat{m} - n(n+1)\hat{m}}{n+2} \\
 &= \frac{\varepsilon - (n^2 + 4n - 3)\hat{m}}{n+2} \\
 &= \frac{\varepsilon - (n^2 + 4n - 3) \cdot \varepsilon / (n^2 + 5n - 1)}{n+2} \\
 &= \varepsilon \cdot \frac{(n^2 + 5n - 1) - (n^2 + 4n - 3)}{(n^2 + 5n - 1)(n+2)} \\
 &= \varepsilon \cdot \frac{n+2}{(n^2 + 5n - 1)(n+2)} \\
 &= \hat{m}.
 \end{aligned}$$

By Lemma 7, the expected potential drop is given by

$$\begin{aligned}
 \mathbb{E}[\Phi(S_0) - \Phi(S_1)] &= \frac{1}{4n} \sum_{v \in V} (|\mathcal{N}_0(v)| + 1) \cdot |m_0(v)|^2 \\
 &= \frac{1}{4n} \left(\frac{n}{2} \left(\frac{n}{4} + 2 \right) + 2 \left(\frac{n}{4} + 3 \right) + \left(\frac{n}{2} - 2 \right) \cdot 4 \right) \hat{m}^2 \\
 &= \frac{1}{4} (n/8 + 7/2 - 2/n) \hat{m}^2 \\
 &= \frac{1}{4} (n/8 + 7/2 - 2/n) (\varepsilon / (n^2/16 + 5n/4 - 1))^2 \\
 &= \Theta(\varepsilon^2/n^3).
 \end{aligned}$$

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

Note that this only proves that the drop in the first step is sufficiently small. The expected drop for the next step could increase after activating a node. Therefore, this theorem does not prove a lower bound for the convergence time. Instead, it shows that the analysis of the step-by-step drop is tight.

In the next section, we see an example of how to change the analysis to circumvent this bound.

3 Special 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 when the social network is a complete graph, while the second holds when, in each step of the HKS, the influence network is the same as the social network.

To prove the result for HKSs with a social network, we prove the following characteristics of these systems.

Lemma 10 *Let $(G = (V, E), \varepsilon, x)$ be any instance of a d -dimensional HKS with complete social network $G = K_n$ and current influence network I . If all edges in I have a length of at most $\varepsilon/2$, each connected component in I is a complete graph.*

Proof Let $V' \subseteq V$ be the set of nodes of a connected component in I . Assume that $v, u \in V'$, 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 can be at most ε . Therefore, 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

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 8, we have

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

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

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 a length at least δ . We project the HKS to one dimension along the longest edge. By Lemma 4, we know that in the projected graph, no edge increases its length, and there exists an edge with a length at least δ .

Let there be k connected components $C_i = (V_i, E_i), i \in \{1, \dots, k\}$, and $\lambda_i(t)$ be the length of the longest edge in

the connected component C_i . We bound the total available movement in this component from below using Lemma 5.

For each connected component C_i with $\lambda_i(t) > 0$ let $e_i = \{u, w\}$ be a longest edge in this component. 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 $E_{\ell,r,i}$ as in Lemma 5. Since the connected component is the complete graph by Lemma 10 each node from $V_{i,\ell}$ is connected to w while each node from $V_{i,r}$ is connected to u . As a consequence, the set $E_{\ell,r,i}$ contains at least $(|V_i| - 1)$ edges of length at least $\lambda_t/2$ and one of them has length λ_t . As a consequence, $\sum_{e \in E_{\ell,r,i}} \|x_t(e)\|_2 \geq |V_i| \lambda_t/2$ and hence, by Lemma 5,

$$\begin{aligned} |V_i| \sum_{v \in V_i} |m_t(v)| &= \sum_{v \in V_i} |\mathcal{N}_t(v)| |m_t(v)| \\ &\geq 2 \sum_{e \in E_{\ell,r,i}} \|x_t(e)\|_2 \geq |V_i| \lambda_i(t) \end{aligned}$$

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 a 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 an HKS 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 of such graphs are the path, where all the nodes are positioned with equal distance of at most ε , as well as the graph from Theorem 9 if the social network for the latter is reduced to the set of edges in \mathcal{E}_0 .

Theorem 11 *Let $(G = (V, E), \varepsilon, x)$ be a HKS 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 8, 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. We prove that there exists a family of instances of HKS $(S_{0,4n})_{n \in \mathbb{N}}$ for which at least $\Omega(\Phi(S_{0,4n})n|E|)$ updates are needed to converge. In this family, we have $|E| = \Theta(n^2)$ and $\Phi(S_0) = \Theta(n)$ for $\varepsilon = 1$ and $\delta < 1/2$. Note that the lower bound holds for any given update sequence. In particular, we also prove that there cannot be a deterministic algorithm that reaches a δ -stable state faster than the proven lower bound.

Theorem 3 *For $\varepsilon = 1$ and $\delta < 1/2$, there exists a family of social 1-dimensional HKSs $(S_{0,4n})_{n \in \mathbb{N}}$ where any given update sequence needs at least $\Omega(\Phi(S_{0,4n})n|E|/\varepsilon^2) = \Omega(n^4)$ steps to reach a δ -stable state.*

Note that this lower bound is tight with regard to the considered family of instances and the parameter n since Theorem 1 states that in expectation at most $O((\Phi(S_{0,4n})n|E|)/\delta^2) = O(n^4)$ updates are needed, to reach a δ -stable state.

We prove this lower bound in three steps. First, we prove that edges cannot be deactivated in the defined family of instances. Then, we upper bound the total available movement of certain nodes in the family by a modified process that is simpler to analyze.

Dumbbell Graph We define the following family of instances of HKS $(S_{0,4n})_{n \in \mathbb{N}} = (G_{4n} = (V, E), \varepsilon, x_0)_{n \in \mathbb{N}}$: Let $G_{4n} = (V, E)$ be the Dumbbell graph defined as follows. The $|V| = 4n$ nodes of the graph are partitioned into the sets $C_\ell, C_r, P, \{\ell, r\} \subseteq V$, such that $|C_\ell| = |C_r| = n, |P| = 2n - 2$.

The set of edges are given such that $C_\ell \cup \{\ell\}$ and $C_r \cup \{r\}$ are cliques, $P = \{p_1, \dots, p_{2n-2}\}$ is a path such that $\{p_i, p_{i+1}\} \in E$ for each $i \in \{1, \dots, 2n - 2\}$, as well as $\{\ell, p_1\}, \{r, p_{2n-2}\} \in E$. In state $S_{0,4n}$ the opinions of the agents that correspond to the nodes V are given as follows:

- $x_0(v) = 0$ for each $v \in C_\ell$
- $x_0(\ell) = \varepsilon/n$
- $x_0(p_i) = i\varepsilon + \varepsilon/n$ for each node $p_i \in P$
- $x_0(r) = (2n - 1)\varepsilon + \varepsilon/n$
- $x_0(v) = (2n - 1)\varepsilon + 2\varepsilon/n$ for each $v \in C_r$.

Note that $\Phi(S_{0,4n}) \in \Omega(n\varepsilon^2)$ and $|E| = n^2$, therefore $\Omega(\Phi(S_0)n|E|/\delta^2) = \Omega(n^4)$.

Lemma 12 *When starting the HKS process with $S_{0,4n}$, it can never happen that an edge is deactivated during the process.*

Proof Assume we activate node $p_i \in P$ in step $t + 1$ and that the edges $\{p_{i-1}, p_i\}$ and $\{p_i, p_{i+1}\}$ are active as well as $x_t(p_{i-1}) \leq x_t(p_i) \leq x_t(p_{i+1})$. Which means that $x_t(p_i) - x_t(p_{i-1}) \leq \varepsilon$ and $x_t(p_{i+1}) - x_t(p_i) \leq \varepsilon$ and hence $|x_t(p_{i+1}) - x_t(p_{i-1})| \leq 2\varepsilon$. As a consequence,

$$\begin{aligned} & x_{t+1}(p_i) - x_{t+1}(p_{i-1}) \\ &= \frac{1}{3}(x_t(p_i) + x_t(p_{i+1}) + x_t(p_{i-1})) - x_t(p_{i-1}) \\ &= \frac{1}{3}((x_t(p_i) - x_t(p_{i-1})) + (x_t(p_{i+1}) - x_t(p_{i-1}))) \\ &\leq \frac{1}{3}(\varepsilon + 2\varepsilon) = \varepsilon \end{aligned}$$

and similarly $x_t(p_{i+1}) - x_{t+1}(p_i) \leq \varepsilon$. Furthermore, we obviously have that $x_t(p_{i-1}) \leq \frac{1}{3}(x_t(p_i) + x_t(p_{i+1}) + x_t(p_{i-1})) \leq x_t(p_{i+1})$ and hence $x_{t+1}(p_{i-1}) \leq x_{t+1}(p_{i+1}) \leq x_{t+1}(p_i)$

For each pair of nodes $v, v' \in C_\ell$ and any time t , we will prove that $|x_t(v') - x_t(v)| \leq \varepsilon/n$, that $|x_t(\ell) - x_t(v)| \leq 2\varepsilon/n$, and that $|x_t(\ell) - x_t(p_1)| \leq \varepsilon$. This claim is true in the start configuration since all nodes in C_ℓ have the same position, and ℓ has distance ε/n to the nodes in C_ℓ .

Assume we activate a node $v \in C_\ell$ in step $t + 1$ and $|x_t(v') - x_t(v'')| \leq \varepsilon/n$ as well as $|x_t(\ell) - x_t(v')| \leq 2\varepsilon/n$ for each pair of nodes $v', v'' \in C_\ell \cup \{\ell\}$. Then it holds for any node $v' \in C_\ell \cup \{\ell\} \setminus \{v\}$ that

$$\begin{aligned} & |x_t(v') - x_{t+1}(v)| \\ &= |x_t(v') - \sum_{v'' \in C_\ell \cup \{\ell\}} x_t(v'') / (|C_\ell \cup \{\ell\}|)| \\ &\leq \sum_{v'' \in C_\ell \cup \{\ell\} \setminus \{v'\}} |x_t(v') - x_t(v'')| / (n + 1) \\ &\leq ((n - 1)(\varepsilon/n) + 2\varepsilon/n) / (n + 1) = \varepsilon/n. \end{aligned}$$

Furthermore, it holds that

$$\begin{aligned} & |x_t(\ell) - x_{t+1}(v)| \\ &= |x_t(\ell) - \sum_{v'' \in C_\ell \cup \{\ell\}} x_t(v'') / (|C_\ell \cup \{\ell\}|)| \end{aligned}$$

$$\begin{aligned} & \leq \sum_{v'' \in C_\ell} |x_t(v') - x_t(v'')| / (n + 1) \\ & \leq (n \cdot (\varepsilon/n)) / (n + 1) \leq 2\varepsilon/n. \end{aligned}$$

Assume we activate the node ℓ in step $t + 1$ and for each pair of nodes $v, v' \in C_\ell$ it holds that $|x_t(v') - x_t(v)| \leq \varepsilon/n$, that $|x_t(\ell) - x_t(v)| \leq 2\varepsilon/n$, and that $|x_t(\ell) - x_t(p_1)| \leq \varepsilon$. Then it holds for any node $v' \in C_\ell$ that

$$\begin{aligned} & |x_t(v') - x_{t+1}(\ell)| \\ &= |x_t(v') - \sum_{v'' \in C_\ell \cup \{\ell, p_1\}} x_t(v'') / (|C_\ell \cup \{\ell, p_1\}|)| \\ &\leq \sum_{v'' \in C_\ell \cup \{\ell, p_1\} \setminus \{v'\}} |x_t(v') - x_t(v'')| / (n + 2) \\ &\leq ((n - 1)(\varepsilon/n) + 2\varepsilon/n + \varepsilon) / (n + 2) \\ &= (2\varepsilon + \varepsilon/n) / (n + 2) \leq 2\varepsilon/n \end{aligned}$$

Furthermore, it holds that

$$\begin{aligned} & |x_t(p_1) - x_{t+1}(\ell)| \\ &= |x_t(p_1) - \sum_{v'' \in C_\ell \cup \{\ell, p_1\}} x_t(v'') / (|C_\ell \cup \{\ell, p_1\}|)| \\ &\leq \sum_{v'' \in C_\ell \cup \{\ell\}} |x_t(p_1) - x_t(v'')| / (n + 2) \\ &\leq (n(\varepsilon/n + \varepsilon) + \varepsilon) / (n + 2) = \varepsilon. \end{aligned}$$

Analogously, we can prove the statement for the nodes in $C_r \cup \{r\}$.

Hence, no edge can disappear when updating any node in the graph at any step t . \square

A necessary condition to terminate is that each edge on the path P has a length of at most δ . Since the distance between ℓ and r is $(2n - 1)\varepsilon$ at state $S_{0,4n}$ and can be at most $(2n - 1)\delta$ in the final state, one of the nodes ℓ and r has to move by at least $(\varepsilon - \delta)(2n - 1)/2$.

Modified Update Process In the following, we will lower-bound the number of required update steps to move ℓ by this distance by defining a modified update process that moves ℓ faster. However, the required number of steps in the modified process can be analyzed more easily. The opinions in the modified process are denoted by $x'_t(v)$ for each $v \in V$ and $t \in \mathbb{N}$. Let $S_t = \sum_{v \in C_\ell \cup \ell} x_t(v) / (n + 1)$ and $S'_t = \sum_{v \in C_\ell \cup \ell} x'_t(v) / (n + 1)$ denote the mass center of the nodes $C_\ell \cup \ell$ in the original and modified process respectively. Note that when activating an agent from C_ℓ in each of the two processes, it will move to the corresponding center of mass.

Consider any update-sequence $(u_t)_{t \in \mathbb{N}_{\geq 0}}$ specifying for any $t \in \mathbb{N}_{\geq 0}$ the node u_t that will be updated to generate

the positions x_{t+1} from x_t . Assume for simplicity of notation, that the nodes in C_ℓ are numbered from one to n , i.e., $C_\ell = \{v_1, \dots, v_n\}$. Given $(u_t)_{t \in \mathbb{N}_{\geq 0}}$, the modified process is defined as follows:

- $x'_0 = x_0$.
- If $u_t \notin C_\ell$ do nothing, i.e., set $x'_{t+1} = x'_t$.
- If $u_t \in C_\ell$, update multiple nodes at once. Generate x'_{t+1} as follows:

- find the node $v_{\min} \in C_\ell$ that has the smallest value $x'_t(v)$ with the smallest index, and set

$$x'_{t+1}(v_{\min}) = S'_t;$$

- move the node ℓ to the right: set

$$x'_{t+1}(\ell) = \frac{\varepsilon}{n} + \sum_{v \in C_\ell} \frac{x'_{t+1}(v)}{n};$$

- if $v_{\min} = v_n$, move all the nodes in C_ℓ to S'_t , i.e., for all $v \in C_\ell$ set

$$x'_{t+1}(v) = S'_t.$$

We prove Theorem 3 in two steps. First, we prove that, indeed, the modified process moves the node ℓ faster than the original process. Afterward, we prove that the modified process needs at least $\Omega(n^4)$ steps to move the clique by the distance $(\varepsilon - \delta)(2n - 1)/2$.

Lemma 13 *For each update sequence $(u_t)_{t \in \mathbb{N}_{>0}}$ there exists for each $t \in \mathbb{N}$ a bijection $f_t : C_\ell \rightarrow C_\ell$ such that $x_t(v) \leq x'_t(f(v))$ for each $v \in C_\ell$. Furthermore in each step it holds that $x_t(\ell) \leq x'_t(\ell)$ and $S_t \leq S'_t$.*

Proof First, note that the modified update process will update the nodes C_ℓ always in the same order from v_1 to v_n since it chooses the node with a smallest opinion and the smallest index. Since, in the beginning, all the nodes have the same opinion, and the opinion only increases when updated, the first n updates will update the nodes in the claimed order. After updating the node v_n , all the nodes will be shifted to the same (increased) opinion. Inductively, the nodes are activated in the same order.

We will prove the claim via induction on the number of steps. Before activating any node, the bijection $f_0 : C_\ell \rightarrow C_\ell, v \mapsto v$ fulfills the required properties, since $x_0(v) = x'_0(v) = 0$. Furthermore it holds that $x_0(\ell) = x'_0(\ell) = \varepsilon/n$ and

$$S_0 = S'_0 = \sum_{v \in C_\ell \cup \ell} x_0(v)/(n+1) = \frac{\varepsilon}{n(n+1)}.$$

Assume that after step t there is a bijection $f_t : C_\ell \rightarrow C_\ell$ such that $x_t(v) \leq x'_t(f_t(v))$ for each $v \in C_\ell$, $S_t \leq S'_t$, and $x_t(\ell) \leq x'_t(\ell)$. We have to consider the following cases: the updated node is in C_ℓ , the updated node is ℓ , the modified process updates the node v_n , and a node not in $C_\ell \cup \{\ell\}$ is updated.

In the case that a node in $V \setminus (C_\ell \cup \{\ell\})$ is updated, the values $x_t(\ell), x'_t(\ell), S_t, S'_t$ stay unchanged. Furthermore, in both processes, no node from C_ℓ will be updated and hence we set the update function $f_{t+1} = f_t$.

If the node ℓ is updated, the modified process will not move the node ℓ . Note that in the original process, the edge connecting the node ℓ with the node p_1 can have length at most ε , since no edge is deactivated when updating any node and hence $x_t(p_1) \leq x_t(\ell) + \varepsilon$. As a consequence, it holds that

$$\begin{aligned} x_{t+1}(\ell) &= \sum_{v \in C_\ell \cup \{\ell, p_1\}} x_t(v)/(n+2) \\ &\leq ((n+1)S_t + (x_t(\ell) + \varepsilon))/(n+2) \\ &\leq ((n+1)S'_t + (x'_t(\ell) + \varepsilon))/(n+2) \\ &= (\sum_{v \in C_\ell} x'_t(v) + 2x'_t(\ell) + \varepsilon)/(n+2) \\ &= \frac{(n+2) \sum_{v \in C_\ell} x'_t(v)/n + (n+2)\varepsilon/n}{(n+2)} \\ &= \sum_{v \in C_\ell} x'_t(v)/n + \varepsilon/n = x'_t(\ell) \end{aligned}$$

and hence $x_{t+1}(\ell) \leq x'_{t+1}(\ell)$. Since in both processes no node from C_ℓ was updated, for the bijection $f_t : C_\ell \rightarrow C_\ell$ with $x_t(v) \leq x'_t(f(v))$ for each $v \in C_\ell$, it still holds that $x_{t+1}(v) \leq x'_{t+1}(f(v))$ and hence we can set $f_{t+1} = f_t$. As a consequence $S_{t+1} \leq S'_{t+1}$ has to hold as well.

Let us assume that we update the clique node $v_{t+1} \in C_\ell$ in the next step in the original process and the node $v'_{t+1} \in C_\ell$ in the modified process. Note that $x_{t+1}(v_{t+1}) = S_t$ and $x'_{t+1}(v'_{t+1}) = S'_t$ by definition of the update process. Furthermore

$$S_{t+1} = S_t + (x_{t+1}(v_t) - x_t(v_t))/(n+1).$$

If $f_t(v_{t+1}) = v'_{t+1}$, we know that $x_{t+1}(v_{t+1}) = S_t \leq S'_t = x'_{t+1}(v'_{t+1}) = x'_{t+1}(f_t(v_{t+1}))$. Since no other node from C_ℓ is moved, we define $f_{t+1} = f_t$ and it holds that $x_{t+1}(v) \leq x'_{t+1}(f(v))$ for each $v \in C_\ell$. Furthermore,

$$\begin{aligned} S_{t+1} &= \sum_{v \in C_\ell \cup \{\ell\}} x_{t+1}(v)/(n+1) \\ &\leq x'_{t+1}(\ell)/(n+1) + \sum_{v \in C_\ell} x'_{t+1}(f_{t+1}(v))/(n+1) \end{aligned}$$

$$\begin{aligned}
 &= \sum_{v \in C_\ell \cup \{\ell\}} x'_{t+1}(v)/(n+1) \\
 &= S'_{t+1}.
 \end{aligned}$$

If $f_t(v_{t+1}) \neq v'_{t+1}$, consider the nodes $f_t(v_{t+1})$ and $f_t^{-1}(v'_{t+1})$. We know that $x_t(v_{t+1}) \leq x'_t(f_t(v_{t+1}))$ and $x_t(f_t^{-1}(v'_{t+1})) \leq x'_{v'_{t+1}}$. Furthermore, by the choice of v'_{t+1} , we know that $x'_{v'_{t+1}} \leq x'_t(f_t(v_{t+1}))$. As a consequence, it holds that $x_t(f_t^{-1}(v'_{t+1})) \leq x'_t(f_t(v_{t+1}))$. When updating the nodes, we get $x_{t+1}(v_{t+1}) = S_t \leq S'_t = x'_{t+1}(v'_{t+1})$. We define

$$f_{t+1}(v) = \begin{cases} v'_{t+1} & \text{if } v = v_{t+1} \\ f_t(v_{t+1}) & \text{if } v = f_t^{-1}(v'_{t+1}) \\ f_t(v) & \text{otherwise} \end{cases}$$

Note that for each $v \in C_\ell$ it holds that $x_{t+1}(v) \leq x'_{t+1}(f_{t+1}(v))$. Furthermore, we have

$$x_{t+1}(\ell) = x_t(\ell) \leq x'_t(\ell) \leq x'_{t+1}(\ell)$$

and

$$\begin{aligned}
 S_{t+1} &= \sum_{v \in C_\ell \cup \{\ell\}} x_{t+1}(v)/(n+1) \\
 &\leq (x'_{t+1}(\ell) + \sum_{v \in C_\ell} x'_{t+1}(f_{t+1}(v)))/(n+1) \\
 &= S'_{t+1}.
 \end{aligned}$$

Finally, if the updated node in the modified processes was v_n , all the nodes in C_ℓ move to the point S'_t . Since $S_t \leq S'_t$, the identity function $f_{t+1} = id$ fulfills the required conditions, since S'_t is monotonically increasing and hence there can be no node with $x_t(v) > S'_t$. Similarly as above it follows that $x_{t+1}(\ell) \leq x'_{t+1}(\ell)$ and $S_{t+1} \leq S'_{t+1}$. \square

Lemma 14 *The modified process needs at least $\Omega(n^4)$ updates before ℓ has moved by at least $(\varepsilon - \delta)(2n - 1)/2$.*

Proof In the modified process, the nodes in C_ℓ are shifted to a common position every n th update step. We define an update round as n updates of nodes in C_ℓ such that, before the first update of the round, all nodes have the same position, and all nodes are shifted to the same position after the last update of the round.

First, we prove that in each update round, the node ℓ moves by at most $\frac{(1+\varepsilon)\varepsilon}{n(n+1)}$. Let p_0 be the common start position of all nodes in C_ℓ . By definition of the modified process, the node ℓ has position $p_0 + \varepsilon/n$. We denote by $S'_t = \sum_{v \in C_\ell \cup \ell} x'_t/(n+1)$ the mass center of the nodes $C_\ell \cup \ell$ and by $\bar{S}'_t = \sum_{v \in C_\ell} x'_t/n$ the mass-center of the clique nodes C_ℓ .

When updating a node $u_{t+1} \in C_\ell$ at step $t + 1$, it moves from position p_0 to position S'_t since in each round, only non-updated nodes will be moved. This update increases the mass-center of the clique \bar{S}'_t to

$$\bar{S}'_{t+1} = \bar{S}'_t + (S'_t - p_0)/n.$$

After updating the node u_{t+1} , the node ℓ is shifted to $x'_t(\ell) = \bar{S}'_{t+1} + \varepsilon/n$, increasing the mass-center of the nodes $C_\ell \cup \ell$ from S_t to

$$\begin{aligned}
 S'_{t+1} &= (\bar{S}'_{t+1} \cdot n + \bar{S}'_{t+1} + \varepsilon/n)/(n+1) \\
 &= \bar{S}'_{t+1} + \frac{\varepsilon}{n(n+1)}.
 \end{aligned}$$

As a consequence, we get that

$$\begin{aligned}
 \bar{S}'_{t+1} &= \bar{S}'_t + (S'_t - p_0)/n \\
 &= \bar{S}'_t + \frac{\bar{S}'_t + \frac{\varepsilon}{n(n+1)} - \bar{S}'_0}{n} \\
 &= \frac{n+1}{n} \bar{S}'_t - \frac{\bar{S}'_0}{n} + \frac{\varepsilon}{n^2(n+1)}
 \end{aligned}$$

Via induction, we show that

$$\bar{S}'_t = \bar{S}'_0 + \left(\left(\frac{n+1}{n} \right)^t - 1 \right) \cdot \frac{\varepsilon}{n(n+1)}$$

It holds that

$$\begin{aligned}
 \bar{S}'_1 &= \frac{n+1}{n} \bar{S}'_0 - \frac{\bar{S}'_0}{n} + \frac{\varepsilon}{n^2(n+1)} \\
 &= \bar{S}'_0 + \frac{\varepsilon}{n^2(n+1)} \\
 &= \bar{S}'_0 + \left(\left(\frac{n+1}{n} \right)^1 - 1 \right) \cdot \frac{\varepsilon}{n(n+1)}
 \end{aligned}$$

Additionally, it holds that

$$\begin{aligned}
 \bar{S}'_{t+1} &= \frac{n+1}{n} \bar{S}'_t - \frac{\bar{S}'_0}{n} + \frac{\varepsilon}{n^2(n+1)} \\
 &= \frac{n+1}{n} \left(\bar{S}'_0 + \left(\left(\frac{n+1}{n} \right)^t - 1 \right) \cdot \frac{\varepsilon}{n(n+1)} \right) \\
 &\quad - \frac{\bar{S}'_0}{n} + \frac{\varepsilon}{n^2(n+1)} \\
 &= \bar{S}'_0 + \left(\left(\frac{n+1}{n} \right)^{t+1} - \frac{n+1}{n} + \frac{1}{n} \right) \frac{\varepsilon}{n(n+1)} \\
 &= \bar{S}'_0 + \left(\left(\frac{n+1}{n} \right)^{t+1} - 1 \right) \cdot \frac{\varepsilon}{n(n+1)}
 \end{aligned}$$

As a consequence

$$\begin{aligned} \bar{S}'_{n-1} &= \bar{S}'_0 + \left(\left(\frac{n+1}{n} \right)^{n-1} - 1 \right) \cdot \frac{\varepsilon}{n(n+1)} \\ &< \bar{S}'_0 + \frac{(e-1)\varepsilon}{n(n+1)} \end{aligned}$$

In the last update step of the round, all nodes are moved to

$$S'_{n-1} = \bar{S}'_{n-1} + \frac{\varepsilon}{n(n+1)} \leq \bar{S}'_0 + \frac{e\varepsilon}{n(n+1)}$$

resulting in $\bar{S}'_n \leq \bar{S}'_0 + \frac{e\varepsilon}{n(n+1)}$. Consequently ℓ is located at or left of position $\bar{S}'_0 + \frac{e\varepsilon}{n(n+1)} + \frac{\varepsilon}{n}$ and hence has moved by at most $\frac{e\varepsilon}{n(n+1)}$ in this round.

Note that after moving each node once, the relative positioning between the nodes in $C_\ell \cup \ell$ is the same as before moving the first node. Hence, it is sufficient to consider the total number of rounds of moving n nodes from the clique.

The node ℓ has to move by $(\varepsilon - \delta)(n - 1)$, but after activating each node of the clique, it has moved by at most $\frac{e\varepsilon}{n(n+1)}$. Hence, we need

$$\frac{(\varepsilon - \delta)(n - 1)}{\frac{e\varepsilon}{n(n+1)}} = \Omega\left(n^3 \frac{\varepsilon - \delta}{\varepsilon}\right) = \Omega(n^3)$$

of these rounds, for δ small enough. Since each round needs n updates, at least $\Omega(n^4)$ updates are needed until the modified process terminates. \square

Proof of Theorem 3 Consider the family $(S_{0,4n})_{n \in \mathbb{N}}$ as defined in this section. By Lemma 12, no edge will be deactivated until the process terminates. As a consequence, for termination, each edge on the path can have a length of at most δ . This implies the the distance between ℓ and r has to shrink from $(2n - 1)\varepsilon$ to $(2n - 1)\delta$. Therefore, one of the nodes ℓ or r must move by at least $(n - 1/2)(\varepsilon - \delta)$ until the process terminates. Let us w.l.o.g. assume that ℓ moves by this distance.

Consider the modified process as defined in this chapter. By Lemma 13 we know that $x_t(\ell) \leq x'_t(\ell)$ for any update step t . Furthermore, by Lemma 14 we know, that at least $t \in \Omega(n^4)$ steps are needed before $x'_t(\ell) - x'_0(\ell) \geq (n - 1/2)(\varepsilon - \delta)$. Consequently, the original process needs at least $\Omega(n^4)$ update steps until it terminates. \square

5 Simulation results

To corroborate our theoretical findings, we performed agent-based simulations of asynchronous Hegselmann–Krause opinion dynamics in one dimension on two types of initial

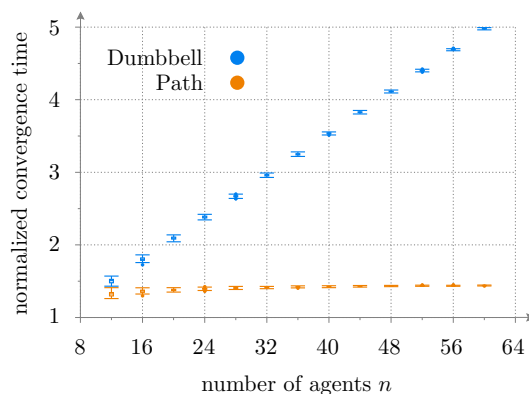


Fig. 2 The plot shows the normalized convergence time: the number of agent activations until a δ -stable state has been reached, divided by n^3 . The data indicate that the convergence time on *Path* instances with equal distances scales as n^3 and on *Dumbbell* instances it scales as n^4

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 with an equal distance of ε so that the influence network forms a path graph with a uniform edge length of ε .
- *Dumbbell*: This is the state constructed in the proof of Theorem 9 using the dumbbell graph, except that the social network contains only the edges that are in \mathcal{E}_0

We fixed $\varepsilon = 100$ and $\delta = 1$ in our simulations. 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. The code for our simulator software and all necessary tools to reproduce our figures are available from our public GitHub repository.¹

We present our simulation results in Fig. 2.

There, the obtained number of activations divided by n^3 is plotted via box plots that summarize the results for each configuration. Since for *Path* instances, the number of activations appears to be constant, we observe that we need $\Theta(n^3)$ activations for *Path* instances. On the other hand, the number of activations seems to grow linearly in n for *Dumbbell* instances. This matches our proofs (upper and lower bound) that in $\Theta(n^4)$ activations *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, the empirically

¹ <https://github.com/dcmx/HKsim>

observed lower bounds on the expected number of steps until convergence match our theoretical analysis for these two graph classes with respect to the dependence on the number of agents.

We highlight that our simulations not only confirm our theoretical results but also allow us to empirically pinpoint the constants hidden in the asymptotic analysis. For the simulations carried out on the path our simulation results indicate a constant of roughly 1.5, giving a total running time of roughly $1.5 \cdot n^3$. For the dumbbell we get a running time of less than $0.1 \cdot n^4$. We remark that in our theoretical analysis we did not attempt to optimize the constants. In fact, our result for the dumbbell graph (see the proof of Theorem 1) gives a bound of at most $\log(100^2)/4 \cdot n^4 \approx 3.32 \cdot n^4$. Interestingly, the total running times are sharply concentrated around their mean. In fact, the plot in Fig. 2 actually shows box plots, but starting with a number of agents as small as only 40 the upper and lower whiskers become almost identical with no outliers detected.

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 $O(n|E|\Phi(S_0)/\delta^2) \leq O(n|E|^2(\varepsilon/\delta)^2)$ 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 the social network indicate that our theoretical upper bound is tight for these instances. For the dumbbell graph, this is not only underlined by our simulations but also proven by presenting a matching lower bound with respect to the number of agents. This 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 by finding a graph class with $\Phi(S_0) \in \Omega(n^2\varepsilon^2)$ with expected convergence time in $\Omega(n^5)$.

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 to 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, allowing us 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.

Funding Open Access funding enabled and organized by Projekt DEAL. This work was supported by DFG Research Group ADYN under grant DFG 411362735.

Declarations

Author contributions All authors worked jointly on writing, reviewing, and polishing the manuscript.

Data availability The source code for our simulator software and all necessary tools to reproduce the simulation results reported in Sect. 5 is available from our public GitHub repository (<https://github.com/dcmx/HKsim>).

Conflict of interest The authors declare no competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Abelson, R., Bernstein, A.: A computer simulation model of community referendum controversies. *Public Opin. Q.* **27**(1), 93–122 (1963)
2. Anagnostopoulos, A., Becchetti, L., Cruciani, E., Pasquale, F., Rizzo, S.: Biased opinion dynamics: when the devil is in the details. In: *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 53–59
3. Angluin, D., Aspnes, J., Eisenstat, D.: A simple population protocol for fast robust approximate majority. *Distrib. Comput.* **21**(2), 87–102 (2008)
4. Auletta, V., Fanelli, A., Ferraioli, D.: Consensus in opinion formation processes in fully evolving environments. In: *Proceedings of the 33rd Conference on Artificial Intelligence (AAAI)*, pp. 6022–6029 (2019)
5. Bankhamer, G., Berenbrink, P., Biermeier, F., Elsässer, R., Hosseinpour, H., Kaaser, D., Kling, P.: Fast consensus via the unconstrained undecided state dynamics. In: *Proceedings of the Symposium on Discrete Algorithms (SODA)*, pp. 3417–3429 (2022)
6. Becchetti, L., Clementi, A.E.F., Natale, E.: Consensus dynamics: an overview. *SIGACT News* **51**(1), 58–104 (2020)
7. Becchetti, L., Clementi, A.E.F., Natale, E., Pasquale, F., Silvestri, R.: Plurality consensus in the gossip model. In: *Proceedings of the 26th Symposium on Discrete Algorithms (SODA)*, pp. 371–390 (2015)
8. Becchetti, L., Clementi, A.E.F., Natale, E., Pasquale, F., Silvestri, R., Trevisan, L.: Simple dynamics for plurality consensus. *Distrib. Comput.* **30**(4), 293–306 (2017)

9. Berenbrink, P., Clementi, A.E.F., Elsässer, R., Kling, P., Mallmann-Trenn, F., Natale, E.: Ignore or comply?: on breaking symmetry in consensus. In: Proceedings of the Symposium on Principles of Distributed Computing (PODC), pp. 335–344 (2017)
10. Berenbrink, P., Friedetzky, T., Giakkoupis, G., Kling, P.: Efficient plurality consensus, or: the benefits of cleaning up from time to time. In: Proceedings of the 43rd International Colloquium on Automata, Languages and Programming (ICALP), pp. 136:1–136:14 (2016a)
11. Berenbrink, P., Giakkoupis, G., Kermarrec, A.-M., Mallmann-Trenn, F.: Bounds on the voter model in dynamic networks. In: Proceedings of the 43rd International Colloquium on Automata, Languages and Programming (ICALP), pp. 146:1–146:15 (2016b)
12. Bhattacharyya, A., Braverman, M., Chazelle, B., Nguyen, H.L.: On the convergence of the Hegselmann–Krause system. In: Kleinberg, R.D. (ed.), Proceedings of the Symposium on Innovations in Theoretical Computer Science (ITCS), pp. 61–66 (2013)
13. Bhattacharyya, A., Shiragur, K.: How friends and non-determinism affect opinion dynamics. In: Proceedings of the 54th IEEE Conference on Decision and Control (CDC). IEEE, pp. 6466–6471 (2015)
14. Bhawalkar, K., Gollapudi, S., Munagala, K.: Coevolutionary opinion formation games. In: Proceedings of the Symposium on Theory of Computing (STOC), pp. 41–50 (2013)
15. Bilò, V., Fanelli, A., Moscardelli, L.: Opinion formation games with dynamic social influences. *Theor. Comput. Sci.* **746**, 73–87 (2018)
16. Bindel, D., Kleinberg, J.M., Oren, S.: How bad is forming your own opinion? *Games Econ. Behav.* **92**, 248–265 (2015)
17. Botan, S., Grandi, U., Perrussel, L.: Multi-issue opinion diffusion under constraints. In: Proceedings of the 18th Conference on Autonomous Agents and Multi-agent Systems (AAMAS), pp. 828–836 (2019)
18. Brederbeck, R., Elkind, E.: Manipulating opinion diffusion in social networks. In: Sierra, C. (ed.), Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI), pp. 894–900 (2017)
19. Brederbeck, R., Jacobs, L., Kellerhals, L.: Maximizing the spread of an opinion in few steps: opinion diffusion in non-binary networks. In: Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI), pp. 1622–1628 (2020)
20. Chazelle, B.: The total s-energy of a multiagent system. *SIAM J. Control. Optim.* **49**(4), 1680–1706 (2011)
21. Cinelli, M., De Francisci Morales, G., Galeazzi, A., Quattrociocchi, W., Starnini, M.: The echo chamber effect on social media. *Proc. Natl. Acad. Sci.* **118**(9), e2023301118 (2021)
22. Clementi, A.E.F., Ghaffari, M., Gualà, L., Natale, E., Pasquale, F., Scornavacca, G.: A tight analysis of the parallel undecided-state dynamics with two colors. In: Proceedings of the 43rd Symposium on Mathematical Foundations of Computer Science (MFCS), pp. 28:1–28:15 (2018)
23. Coates, A., Han, L., Kleerekoper, A.: A unified framework for opinion dynamics. In: Proceedings of the 17th Conference on Autonomous Agents and Multi-agent Systems (AAMAS), pp. 1079–1086 (2018a)
24. Coates, A., Han, L., Kleerekoper, A.: A Unified opinion framework simulator. In: Proceedings of the 17th Conference on Autonomous Agents and Multi-agent Systems (AAMAS), pp. 1803–1805 (2018b)
25. Cooper, C., Elsässer, R., Ono, H., Radzik, T.: Coalescing random walks and voting on graphs. In: Proceedings of the Symposium on Principles of Distributed Computing (PODC), pp. 47–56 (2012)
26. Cooper, C., Elsässer, R., Radzik, T.: The power of two choices in distributed voting. In: Proceedings of the 41st International Colloquium on Automata, languages and programming (ICALP), pp. 435–446 (2014)
27. Cooper, C., Elsässer, R., Radzik, T., Rivera, N., Shiraga, T.: Fast consensus for voting on general expander graphs. In: Proceedings of the 29th Symposium on Distributed Computing (DISC), pp. 248–262 (2015)
28. Cooper, C., Radzik, T., Rivera, N., Shiraga, T.: Fast plurality consensus in regular expanders. In: Proceedings of the 31st Symposium on Distributed Computing (DISC), pp. 13:1–13:16 (2017)
29. De, A., Bhattacharya, S., Ganguly, N.: Shaping opinion dynamics in social networks. In: Proceedings of the 17th Conference on Autonomous Agents and Multi-agent Systems (AAMAS), pp. 1336–1344 (2018)
30. DeGroot, M.H.: Reaching a consensus. *J. Am. Stat. Assoc.* **69**(345), 118–121 (1974)
31. Downs, A.: *An Economic Theory of Democracy*. Harper & Row, New York (1957)
32. Epiropou, M., Fotakis, D., Hofer, M., Skoulakis, S.: Opinion formation games with aggregation and negative influence. *Theory Comput. Syst.* **63**(7), 1531–1553 (2019)
33. Etesami, S.R., Başar, T.: Game-theoretic analysis of the Hegselmann–Krause model for opinion dynamics in finite dimensions. *IEEE Trans. Autom. Control.* **60**(7), 1886–1897 (2015)
34. Evans, T., Fu, F.: Opinion formation on dynamic networks: identifying conditions for the emergence of partisan echo chambers. *R. Soc. Open Sci.* **5**(10), 181122 (2018)
35. Faliszewski, P., Gonen, R., Koutecký, M., Talmon, N.: Opinion diffusion and campaigning on society graphs. In: Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI), pp. 219–225 (2018)
36. Fortunato, S.: On the consensus threshold for the opinion dynamics of Krause–Hegselmann. *Int. J. Modern Phys. C* **16**(02), 259–270 (2005)
37. Fotakis, D., Palyvos-Giannas, D., Skoulakis, S.: Opinion dynamics with local interactions. In: Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI), pp. 279–285 (2016)
38. Friedkin, N.E., Johnsen, E.C.: Social influence and opinions. *J. Math. Soc.* **15**(3–4), 193–206 (1990)
39. Fu, F., Wang, L.: Coevolutionary dynamics of opinions and networks: from diversity to uniformity. *Phys. Rev. E* **78**(1), 016104 (2008)
40. Ghaffari, M., Lengler, J.: Nearly-tight analysis for 2-choice and 3-majority consensus dynamics. In: Proceedings of the Symposium on Principles of Distributed Computing (PODC), pp. 305–313 (2018)
41. Ghaffari, M., Parter, M.: A polylogarithmic gossip algorithm for plurality consensus. In: Proceedings of the Symposium on Principles of Distributed Computing (PODC), pp. 117–126 (2016)
42. Hassin, Y., Peleg, D.: Distributed probabilistic polling and applications to proportionate agreement. *Inf. Comput.* **171**(2), 248–268 (2001)
43. Hegselmann, R., Krause, U.: Opinion dynamics and bounded confidence models, analysis, and simulation. *J. Artif. Soc. Soc. Simul.* **5**(3) (2002)
44. Lengler, J.: Drift analysis. In: Doerr, B., Neumann, F. (eds.), *Theory of Evolutionary Computation: Recent Developments in Discrete Optimization*. Chapter 2, pp. 89–131 (2020)
45. Martinsson, A.: An improved energy argument for the Hegselmann–Krause model. *J. Differ. Eq. Appl.* **22**(4), 513–518 (2016)
46. McPherson, M., Smith-Lovin, L., Cook, J.M.: Birds of a feather: homophily in social networks. *Ann. Rev. Sociol.* **27**(1), 415–444 (2001)
47. Nakata, T., Imahayashi, H., Yamashita, M.: A probabilistic local majority polling game on weighted directed graphs with an application to the distributed agreement problem. *Networks* **35**(4), 266–273 (2000)

48. Parasis, R., Franceschetti, M., Touri, B.: On the convergence properties of social Hegselmann–Krause dynamics. [arXiv:1909.03485](https://arxiv.org/abs/1909.03485) [math.OC] (2019)
49. Pariser, E.: *The Filter Bubble: What the Internet is Hiding From You*. Penguin, London (2011)
50. Touri, , Nedic, A: Discrete-time opinion dynamics. In: M.B. Matthews (Ed.), *Conference Record of the 45th Asilomar Conference on Signals, Systems and Computers (ACSCC)*, pp. 1172–1176 (2011)
51. Wedin, E., Hegarty, P.: A quadratic lower bound for the convergence rate in the one-dimensional Hegselmann–Krause bounded confidence dynamics. *Discret. Comput. Geom.* **53**(2), 478–486 (2015)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.