

Opinion Dynamics with Local Interactions

Dimitris Fotakis and Dimitris Palyvos-Giannas and Stratis Skoulakis

National Technical University of Athens

fotakis@cs.ntua.gr, dpalyvos@corelab.ntua.gr, sskoul@corelab.ntua.gr

Abstract

We study convergence properties of opinion dynamics with local interactions and limited information exchange. We adopt a general model where the agents update their opinions in rounds to a weighted average of the opinions in their neighborhoods. For fixed neighborhoods, we present a simple randomized protocol that converges in expectation to the stable state of the Friedkin-Johnsen model. For opinion-dependent neighborhoods, we show that the Hegselmann-Krause model converges to a stable state if each agent’s neighborhood is restricted either to a subset of her acquaintances or to a small random subset of agents. Our experimental findings indicate that for a wide range of parameters, the convergence time and the number of opinion clusters of the neighborhood-restricted variants are comparable to those of the standard Hegselmann-Krause model.

“Most people are other people. Their thoughts are someone else’s opinions, their lives a mimicry, their passions a quotation.” – Oscar Wilde

1 Introduction

In everyday life, people form their opinions by exchanging information about almost everything. Information exchange is mostly *local*, in the sense that socially connected people (e.g., family, friends, colleagues) interact more often and affect each other’s opinion more strongly, and *dynamic*, i.e., discussions go on until opinions converge. Due to their ubiquity and importance, opinion dynamics have been studied extensively for decades (see e.g., [Jackson, 2008]). The interest has been intensified recently by the rapid growth of online social networks (e.g., Facebook, Twitter), which provide the basis for a better understanding (and possibly for applications, see e.g., [Lever, 2010]) of opinion formation.

Over the years, several models of opinion formation have been proposed. The standard assumption is that people (or agents, in general) update their opinions to a weighted average of the opinions in their neighborhood. The updates proceed in synchronous parallel rounds, until opinions reach a stable state. A distinctive feature of the models is whether the agent neighborhoods are determined by a fixed social network, decoupled from the opinion formation process, or they

reflect the affinity of the agent opinions and are updated during the process. Among the most influential models are those of [DeGroot, 1974] and [Friedkin and Johnsen, 1990], in the former case, and the *bounded confidence* model of [Hegselmann and Krause, 2002], in the latter. In this work, we investigate how practical restrictions on the locality of agent interaction and on the amount of information exchanged in each round affect the convergence properties of these models.

Related Work. The convergence properties of the Friedkin-Johnsen (FJ) and the Hegselmann-Krause (HK) models have received considerable attention. Both are known to converge to a stable state relatively fast. In the FJ model, the stable state is unique [Friedkin and Johnsen, 1990] and approximately minimizes a certain notion of opinion disagreement in the network [Bindel *et al.*, 2011]. Moreover, the convergence time is determined by the spectral radius of the adjacency matrix [Ghaderi and Srikant, 2014] (see also [Ferraioli *et al.*, 2012; Yildiz *et al.*, 2013] for discrete variants of the FJ model).

The HK model leads to *opinion clustering*, where the agents form groups with consensus in each group (see e.g., [Blondel *et al.*, 2009]). The single-dimensional HK model is known to converge in at least $\Omega(n^2)$ and at most $O(n^3)$ rounds, where n is the number of agents [Bhattacharyya *et al.*, 2013; Wedin and Hegarty, 2015]. [Lorenz, 2005] proved that generalizations of the HK model converge to a stable state under symmetric (and possibly time-dependent) local agent interactions. E.g., this implies convergence of the Deffuant-Weisbuch (DW) model [Deffuant *et al.*, 2001; Weisbuch *et al.*, 2002], where two random agents meet and exchange opinions in each round. Moreover, there has been significant experimental work on the convergence properties of variants of the HK model and on confidence levels that are sufficient or necessary for consensus (see e.g., [Fortunato, 2005; Lorenz and Urbig, 2007]). Nevertheless, the convergence properties of heterogeneous and/or asymmetric variants of the HK model are far from being well understood. E.g., only recently, [Chazelle and Wang, 2015] proved that the HK model with partially stubborn agents converges.

Recent work has also studied the convergence properties of opinion dynamics that combine features of the DeGroot/FJ and the HK models and allow for *coevolution* of the opinions and the network under mutual influence between the agents. E.g., motivated by applications of opinion dynamics to expert opinion aggregation, [Carvalho and Larson, 2013] proved

convergence to consensus for a generalization of DeGroot’s model, where the influence between two agents is inversely proportional to the distance of their opinions. Subsequently, [Tsang and Larson, 2014] analyzed experimentally the convergence properties of a similar model where the agents are skeptical towards opinions far away from theirs. On the theoretical side, [Bhawalkar *et al.*, 2013] investigated the existence and the efficiency of stable states for discrete and continuous coevolutionary opinion formation processes.

Motivation and Contribution. An important assumption in most of the previous work is that each agent exchanges opinions with all its neighbors in every round. In the HK model, in particular, determining each agent’s neighborhood requires access to the opinions of all agents! This assumption is crucial for establishing existence of and convergence to a stable state in the FJ model [Ghaderi and Srikant, 2014] and fast convergence to opinion clustering in the HK model [Bhattacharyya *et al.*, 2013]. However, this assumption is questionable in real life and in modern online social networks. Instead, it is far more reasonable to assume that agents act within a social network and update their opinions by consulting the opinions of a small (possibly random) subset of their neighbors. Motivated by this natural observation and by experimental work on opinion dynamics with communication regimes [Urbig and Lorenz, 2004], we introduce variants of the FJ and the HK model where information exchange in each round is limited (for both models) and local (for the HK model). We thoroughly investigate, both theoretically and experimentally, to which extent the convergence properties of the FJ and the HK models are affected by such practical considerations in the opinion exchange between the agents.

For the FJ model, we present a simple randomized protocol where each agent consults the opinion of one of her social neighbors in each round. We show that it converges in expectation to the stable state of the FJ model, albeit in exponential time. Our result indicates the importance of lively communication for the fast convergence of opinion dynamics.

Our main contribution concerns variants of the HK model with local agent interactions and limited (and possibly asymmetric) communication in each round. To decouple the two issues, we consider two variants of the HK model, one accompanied by a fixed social network and another with agent neighborhoods determined by random sampling in each round. In the former variant, the agents are embedded in an undirected network and exchange opinions only with their social neighbors in each round. Due to this restriction, the relative order of agent opinions is not preserved and standard convergence proofs for the HK model cannot be applied. Inspired by the notion of temporal network connectivity [Kempe *et al.*, 2002] and using a recent result of [Chazelle, 2011], we prove convergence of this variant to a stable state, where opinion clustering occurs within social neighborhoods. Our approach provides a simpler and more versatile proof of the main results in [Hendrickx and Blondel, 2006; Lorenz, 2005]. We also extend this convergence result to multidimensional opinions. In the variant with random sampling, each agent consults the opinions of a small random subset of other agents, chosen independently in each round. Extending our temporal connectivity approach, we show that this vari-

ant converges to a stable state. Conceptually, HK with random sampling can be regarded as a significant generalization of the DW model. Technically, the DW model employs symmetric opinion exchange between random agent pairs. Hence, its convergence follows directly from the result above. On the other hand, HK with random sampling employs asymmetric opinion exchange (i.e., it may be that an agent i consults the opinion of an agent j , but j does not consult i ’s opinion). Thus, convergence proofs for symmetric variants of the HK model do not apply to HK with random sampling (e.g., condition (2) in [Lorenz, 2005, Theorem 2] is violated). Moreover, such asymmetric variants of the HK model either do not converge to a stable state, or if they converge, it is notoriously difficult to be proven (see e.g., [Bhawalkar *et al.*, 2013]).

Our experimental findings indicate that if the size of random neighborhoods is roughly $\log n/\varepsilon$, where n is the number of agents and ε is the agent confidence, the HK model with random sampling converges at comparable speed and to almost the same number of opinion clusters as the standard HK model. On the other hand, the HK model on random $\mathcal{G}_{n,p}$ and power-law networks and on real-world social networks converges noticeably slower and to a larger number of clusters than the standard version. This behavior is amplified by small confidence values. The delay in convergence is justified by the existence of opinion clusters with relatively large diameter in the social network, while the cluster proliferation is justified by the existence of many small opinion clusters in relatively isolated parts of the network.

Due to the space limitations, some proofs and a significant part of our experimental justification are omitted.

2 Models of Opinion Dynamics

We consider a population of n agents, each maintaining an opinion $x_i(t) \in [0, 1]$ in each time step $t \in \mathbb{N}$. We let $\vec{x}(t) \in [0, 1]^n$ denote the vector of agents’ opinions at time t .

We use $\|\vec{x}\|_1$ and $\|\vec{x}\|_\infty$ to denote the 1-norm and the infinity-norm of a vector \vec{x} , respectively. We say that a sequence $\{\vec{x}(t)\}_{t \in \mathbb{N}}$ converges to a stable state \vec{x}^* if for all $\gamma > 0$, there is a t_γ , so that for all $t \geq t_\gamma$, $\|\vec{x}(t) - \vec{x}^*\|_\infty \leq \gamma$.

The Friedkin-Johnsen Model. We are given a weighted network $G(V, E)$ and the agents’ initial opinions $\vec{x}(0)$. Each agent i corresponds to vertex $i \in V$, has weight $w_{ii} \in (0, 1]$ for her initial opinion and weight $w_{ij} \in [0, 1)$ for the current opinion of each other agent j . We assume that $w_{ij} = w_{ji}$, for all i, j , and that $\sum_{j \in V} w_{ij} = 1$. At any round $t \geq 1$, each agent i updates her opinion to

$$x_i(t) = \sum_{j \neq i} w_{ij} x_j(t-1) + w_{ii} x_i(0) \quad (1)$$

We let A denote the adjacency matrix of G with 0 on its main diagonal and let B be the diagonal matrix with $B_{ii} = w_{ii}$, for all i . Then, (1) can be written in matrix form as:

$$\vec{x}(t) = A\vec{x}(t-1) + B\vec{x}(0) \quad (2)$$

Since $w_{ii} > 0$, A is a substochastic matrix and has spectral radius $\rho(A) < 1$. Thus, (2) converges to a unique stable state $\vec{x}^* = (I - A)^{-1} B\vec{x}(0)$, where I is the $n \times n$ identity matrix, in $O\left(\frac{\ln(n/\gamma)}{1-\rho(A)}\right)$ time steps [Ghaderi and Srikant, 2014].

The Hegselmann-Krause Model. We are given the agents' confidence $\varepsilon > 0$ and their initial opinions $\vec{x}(0)$. At any round $t \geq 1$, agent i computes her opinion-dependent neighborhood

$$S_i(t, \varepsilon) = \{j : |x_i(t-1) - x_j(t-1)| \leq \varepsilon\}$$

and updates her opinion to the average of the opinions in $S_i(t, \varepsilon)$, i.e., $x_i(t) = \sum_{j \in S_i(t, \varepsilon)} x_j(t-1) / |S_i(t, \varepsilon)|$.

We consider two variants of the HK model with local agent interaction and limited information exchange. In the *network-HK* model, we are also given a undirected connected social network $G(V, E)$. The neighborhood of each agent i is now restricted to her social neighbors:

$$S_i(G, t, \varepsilon) = \{j : \{i, j\} \in E \text{ and } |x_i(t-1) - x_j(t-1)| \leq \varepsilon\}$$

In the *random-HK* model, we let k be the size of the sample. At any round $t \geq 1$, each agent i samples (from all agents, with replacement) a random subset K_i of k agents and lets

$$S_i(K_i, t, \varepsilon) = \{j \in K_i \text{ and } |x_i(t-1) - x_j(t-1)| \leq \varepsilon\}$$

We require that $\{i, i\} \in E$ and $i \in K_i$, for all i , so that i is always included in S_i , in both variants. Then, agent i updates her opinion to the average of the opinions in $S_i(G, t, \varepsilon)$ and $S_i(K_i, t, \varepsilon)$, respectively.

To establish that the two variants converge, we use the following equation describing the agent opinions at round $t \geq 1$:

$$\vec{x}(t) = A_t A_{t-1} \cdots A_1 \vec{x}(0) \quad (3)$$

Each A_r is an $n \times n$ matrix with $1/|S_i|$ in position $A_r(i, j)$, if $j \in S_i$, and 0, otherwise (S_i is $S_i(G, t, \varepsilon)$ for network-HK and $S_i(K_i, t, \varepsilon)$ for random-HK). A_r corresponds to the adjacency matrix of an unweighted graph (undirected for network-HK and directed for random-HK), but with normalized entries so that A_r is stochastic.

In the convergence proofs, we use the *coefficient of ergodicity* $\tau(A) = \frac{1}{2} \max_{i,j} \|A^T(\vec{e}_i - \vec{e}_j)\|_1$ of a stochastic square matrix A , where A^T is A 's transpose and \vec{e}_i is the vector with 1 in coordinate i and 0 elsewhere (see e.g., [Seneta, 1979]). We know that for any stochastic matrices A, B , (i) $\tau(A) \leq 1$, (ii) $\tau(AB) \leq \tau(A)\tau(B)$, and (iii) $\tau(A) = 0$ iff $\text{rank}(A) = 1$. Moreover, if A has strictly positive elements only, $\tau(A) < 1$.

We also let $\Pr[E]$ denote the probability of an event E and $\mathbb{E}[X]$ denote the expectation of a random variable X .

3 The FJ Model with Limited Information

A natural question about the FJ model is whether we can simulate the opinion formation process by simple protocols where agents consult the opinions of a small subset of their neighbors in each round. To this end, we present such the *Limited Information Protocol*, or *LIP-FJ*, in brief, and discuss its convergence properties.

Let $(G(V, E), \vec{x}(0))$ be an instance of the FJ model and let $\lambda(t) \in (0, 1)$. In LIP-FJ, at any round $t \geq 1$, each agent i selects one index $j \in V$ with probability w_{ij} and sets $s_i(t) = x_j(t-1)$, if $j \neq i$, and $s_i(t) = x_i(0)$, if $j = i$. Then, agent i updates her opinion to $x_i(t) = \lambda(t)x_i(t-1) + (1-\lambda(t))s_i(t)$.

We observe that (i) for any fixed instance and any fixed round t , the set Ω^t of possible values of the random variable $\vec{x}(t)$ is finite; and that (ii) for any possible value \vec{y} of $\vec{x}(t-1)$,

$$\mathbb{E}[\vec{x}(t) | \vec{x}(t-1) = \vec{y}] = \lambda(t)\vec{y} + (1-\lambda(t))(A\vec{y} + B\vec{x}(0))$$

For brevity, let $p_{t-1}(\vec{y}) \equiv \Pr[\vec{x}(t-1) = \vec{y}]$ be the probability that the opinions at round $t-1$ are as in \vec{y} . Then, using (i) and (ii), we obtain that for any fixed $t \geq 1$,

$$\begin{aligned} \mathbb{E}[\vec{x}(t)] &= \sum_{\vec{y} \in \Omega^{t-1}} p_{t-1}(\vec{y}) \mathbb{E}[\vec{x}(t) | \vec{x}(t-1) = \vec{y}] \\ &= \lambda(t) \sum_{\vec{y} \in \Omega^{t-1}} p_{t-1}(\vec{y}) \vec{y} + \\ &\quad + (1-\lambda(t)) \sum_{\vec{y} \in \Omega^{t-1}} p_{t-1}(\vec{y}) (A\vec{y} + B\vec{x}(0)) \end{aligned}$$

Using linearity of expectation, we conclude that any $t \geq 1$,

$$\mathbb{E}[\vec{x}(t)] = \lambda(t)\mathbb{E}[\vec{x}(t-1)] + (1-\lambda(t))(A\mathbb{E}[\vec{x}(t-1)] + B\vec{x}(0))$$

Then, using induction on t and standard properties of the matrix infinity norm $\nu(A) \equiv \|A\|_\infty \in (0, 1)$, we show that for an appropriate choice of $\lambda(t)$, LIP-FJ converges in expectation to the stable state $\vec{x}^* = (I-A)^{-1}B\vec{x}(0)$ of the FJ model.

Theorem 1. *For any instance $(G(V, E), \vec{x}(0))$ and any round $t \geq 1$, the opinions maintained by LIP-FJ satisfy*

$$\|\mathbb{E}[\vec{x}(t)] - \vec{x}^*\|_\infty \leq e^{-(1-\nu(A))\sum_{q=1}^t(1-\lambda(q))} \|\vec{x}(0) - \vec{x}^*\|_\infty$$

where \vec{x}^* is the stable state of the FJ model.

If $\lambda(t) = \frac{t}{t+1}$, then $\|\mathbb{E}[\vec{x}(t)] - \vec{x}^*\|_\infty \leq e^{1-\nu(A)}/t^{1-\nu(A)}$, by Theorem 1, and LIP-FJ converges in expectation in time exponential in $1/(1-\nu(A))$. The experiments indicate that for such values of $\lambda(t)$, LIP-FJ indeed converges asymptotically to \vec{x}^* at a very slow rate. For larger values of $\lambda(t)$, e.g., for $\lambda(t) = 1 - \frac{1}{t^2}$, $\sum_{q=1}^t(1-\lambda(q))$ converges to a constant value. Therefore, the expected distance to \vec{x}^* stops decreasing after a finite number of rounds. If we set $\lambda(t)$ to some constant, aiming at improving the convergence time, the experiments indicate that LIP-FJ does not converge asymptotically to \vec{x}^* , due to the high variance of the stochastic process.

4 Convergence of Network-HK

We proceed to show that the HK model on social networks converges to a stable state. We consider any fixed instance $(G(V, E), \varepsilon, \vec{x}(0))$ of the network-HK model. As in Section 2, we let A_t be the stochastic matrix encoding the agent neighborhoods in round t and let G_t be the corresponding undirected network. Then, $\vec{x}(t) = A_t \vec{x}(t-1)$ for each $t \geq 1$.

Our approach is motivated by temporal graph connectivity. Specifically, we say that a set of agents $S \subseteq V$ is *weakly connected* if for any non-empty $S' \subset S$ and any $t_0 \in \mathbb{N}$, there is a round $t \geq t_0$ so that G_t includes at least one edge connecting an agent in S' to some agent in $S \setminus S'$ (we highlight that the sequence $\{G_t\}_{t \geq 0}$ and weak connectivity depend on the opinion formation process). Intuitively, all agents in a weakly connected set influence each other with their opinions.

Lemma 1. *Let $(G(V, E), \varepsilon, \vec{x}(0))$ be an instance of network-HK, where the opinion dynamics keep V weakly connected. Then, all agents converge to a single opinion x^* .*

Proof. To prove the lemma, it suffices to show that there is a round t_0 so that the coefficient of ergodicity of the matrix

$C = A_{t_0}A_{t_0-1} \cdots A_1$ is $\tau(C) \leq \varepsilon/2$. Given this, we have that $\bar{x}(t_0) = C\bar{x}(0)$, by (3), and that for all agents i and j ,

$$\begin{aligned} |x_i(t_0) - x_j(t_0)| &= |(C_i - C_j)\bar{x}(0)| \\ &\leq \|C_i - C_j\|_1 \leq 2\tau(C) \leq \varepsilon \end{aligned} \quad (4)$$

where C_i is the i -th row of matrix C . Since at t_0 all opinions are within distance ε , in any round $t > t_0$, all agents take the average of all opinions in their social neighborhood (including their opinion). Hence, after round t_0 , we have essentially an instance of DeGroot's model on the undirected connected network G (enhanced with self-loops). Since G (with self-loops) defines an irreducible and aperiodic process, by [Jackson, 2008, Cor. 8.3.1], all agents converge to a single opinion.

So, we need to show that for any $\delta > 0$, there is a t_0 such that $\tau(A_{t_0}A_{t_0-1} \cdots A_1) \leq \delta$. To this end, we first show that since V remains weakly connected, for any round t , there is a round $\ell(t)$, such that the matrix $C_t^{\ell(t)} = A_{\ell(t)}A_{\ell(t)-1} \cdots A_t$ has all its elements positive, and thus, $\tau(C_t^{\ell(t)}) < 1$.

We recall that any matrix A_t has only positive diagonal elements. Moreover, an element $C_t^{\ell(t)}(i, j)$ is positive iff there is a (time-respecting) walk $(i, i_1, \dots, i_{m-1}, j)$ from agent i to agent j such that (i) the first edge $\{i, i_1\}$ exists in some network $G_{t'}$ with $t' \geq t$, (ii) for each index q , $1 \leq q \leq m-1$, if the edge $\{i_{q-1}, i_q\}$ exists in some $G_{t'}$, then the edge $\{i_q, i_{q+1}\}$ exists in some network $G_{t''}$ with $t'' > t'$, and (iii) the last edge $\{i_{m-1}, j\}$ exists in some $G_{t'}$ with $t' \leq \ell(t)$ (since any matrix A_t has positive diagonal elements, the walk can wait at each intermediate agent until the next edge appears). The existence of such a walk between any pair of agents in V follows from the definition of weak connectivity.

Then, [Chazelle, 2011, Theorem 2.5] implies that there is a fixed $\eta > 0$, so that for any round t and the corresponding round $\ell(t)$, $\tau(A_{\ell(t)} \cdots A_t) \leq 1 - \eta$. Now, we can concatenate an appropriately large number of non-overlapping sequences $A_t, \dots, A_{\ell(t)}$ and obtain a sequence A_1, \dots, A_{t_0} with $\tau(A_{t_0} \cdots A_1) \leq \delta$, for any $\delta > 0$. \square

If V is not weakly connected, we can show that it admits a unique partition into weakly connected components. Applying Lemma 1, we get that the agents in each component converge to a single opinion. Thus, we obtain the following:

Theorem 2. *Network-HK converges to a stable state.*

Extension to d -Dimensional Opinions. We can generalize Theorem 2 to the case where each agent i maintains a d -dimensional opinion $\vec{x}_i(t) \in [0, 1]^d$ by the d -dimensional HK model (see e.g., [Bhattacharyya *et al.*, 2013]) on a social network G . The proof is essentially identical, with the only difference that in (4), we need that $\tau(C) \leq \varepsilon/(2\sqrt{d})$.

5 Convergence of Random-HK

Next, we show that the random-HK model, where each agent contacts only k random agents in each round, converges to a stable state. Let $(k, \varepsilon, \vec{x}(0))$ be any fixed instance. As before, we let A_t be the stochastic matrix that encodes the agent neighborhoods in round t and let G_t be the corresponding network. Now A_t is a random matrix and G_t is a random directed network. Again, $\vec{x}(t) = A_t\vec{x}(t-1)$, for any $t \geq 1$.

Since the process is now stochastic, we prove that each agent converges to a specific opinion with probability arbitrarily close to 1.

We first prove the equivalent of Lemma 1, but now with the notion of ε -connectivity. For any two disjoint sets of agents S_1, S_2 , we let $d^t(S_1, S_2) = \min_{i \in S_1, j \in S_2} |x_i(t) - x_j(t)|$ be their distance at round t . We say that a set of agents S is ε -connected at round t , if for any non-empty $S' \subset S$, $d^t(S', S \setminus S') \leq \varepsilon$. A set of agents S breaks at round t if S is ε -connected at round $t-1$ and is not ε -connected at round t . We note that once S' and $S \setminus S'$ break, they behave as independent subinstances in the future. We also say that agents i and j are (ε, t) -connected if there is a "path" $(i, i_1, \dots, i_{k-1}, j)$ so that for each "step" q , $|x_{i_q}(t) - x_{i_{q+1}}(t)| \leq \varepsilon$. The diameter at some time step t , denoted $\text{diam}(t)$, is the maximum distance $|x_i(t) - x_j(t)|$ overall (ε, t) -connected pairs of agents i, j . We observe that if at time step t , $\text{diam}(t) \leq \varepsilon$, then no break occurs in the future. We also let Γ_l be the set of all opinion vectors \vec{y} such that for all rounds $t \geq 0$, $\Pr[\text{at most } l \text{ breaks occur in } \{0, t\} \mid \vec{x}(0) = \vec{y}] = 1$. Namely, Γ_l consists of all vectors \vec{y} such that if the initial opinions are \vec{y} , then no matter the random choices, at most l breaks occur. E.g., all vectors with a single opinion belong to Γ_0 . We next show that in a set that remains ε -connected during the (random) opinion formation process, all agents converge to a single opinion with probability that tends to 1.

Lemma 2. *Let $(k, \varepsilon, \vec{x}(0))$ be an instance of the random-HK model with $\vec{x}(0) \in \Gamma_0$. For any $\gamma, \delta > 0$, there is a round t_0 such that*

$$\Pr[\text{diam}(t) \leq \gamma] \geq 1 - \delta$$

Proof sketch. Without loss of generality, we assume that there exists a single ε -connected component (otherwise, the lemma applies to each ε -connected component separately). Since $\vec{x}(0) \in \Gamma_0$ no break occurs and the agents are (ε, t) -connected for all t and all random choices.

Let $p = 1 - (1 - 1/n)^k$ be the probability that an agent j is not in the sample set of agent i in a round t . For any round ℓ , we denote $C_\ell = A_{\ell+2n^2/p} \cdots A_\ell$ and $D_\ell = A_{\ell-1} \cdots A_1$. The important step is to show that there is some fixed $\eta > 0$ such that for any fixed value of D_ℓ , $\mathbb{E}[\tau(C_\ell) \mid D_\ell] \leq 1 - \eta/2$.

For any round $t' \geq \ell$, $\text{pos}(t')$ (resp. $\text{pos}_i(t')$) denotes the number of positive entries in (resp. the i -th row of) matrix $A_{t'} \cdots A_\ell$. We have $0 \leq \text{pos}(t') \leq n^2$ and $\text{pos}(t'+1) \geq \text{pos}(t')$. As long as $\text{pos}(t') < n^2$, there is some agent i with $\text{pos}_i(t') < n$. As in the proof of Lemma 1, $\text{pos}_i(t')$ is the number of agents reachable from i , between rounds ℓ and t' , by time-respecting walks. Since V is ε -connected and no break occurs, if $\text{pos}_i(t') < n$, there is at least one new agent reachable from i in round $t'+1$, with probability at least p . Hence, for any round t' with $\text{pos}(t') < n^2$, $\text{pos}(t'+1) > \text{pos}(t')$ with probability at least p , and the expected number of rounds before it becomes $\text{pos}(t') = n^2$ for the first time is at most n^2/p . By Markov's inequality, $\Pr[\text{pos}(\ell + 2n^2/p) < n^2 \mid D_\ell] \leq 1/2$. Moreover, since $\text{pos}(\ell + 2n^2/p) = n^2$ implies that $\tau(C_\ell) < 1$,

$$\Pr[\tau(C_\ell) < 1 \mid D_\ell] \geq 1/2$$

As in the proof of Lemma 1, since C_ℓ is the product of $2n^2/p$ matrices, there exists a fixed fractional $\eta > 0$ such

that if $\tau(C_\ell) < 1$, then $\tau(C_\ell) \leq 1 - \eta$. Thus, we obtain that for any fixed value of D_ℓ , $\mathbb{E}[\tau(C_\ell)|D_\ell] \leq 1 - \eta/2$.

Now, we can work as in the proof of Lemma 1. Taking an appropriately large number of rounds, we obtain a t_0 and a matrix $C = A_{t_0} \cdots A_1$ such that $\tau(C) \leq \gamma/2$ with probability at least $1 - \delta$. Then, the Lemma follows from the properties of the coefficient of ergodicity, similarly to (4). \square

We proceed to show that the random-HK model converges asymptotically, with probability that tends to 1. We recall that if there exists a round t^* with $\text{diam}(t^*) \leq \varepsilon$, then $\vec{x}(t^*) \in \Gamma_0$ and Lemma 2 again implies convergence in each ε -connected component separately. The following lemma establishes the existence of such a round t^* with probability that tends to 1.

Lemma 3. *Let $(k, \varepsilon, \vec{x}(0))$ be any instance of the random-HK model. For any $\delta > 0$ there is a round t^* such that*

$$\Pr[\text{diam}(t^*) \leq \varepsilon] \geq 1 - \delta$$

Proof sketch. Intuitively, if $\vec{x}(0) \in \Gamma_l$, there are constants p and t_0 such that $\Pr[\vec{x}(t_0) \in \Gamma_{l-1}] \geq p$. Moreover, there is a constant m such that $\Pr[\vec{x}(mt_0) \in \Gamma_0] \approx 1$, i.e., with almost certainty, all possible breaks have occurred by round mt_0 . Then, the proof follows easily from Lemma 2.

In the following, we let t_0 be the number of rounds in Lemma 2 for $\gamma = \varepsilon$. Namely, t_0 is such that if $\vec{x}(0) \in \Gamma_0$ then $\Pr[\text{diam}(t_0) \leq \varepsilon] \geq 1 - \delta$. For brevity, we let $p = 1/n^{knt_0}$ and let $P(\vec{y}, m) = \Pr[\text{diam}((m+1)t_0) \leq \varepsilon | \vec{x}(0) = \vec{y}]$. Namely, $P(\vec{y}, m)$ is the probability that the diameter is at most ε after $(m+1)t_0$ rounds, given that the initial opinion vector is \vec{y} .

At first, we consider the case where $\vec{x}(0) \in \Gamma_1$ and prove inductively that

$$P(\vec{x}(0), m) \geq (1 - \delta)(1 - (1 - p)^m) \quad (5)$$

We can verify (5) is true for $m = 1$. We inductively assume that m satisfies (5) and consider the following cases for $m+1$.

$\Pr[\vec{x}(t_0) \in \Gamma_0] = 0$. Therefore, since $\vec{x}(0) \in \Gamma_1$, no break occurs in $\{0, t_0\}$ for all random choices. Thus, $\vec{x}(0)$ satisfies the hypothesis of Lemma 2 and $P(\vec{x}(0), 0) \geq 1 - \delta$. As a result, $P(\vec{x}(0), m+1) \geq 1 - \delta$.

$\Pr[\vec{x}(t_0) \in \Gamma_0] > 0$. There is an opinion vector $\vec{y} \in \Gamma_0$ such that $\Pr[\vec{x}(t_0) = \vec{y}] \geq p$. Since $\vec{x}(0) \in \Gamma_1$, if $\vec{x}(t_0) \neq \vec{y}$, then $\vec{x}(t_0) \in \Gamma_1$. Hence, we obtain that:

$$\begin{aligned} P(\vec{x}(0), m+1) &= \\ &\Pr[\vec{x}(t_0) = \vec{y}] P(\vec{y}, m) + \sum_{\vec{a} \in \Gamma_1} \Pr[\vec{x}(t_0) = \vec{a}] P(\vec{a}, m) \\ &\geq (1 - \delta)[p + (1 - p)(1 - (1 - p)^m)] \\ &\geq (1 - \delta)(1 - (1 - p)^{m+1}) \end{aligned}$$

Now we extend the proof to the case where $\vec{x}(0) \in \Gamma_l$, for any $2 \leq l \leq n - 1$. We recursively define the functions $f_l(m) = pf_{l-1}(m-1) + (1-p)f_l(m-1)$, for all $m, l \geq 2$, with $f_1(m) = (1 - \delta)(1 - (1 - p)^m)$. Using induction and the same arguments as above, we can show that if $\vec{x}(0) \in \Gamma_l$, then $P(\vec{x}(0), m) \geq (1 - \delta)f_l(m)$. We observe that $\lim_{m \rightarrow \infty} f_1(m) = 1$. Then, by the definition of f_l , we can show inductively that $\lim_{m \rightarrow \infty} f_{n-1}(m) = 1$. Since at most $n - 1$ breaks can occur, we conclude that $P(\vec{x}(0), m) \geq (1 - \delta)f_{n-1}(m)$. \square

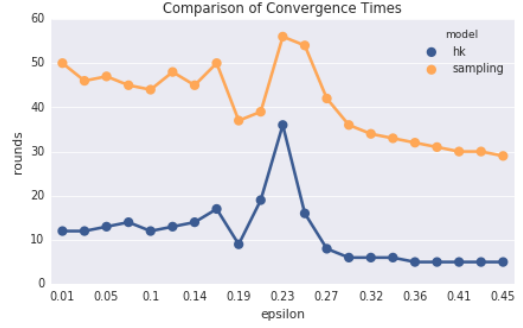


Figure 1: Comparison of average convergence time of the standard HK and the random-HK models.

6 Experimental Evaluation

We performed computer simulations so as to obtain a better understanding of the convergence properties and the number of clusters of network-HK and random-HK. We sketch the main conclusions of the experimental evaluation, which confirm and enhance our results in sections 4 and 5.

The Random-HK Model. We simulate the random-HK and the standard HK models for 625 agents and 21 values of $\varepsilon \in [0.01, .45]$, repeating each run for 100 different initial opinion vectors to account for the randomness in local interaction. Opinions are selected uniformly at random from $[0, 1]$ in all our simulations. We choose $k = \min(n/10, \log n/\varepsilon)$ to ensure that with high probability every agent has some neighbors in each cluster.

Figure 1 shows the average convergence time for different values of ε . The curve for random-HK is very similar to that for the standard HK model. Naturally, random-HK results in larger convergence times, as every agent has access to less information. The peak at .25 is where we move from two clusters to one cluster in the stable state.

The number of opinion clusters is an important property of the HK model. We calculate it by sorting the final opinion values and separating the groups which differ more than ε from each other. Figure 2 shows the average number of clusters created by each model as ε changes. At a first glance, the limited amount of information in random-HK does not hinder its ability to simulate the standard HK model effec-

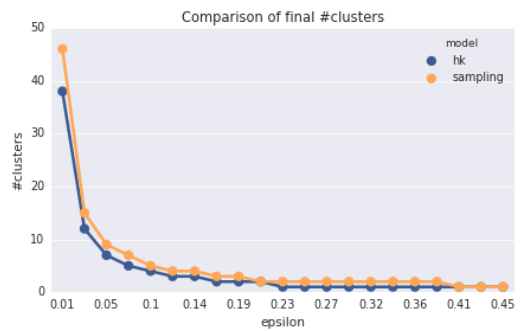


Figure 2: Comparison of the number of clusters created by the standard HK and the random-HK models.




Figure 3: Convergence time of the network-HK model for various values of ε on two types of networks.

tively regarding the number of clusters. The small difference in the number of clusters is rather insignificant, especially if one considers the rather unstable behavior of the HK model. Specifically, even small perturbations during the opinion formation can cause noticeable differences in the final clustering of the agents. To demonstrate this, we first simulate the standard HK model, and then, we repeat the process on the same initial opinion vector, but with some small random perturbations for few randomly selected rounds. The perturbations consists of increasing ε to 2ε for the selected round. Our simulations show that for $\varepsilon \leq 0.25$, these small perturbations may significantly change the resulting number of clusters.

The Network-HK Model. We start with the convergence time for different types of networks. We mostly focus on random $\mathcal{G}_{n,p}$ networks (see e.g., [Bollobas, 2001]) and Barabási-Albert (preferential attachment) graphs [Barabási and Albert, 1999]. We select 64 different values of $n \in [256, 1024]$ and create 10 relatively sparse networks for each value. We run the network-HK and the standard HK models on each of these networks for 10 values of $\varepsilon \in [.001, .4]$.

The resulting convergence times are shown in Figure 3. For both network types, the variance drops as ε increases. This happens because for large values of ε , the agents can more easily find other agents at opinion-distance at most ε , regardless of the network structure. As a result, the opinion formation process behaves more uniformly. The lower variance in Barabási-Albert networks (compared to $\mathcal{G}_{n,p}$ networks) is mostly due to their smaller diameter. Our simulations show

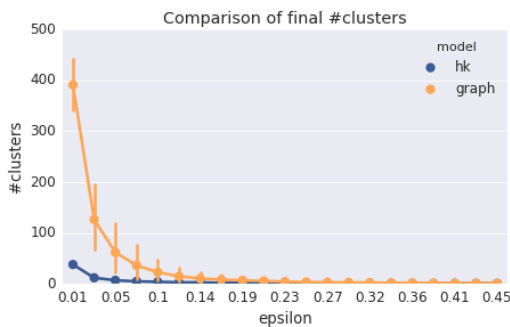


Figure 4: Comparison of the number of clusters created by the network-HK and the standard HK models.




Figure 5: Network-HK leads to many singleton clusters.

that the convergence time depends more on the value of ε , less on the network type, and not so much on the number of agents n . A notable exception is chain networks where, for large values of ε , the convergence time heavily depends on n .

As for the number of clusters, we repeat the simulations that we performed for the random-HK model. We now create random $\mathcal{G}_{n,p}$ networks with $p \in [\log n/n, .1]$. To count the clusters at the stable state, we delete every edge that connects agents at opinion-distance more than ε and count the number of connected components. The results are depicted in Figure 4. The absolute difference in the number of clusters is now quite large. It is larger than 300, for $\varepsilon = .01$, and slowly drops to 0, when both models reach consensus. We observe a similar behavior for sparse Barabási-Albert networks. To explain the cluster proliferation, we run both models for $\varepsilon = .15$ and a random $\mathcal{G}_{1000,.05}$ network. Figure 5 shows that the large number of clusters for network-HK is due to the fact that many agents (and especially agents with extreme opinions) fail to find any neighbors at opinion-distance at most ε , which leads to many singleton clusters.

Finally, we test the network-HK model on the Facebook “circles” network [Leskovec, 2012], with 4039 nodes and 88234 edges (see Figure 6a). Compared to the computer-generated networks before, this instance takes much longer to converge. To obtain a better understanding of the increased convergence time, we draw, at round $t = 40$, the agents who will eventually form a cluster with opinions in $[\cdot 2, \cdot 4]$. In Figure 6b, we see that the future cluster consists of tightly bound communities which communicate with each other through a small number of edges. As a result, information exchange is slow, which crucially affects the convergence time.

Acknowledgments: Stratis Skoulakis is partially supported by a scholarship from the Onassis Foundation.




Figure 6: Network-HK on the Facebook network.

References

- [Barabási and Albert, 1999] A.L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286(5439):509-512, 1999.
- [Bhattacharyya *et al.*, 2013] A. Bhattacharyya, M. Braverman, B. Chazelle, and H.L. Nguyen. On the convergence of the Hegselmann-Krause system. In *Proc. of 4th Innovations in Theoretical Computer Science Conference (ITCS '13)*, pp. 61-66, 2013.
- [Bhawalkar *et al.*, 2013] K. Bhawalkar, S. Gollapudi, and K. Munagala. Coevolutionary opinion formation games. In *Proc. of the 45th ACM Symposium on Theory of Computing (STOC '13)*, pp. 41-50, 2013.
- [Bindel *et al.*, 2011] D. Bindel, J.M. Kleinberg, and S. Oren. How bad is forming your own opinion? In *Proc. of the 52nd IEEE Symposium on Foundations of Computer Science (FOCS '11)*, pp. 57-66, 2011.
- [Blondel *et al.*, 2009] V.D. Blondel, J.M. Hendrickx, and J.N. Tsitsiklis. On Krause's Multi-Agent Consensus Model With State-Dependent Connectivity. *IEEE Transactions on Automatic Control*, 54(11):2586-2597, 2009.
- [Bollobas, 2001] B. Bollobas. *Random Graphs*. Cambridge University Press, 2001.
- [Carvalho and Larson, 2013] A. Carvalho and K. Larson. A consensual linear opinion pool. In *Proc. of the 23rd International Joint Conference on Artificial Intelligence (IJCAI '13)*, pp. 2518-2524, 2013.
- [Chazelle and Wang, 2015] B. Chazelle and C. Wang. Inertial Hegselmann-Krause systems. *CoRR*, abs/1502.03332, 2015.
- [Chazelle, 2011] B. Chazelle. The Total s -Energy of a Multi-agent System. *SIAM Journal on Optimization and Control*, 49(4):16801706, 2011.
- [Deffuant *et al.*, 2001] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch. Mixing beliefs among interacting agents. *Advances in Complex Systems*, 3:87-98, 2001.
- [DeGroot, 1974] M.H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69:118-121, 1974.
- [Ferraioli *et al.*, 2012] D. Ferraioli, P.W. Goldberg, and C. Ventre. Decentralized dynamics for finite opinion games. In *Proc. of the 5th Symposium on Algorithmic Game Theory (SAGT '12)*, LNCS 7615, pp. 144-155, 2012.
- [Fortunato, 2005] S. Fortunato. On the Consensus Threshold for the Opinion Dynamics of Krause-Hegselmann. *International Journal of Modern Physics C*, 16(2):259-270, 2005.
- [Friedkin and Johnsen, 1990] N.E. Friedkin and E.C. Johnsen. Social influence and opinions. *Journal of Mathematical Sociology*, 15(3-4):193-205, 1990.
- [Ghaderi and Srikant, 2014] J. Ghaderi and R. Srikant. Opinion dynamics in social networks with stubborn agents: Equilibrium and convergence rate. *Automatica*, 2014.
- [Hegselmann and Krause, 2002] R. Hegselmann and U. Krause. Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of Artificial Societies and Social Simulation*, 5, 2002.
- [Hendrickx and Blondel, 2006] J.M. Hendrickx and V.D. Blondel. Convergence of different linear and non-linear Vicsek models. In *Proc. of the 17th International Symposium on Mathematical Theory of Networks and Systems (MTNS '06)*, page 12291240, 2006.
- [Jackson, 2008] M.O. Jackson. *Social and Economic Networks*. Princeton University Press, 2008.
- [Kempe *et al.*, 2002] D. Kempe, J.M. Kleinberg, and A. Kumar. Connectivity and inference problems for temporal networks. *Journal of Computer and System Sciences*, 64(4):820-842, 2002.
- [Leskovec, 2012] J. Leskovec. Stanford network analysis package (snap). <http://snap.stanford.edu>, 2012.
- [Lever, 2010] C.R. Lever. *Strategic Competition over Networks*. PhD Thesis, Stanford University, 2010.
- [Lorenz and Urbig, 2007] J. Lorenz and D. Urbig. About the power to enforce and prevent consensus by manipulating communication rules. *Advances in Complex Systems*, 10(2):251269, 2007.
- [Lorenz, 2005] J. Lorenz. A stabilization theorem for dynamics of continuous opinions. *Physica A*, 355(1):217-223, 2005.
- [Seneta, 1979] E. Seneta. Coefficients of ergodicity: Structure and applications. *Advances in Applied Probability*, 11(3):576-590, 1979.
- [Tsang and Larson, 2014] A. Tsang and K. Larson. Opinion dynamics of skeptical agents. In *Proc. of the 13th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS '14)*, pp. 277-284, 2014.
- [Urbig and Lorenz, 2004] D. Urbig and J. Lorenz. Communication regimes in opinion dynamics: Changing the number of communicating agents. In *Proc. of the 2nd Conference of the European Social Simulation Association (ESSA '04)*, 2004.
- [Wedin and Hegarty, 2015] E. Wedin and P. Hegarty. A Quadratic Lower Bound for the Convergence Rate in the One-Dimensional Hegselmann-Krause Bounded Confidence Dynamics. *Discrete and Computational Geometry*, 53(2):478-486, 2015.
- [Weisbuch *et al.*, 2002] G. Weisbuch, G. Deffuant, F. Amblard, and J.-P. Nadal. Interacting Agents and Continuous Opinions Dynamics. *Lecture Notes in Economics and Mathematical Systems*, 521:225-242, 2002.
- [Yildiz *et al.*, 2013] E. Yildiz, A. Ozdaglar, D. Acemoglu, A. Saberi, and A. Scaglione. Binary opinion dynamics with stubborn agents. *ACM Transactions on Economics and Computation*, 1(4):19:1-19:30, 2013.