

Convergence in (Social) Influence Networks

Silvio Frischknecht, Barbara Keller, and Roger Wattenhofer

Computer Engineering and Networks Laboratory (TIK), ETH Zürich, Switzerland
{fsilvio, barkelle, wattenhofer}@tik.ee.ethz.ch

Abstract. We study the convergence of influence networks, where each node changes its state according to the majority of its neighbors. Our main result is a new $\Omega(n^2/\log^2 n)$ bound on the convergence time in the synchronous model, solving the classic “Democrats and Republicans” problem. Furthermore, we give a bound of $\Theta(n^2)$ for the sequential model in which the sequence of steps is given by an adversary and a bound of $\Theta(n)$ for the sequential model in which the sequence of steps is given by a benevolent process.

Keywords: Social Networks, Stabilization, Democrats and Republicans, Majority Function, Equilibrium.

1 Introduction

What do social networks, belief propagation, spring embedders, cellular automata, distributed message passing algorithms, traffic networks, the brain, biological cell systems, or ant colonies have in common? They are all examples of “networks”, where the entities of the network are continuously influenced by the states of their respective neighbors. All of these examples of *influence networks* (INs) are known to be difficult to analyze. Some of the applications mentioned are notorious to have long-standing open problems regarding convergence.

In this paper we deal with a generic version of such networks: The network is given by an arbitrary graph $G = (V, E)$, and all nodes of the graph switch simultaneously to the state of the majority of their respective neighbors. We are interested in the stability of such INs with a binary state. Specifically, we would like to determine whether an IN converges to a stable situation or not. We are interested in how to specify such a stable setting, and in the amount of time needed to reach such a stable situation. We study several models how the nodes take turns, synchronous, asynchronous, adversarial, benevolent.

Our main result is for synchronous INs: Each node is assigned an initial state from the set $\{R, B\}$, and in every round, all nodes switch their state to the state of the majority of their neighbors simultaneously. This specific problem is commonly referred to as “Democrats and Republicans”, see e.g. Peter Winkler’s CACM column [Win08]. It is well known that this problem stabilizes in a peculiar way, namely that each node eventually is in the same state every second round [GO80]. This result can be shown by using a potential bound argument, i.e., until stabilization, in each round at least one more edge becomes “more stable”. This directly gives a $\mathcal{O}(n^2)$ upper bound for the

convergence time. On the other hand, using a slightly adapted linked list topology, one can see that convergence takes at least $\Omega(n)$ rounds. But what is the correct bound for this classic problem? Most people that worked on this problem seem to believe that the linear lower bound should be tight, at least asymptotically. Surprisingly, in the course of our research, we discovered that this is not true. In this paper we show that the upper bound is in fact tight up to a polylogarithmic factor. Our new lower bound is based on a novel graph family, which has interesting properties by itself. We hope that our new graph family might be instrumental to research concerning other types of INs, and may prove useful in obtaining a deeper understanding of some of the applications mentioned above.

We complement our main result with a series of smaller results. In particular, we look at asynchronous networks where nodes update their states sequentially. We show that in such a sequential setting, convergence may take $\Theta(n^2)$ time if given an adversarial sequence of steps, and $\Theta(n)$ if given a benevolent sequence of steps.

2 Related Work

Influence networks have become a central field of study in many sciences. In biology, to give three examples from different areas, [RT98] study networks in the context of brain science, [AAB⁺11] study cellular systems and their relation to distributed algorithms, and [AG92] study networks in the context of ant colonies. In optimization theory, belief propagation [Pea82, BTZ⁺09] has become a popular tool to analyze large systems, such as Bayesian networks and Markov random fields. Nodes are continuously being influenced by their neighbors; repeated simulation (hopefully) quickly converges to the correct solution. Belief propagation is commonly used in artificial intelligence and information theory and has demonstrated empirical success in numerous applications such as coding theory. A prominent example in this context are the algorithms that classify the importance of web pages [BP98, Kle99]. In physics and mechanical engineering, force-based mechanical systems have been studied. A typical model is a graph with springs between pairs of nodes. The entire graph is then simulated, as if it was a physical system, i.e. forces are applied to the nodes, pulling them closer together or pushing them further apart. This process is repeated iteratively until the system (hopefully) comes to a stable equilibrium, [KK89, Koh89, FR91, KW01]. Influence networks are also used in traffic simulation, where nodes (cars) change their position and speed according to their neighboring nodes [NS92]. Traffic networks often use cellular automata as a basic model. A cellular automaton [Neu66, Wol02] is a discrete model studied in many fields, such as computability, complexity, mathematics, physics, and theoretical biology. It consists of a regular grid of cells, each in one of a finite number of states, for instance 0 and 1. Each cell changes its state according to the states of its neighbors. In the popular game of life [Gar70], cells can be either dead or alive, and change their states according to the number of alive neighbors.

Our synchronous model is related to cellular automata, on a general graph; however, nodes change their opinion according to the majority of their neighbors. As majority

functions play a central role in neural networks and biological applications this model was already studied during the 1980s. Goles and Olivos [GO80] have shown that a synchronous binary influence network with a generalized threshold function always leads to a fixed point or to a cycle of length 2. This means that after a certain amount of synchronous rounds, each participant has either a fixed opinion or changes its mind in every round. Poljak and Sura [PS83] extended this result to a finite number of opinions. In [GT83], Goles and Tchuente show that an iterative behavior of threshold functions always leads to a fixed point. Sauerwald and Sudholt [SS10] study the evolution of cuts in the binary influence network model. In particular, they investigate how cuts evolve if unsatisfied nodes flip sides probabilistically. To some degree, one may argue that we look at the deterministic case of that problem instead.

In sociology, understanding social influence (e.g. conformity, socialization, peer pressure, obedience, leadership, persuasion, sales, and marketing) has always been a cornerstone of research, e.g. [Kel58]. More recently, with the proliferation of online social networks such as Facebook, the area has become en vogue, e.g. [MMG⁺07, AG10]. Leskovec et al. [LHK10] for instance verify the balance theory of Heider [Hei46] regarding conformity of opinions; they study how positive (and negative) influence links affect the structure of the network. Closest to our paper is the research dealing with influence, for instance in the form of sales and marketing. For example, [LSK06] investigate a large person-to-person recommendation network, consisting of four million people who made sixteen million recommendations on half a million products, and then analyze cascades in this data set. Cascades can also be studied in a purely theoretical model, based on random graphs with a simple threshold model which is close to our majority function [Wat02]. Rumor spreading has also been studied algorithmically, using the random phone call model, [KSSV00, SS11, DFF11]. Using real data from various sources, [ALP12] show that networks generally have a core of influential (elite) users. In contrast to our model, nodes cannot change their state back and forth, once infected, a node will stay infected. Plenty of work was done focusing on the prediction of influential nodes. One wants to find subset of influential nodes for viral marketing, e.g. [KKT05, CYZ10]. In contrast, [KOW08] studies the case of competitors, which is closer to our model since nodes can have different opinions. However, also in [KOW08] nodes only change their opinion once. However, in all these social networks the underlying graph is fixed and the dynamics of the stabilization process takes place on the changing states of the nodes only. An interesting variant changes the state of the edges instead. A good example for this is matching. A matching is (hopefully) converging to a stable state, based on the preferences of the nodes, e.g. [GS62, KPS10, FKPS10]. Hoefer takes these edge dynamics one step further, as not only the state of the edge changes, but the edge itself [Hoe11].

3 Model Definition

An *influence network* (IN) is modeled as a graph $G = (V, E, o_0)$. The set of nodes V is connected by an arbitrary set of edges E . Each node has an initial opinion (or state) $o_0(v) \in \{R(ed), B(lue)\}$. A node only changes its opinion if a majority of its

neighbors has a different opinion. One may consider several options to breaking ties, e.g., using the node's current opinion as a tie-breaker, or weighing the opinions of individual neighbors differently. As it turns out, for many natural tie-breakers, graphs can be reduced to equivalent graphs in which no tie breaker is needed. For instance, using a node's own opinion as a tie-breaker is equivalent to cloning the whole graph, and connecting each node with its clone and the neighbors of its clone.

In this paper we study both synchronous and asynchronous INs. The state of a synchronous IN evolves over a series of rounds. In each round every node changes its state to the state of the majority of its neighbors simultaneously. The opinion of a node v in round t is denoted as $o_t(v)$.

As will be explained in Section 5, the only interesting asynchronous model is the sequential model. In this model, we call the change of opinion of one node a *step*. The opinion of node v after t steps is defined as $o_t(v)$. In general, more than one node may be ready to take a step. Depending on whether we want convergence to be fast or slow, we may choose different nodes to take the next step. If we aim for fast convergence, we call this the *benevolent sequential model*. Slow conversion on the other hand we call the *adversarial sequential model*.

We say that an IN stabilizes if it reaches a state where no node will ever change its opinion again, or if each node changes its opinion in a cyclic pattern with periodicity q . In other words, a state can be stable even though some nodes still change their opinion.

Definition 1. An IN $G = (V, E, o_0)$ is stable at time t with periodicity q , if for all vertices $v \in V : o_{t+q}(v) = o_t(v)$. A fixed state of an IN G is a stable state with periodicity 1. The convergence time c of an IN G is the smallest t for which G is stable.

Note that since INs are deterministic an IN which has reached a stable state will stay stable.

In this paper we investigate the stability, the convergence time c and the periodicity q of INs in the described models. Clearly, the convergence process depends not only on the graph structure, but also on the initial opinions of the nodes. We investigated graphs and initial opinions that maximize convergence time. In the benevolent sequential in particular, we investigate graphs and sets of initial opinions leading to the worst possible convergence time, given the respectively best sequence of steps.

4 Synchronous IN

A synchronous IN may stabilize in a state where some nodes change their opinion in every round. For example, consider the graph K_2 (two nodes, connected by an edge) where the first node has opinion B and the second node has opinion R . After one round, both vertices have changed their state, which leads to a symmetric situation. This IN remains in this stable state forever with a period of length 2. As has already been shown in [GO80, Win08], a synchronous IN always reaches a stable state with a periodicity of at most 2 after $\mathcal{O}(n^2)$ rounds.

Theorem 1 ([Win08]). *A synchronous influence network reaches a stable state after at most $\mathcal{O}(n^2)$ rounds.*

Theorem 2 ([GO80]). *The periodicity of the stable state of a synchronous influence network is at most 2.*

We prove this bound to be almost tight.

Theorem 3. *There exists a family of synchronous influence networks with convergence time of $\Omega\left(\frac{n^2}{(\log \log n)^2}\right)$.*

Due to page limitations the technical proof of Theorem 3 does not fit here, but can be found in the full version of this paper. In this section, we instead present a simpler IN with convergence time $\Omega(n^{3/2})$.

The basic idea is to construct a mechanism which forces vertices on a simple path graph to change their opinion one after the other. Every time the complete path has changed, the mechanism should force the vertices of the path to change their opinions back again in the same order. To create this mechanism, we introduce an auxiliary structure called transistor, which is depicted in Figure 1.




Fig. 1. A transistor $T(4)$. The dotted lines indicate how the transistor will be connected.

Definition 2. *A transistor of size k , denoted as $T(k)$, is an undirected graph consisting of three base vertices $B = \{B^1, B^2, B^3\}$, k collector vertices $C = \{C^i \mid 0 \leq i \leq k - 1\}$ and k emitter vertices $E = \{E^i \mid 0 \leq i \leq k - 1\}$. All edges between collector and emitter vertices, all edges between any two base vertices, and all edges between collector vertices and the third base vertex exist. Formally:*

$$\begin{aligned}
 T(s) &= (V, E) \\
 V &= C \cup E \cup B \\
 E &= \{ \{u, v\} \mid u \in C, v \in E \} \cup \{ \{u, B^3\} \mid u \in C \} \cup \\
 &\quad \{ \{u, v\} \mid u, v \in B, u \neq v \}
 \end{aligned}$$

All nodes in a transistor are initialized with the same opinion $X \in \{R = 1, B = -1\}$. The $3 + k + k^2$ collector edges (dotted edges pointing to the top of Figure 1, including those originating from $\mathcal{B}^1, \mathcal{B}^2$ and \mathcal{B}^3) are connected to vertices with the constant opinion $-X$, while up to $k^2 - k$ emitter edges (dotted edges pointing to the bottom) and the 2 base edges (dotted edges pointing to the left) may be connected to any vertex. As soon as both base edges advertise opinion $-X$, the transistor will flip to opinion $-X$ in 4 rounds regardless of what is advertised over the emitter edges, i.e., the following sets of vertices will all change their opinion to $-X$ in the given order: $\{\mathcal{B}^1\}, \{\mathcal{B}^2, \mathcal{B}^3\}, \mathcal{C}, \mathcal{E}$.




Fig. 2. Path with 4 vertices connected to one transistor $T(3)$




Fig. 3. Path with 4 vertices connected to 3 transistors $T(3)$. Note that transistors at bottom of figures are always upside down.




Fig. 4. Two copies of Figure 3 with inverse opinions




Fig. 5. In this graph, every time the path has run through completely the next transistor will flip, causing the path to run again

Note that $T(k)$ contains only $\mathcal{O}(k)$ many vertices, yet its emitter vertices can potentially be connected to $\Omega(k^2)$ other vertices. Given a path graph of length $\mathcal{O}(k^2)$ and a transistor $T(k)$, the emitter vertices of the transistor are connected to the path in the following way: The first vertex in the path is connected to exactly two emitter vertices, the last is connected to none and each of the remaining nodes of the path is connected to exactly one emitter vertex. Furthermore, the collector edges of transistors of opinion X are always connected to constant reservoirs of opinion $-X$. Such a reservoir can be implemented as a clique. An illustration of this graph with $k = 3$ is given in Figure 2. Without loss of generality, we set the initial state of the nodes of the path to B , and that




Fig. 6. Final graph in which the paths run 3 times. For an illustration with colors please refer to the electronic version of this paper.

of the transistor to R . As long as the transistor remains red, the path will turn red one vertex at a time. As soon as the transistor flips its opinion to blue (as a result of both base edges having advertised blue) the path will turn blue again, one vertex at a time. To force the path to change k times, k transistors are needed. Each of these transistors (note that we make use of red as well as blue transistors) is connected with the path in the same way as the first transistor. The resulting graph is given in Figure 3. A series of k switches of the complete path can now be provoked by switching transistors of alternating opinions in turns. For the example depicted in the Figures, the switching order of the transistors is given by their respective indices.

Now, a way is needed to flip the next transistor every time the last vertex of the path has changed its opinion. Assume the last vertex has changed to red. It is necessary to flip a red transistor to blue in order to change the path to blue; however, the path changing to red can only cause a blue transistor to turn red. To this end, the graph is extended by a copy of itself with all opinions inverted. The resulting graph is given in Figure 4. As in every round each vertex in the copy is of the opposite opinion than its original, the copy of the last vertex in the path enables us to flip a red transistor to blue as desired. The edges necessary to achieve this (highlighted in green in Figure 5) connect the end of a path to B^1 of each transistor in the other half of the graph. To ensure that the transistors flip in the required order, additional edges (highlighted in magenta in Figure 5) are introduced, connecting an emitter node of each transistor T_i^X to the node B^1 of transistor T_{i+1}^X .

The green edges cause an unwanted influence on the last vertex of the paths. This influence can be negated by introducing additional edges (highlighted in cyan in Figure 6). These edges connect the last vertex of each path with an emitter vertex of each transistor not yet connected to that vertex.

The resulting graph contains $\mathcal{O}(k^2)$ vertices, yet has a convergence time of $\Omega(k^3)$. In terms of the number of vertices n , the convergence time is $n^{3/2}$. The detailed proof in the full version of this paper shows that this technique can be applied to run the entire graph repeatedly, just as the graph in this section runs two paths repeatedly. This leads to a convergence time of $\Omega(n^{7/4})$. In this new graph, the transistors change back and fourth repeatedly, always taking on the opinion advertised over the collector edges, just like real transistors. When applied recursively $\log \log n$ times, an asymptotic convergence time of $\Omega(n^2/(\log \log n)^2)$ is reached. Since the full proof is long and involved, to complement our formal proof, we also simulated this recursively constructed networks for path lengths of up to 100. Table 1 and Figure 7 show the outcome of this simulation.

Table 1. Table summarizing the simulated results

path length	#nodes	convergence time
1	10	1
2	12	2
3	96	22
10	494	310
20	1614	3331
30	2010	5701
100	5518	45985




Fig. 7. Shows how our simulation results compare to a quadratic curve. The point clusters arise when for several consecutive path lengths no new transistor is created. Small jumps in the number of vertices indicate that a new transistor was added; big jumps indicate that a new layer of transistors was added.

5 Sequential IN

To complement our results for the synchronous model, we consider an asynchronous setting in this section. In an asynchronous setting, nodes can take steps independently of each other, i.e. subsets of nodes may reassess and change their opinion concurrently. Unfortunately, in such a setting, convergence time is not well defined. To see this, consider a star-graph where the center has a different initial opinion than the leaves. An adversary may arbitrarily often chooses the set of all nodes to reassess their opinion. After r such rounds the adversary chooses only the center node. Now this IN stabilizes, after r rounds for an arbitrary $r \rightarrow \infty$. In other words, asynchrony in its most general form is not well defined, and we restrict ourselves to sequential steps only, whereas a step is a single node changing its opinion. The sequence of steps is chosen by an adversary which tries to maximize the convergence time. Note that the convergence upper bound presented in Lemma 1 implies immediately that the IN stabilizes in a fixed state.

Lemma 1. *A sequential IN reaches a fixed state after at most $\mathcal{O}(n^2)$ steps.*

Proof. Divide the nodes into the following two sets according to their current opinion: $S_R = \{v \mid o(v) = R\}$ and $S_B = \{v \mid o(v) = B\}$. If a node changes its opinion, it has more neighbors in the opposite set than in its current set. Therefore the number of edges $X = \{\{u, v\} \mid u \in S_R, v \in S_B\}$ between nodes in set S_R and set S_B is strictly decreasing. Each change of opinion reduces the number of edges of X by at least one. Therefore the number of steps is bounded by the number of edges in X . In a graph G with n nodes $|X|$ is at most $n^2/4$, therefore at most $\mathcal{O}(n^2)$ steps can take place until the IN reaches a fixed state. \square

It is more challenging to show that this simple upper bound is tight. We show a graph and a sequence of steps in which way an adversary can provoke $\Omega(n^2)$ convergence time.

Lemma 2. *There is a family of INs with n vertices such that a fixed state is reached after $\Omega(n^2)$ steps.*

Algorithm 1. Adversarial Sequence

```

 $S \leftarrow ()$ 
for  $i = 0$  to  $n/3$  do
     $S = \text{reverse}(S)$ ;
     $S \leftarrow (i, S)$ ;
    for all  $x \in S$  do
        take step  $x$ ;
    end for
end for

```

Proof. Consider the following graph G with n nodes. The nodes are numbered from 0 to $n - 1$, whereas nodes with an even id are initially assigned opinion B and nodes with an odd id are assigned opinion R . See also Figure 8. All even nodes with $id \leq n/3$ are connected to all odd nodes. All odd nodes with $id \leq n/3$ are connected to all even nodes respectively. In addition an even node with $id \leq n/3$ is connected to nodes $\{0, 2, 4, \dots, n - 2 \cdot id - 2\}$, respectively an odd node with $id \leq n/3$ is connected to nodes $\{1, 3, 5, \dots, n - 2 \cdot id - 3\}$. For example, node 0 is a neighbor of all nodes, whereas node 1 is neighbor of all nodes except the nodes $n - 1$ and $n - 3$. Note that each node i with $i \leq n/3$ is connected to all other nodes with $id \leq n/3$. For each node v the change potential $P(v)$ is defined as:

$$P(v) = |\{u \mid o(u) \neq o(v)\}| - |\{u \mid o(u) = o(v)\}|$$

Put differently, if the change potential of a node is larger than 0, and it is requested to reassess its opinion, it takes a step. A large change potential of a node v , means that many neighbors of v have the opposite opinion from v . If a neighbor of v with the same




Fig. 8. In this graph an adversary can provoke $\Omega(n^2)$ changes of opinion

opinion takes a step, v 's change potential $P(v)$ is increased by 2. On the other hand, if a neighbor changes from the opposite opinion to the same opinion as node v , $P(v)$ is decreased by 2. If v itself changes its opinion, its change potential turns from p to $-p$. The change potential of v is basically the number of edges by which the total number of edges between set S_B and set S_R is reduced if v changes its opinion. As the total amount of steps is bounded by the number of edges between set S_B and S_R , a node v with $P(v) = p$ reduces the remaining number of possible changes by p if it takes a step. E.g. in the previously constructed graph G , the first nodes have the following change potential: $P(0) = 1, P(1) = 3, P(2) = 3, P(3) = 5$. Generally, node i has a change potential $P(i) = n/2 - (n/2 - i - 1) = i + 1$ if i is even respectively $P(i) = i + 2$ if i is odd. In order to provoke as many steps as possible, the adversary selects the nodes which have to reassess their opinion according to the following rule: He chooses the node with the smallest id for which $P(v) = 1$. Therefore each step reduces the remaining number of possible steps by 1. G is constructed in such a way, that a step from a node triggers a cascade of steps from nodes which have already changed their opinion whereas each change reduces the overall potential by 1.

The adversary chooses the nodes in phases according to algorithm 1. Phase i starts with the selection of node i followed by the selections of all nodes with $id < i$, where the adversary chooses the nodes in the reverse order than it did in round $i - 1$. Phase 0 consists of node 0 changing its opinion, in phase 1 node 1 and then node 0 make steps, and in phase 2 the nodes change in the sequence 2, 0, 1. As a node v can only change its opinion if $P(v) > 0$, we need to show that this is the case for each node v which is selected by the adversary. It is sufficient to show that each node which is selected has a change potential of 1.

We postulate:

- (i) At the beginning of phase i , it holds that: $P(i) = 1$ and $\forall v < i : o(v) = o(i)$.
- (ii) Each node the adversary selects has change potential 1 and each node with $id \leq i$ is selected eventually in phase i .
- (iii) At the end of phase i , all nodes with $id \leq i$ have opinion R if i is even and opinion B if i is odd.

We prove (i), (ii) and (iii) by induction. Initially, part (i) holds, as no node with $id < 0$ exists and as node 0 is connected to $n/2$ nodes with opinion R and to $n/2 - 1$ nodes with opinion B and therefore has change potential 1. In phase 0 only node 0 is selected, therefore part (ii) of holds as well. Node 0 changed its opinion and has therefore at the end of phase 0 opinion R , therefore part (iii) holds as well.

Now the induction step: To simplify the proof of part (i) of we consider odd and even phases separately. Consider an odd phase i . At the start of phase i , no node with $id \geq i$ has changed its opinion yet. Therefore node i still has its initial opinion $o(i) = R$. According to (iii), each node with $id \leq i - 1$ has at the end of phase $i - 1$ opinion $R = o(i)$. So $(i + 1)/2$ neighbors of i have compared to the initial state, changed their opinion from B to R . If a neighbor u of a node v with a different opinion than v changes it, v 's change potential is decreased by 2. Therefore node i 's initial change potential $P_{t_0}(i) = n/2 - (n/2 - i - 2) = i + 2$ is decreased by $2 \cdot (i + 1)/2 = i + 1$ and is therefore $P(i) = i + 2 - (i + 1) = 1$ at the beginning of phase i . Therefore (i) holds before an odd phase.

Now consider an even phase i . At its start, all nodes with $id \geq i$ still have their initial opinion. Therefore node i has opinion $o(i) = B$. According to (iii) each node with $id \leq i - 1$ has at the end of phase $i - 1$ opinion $B = o(i)$. As node i 's initial change potential was $P_{t_0}(i) = n/2 - (n/2 - i - 1) = i + 1$ and $i/2$ neighbors of i changed from opinion R to opinion B compared to the initial state, i 's new change potential is calculated as $P(i) = i + 1 - 2 \cdot i/2 = 1$. Therefore (i) holds before an even phase, hence (i) holds.

To prove part (ii) let v be the last node which was selected in phase $i - 1$. As v was selected, it had according to (ii) a change potential of 1. If a node changes its opinion, its change potential gets inversed. Therefore node v had at the beginning of phase i a change potential of -1 . In addition, node v is by construction a neighbor of node i and has according to (i) at the start of phase i the same opinion as node i . As node i changes its opinion, node v 's change potential is increased by 2. Therefore v 's new change potential is again $-1 + 2 = 1$, when it is selected by the adversary. The same argument holds for the second last selected node u . After it was selected in phase $i - 1$ its change potential was -1 . Then v has changed its opinion which led to $P(u) = -3$. As node i and node v changed their opinions in phase i , $P(u)$ was again 1. Hence if the adversary selects the nodes in the inverse sequence as in phase $i - 1$, each selected node has a change potential of 1 and is selected eventually. Therefore (ii) holds.

As node i and all nodes with $id \leq i - 1$ had at the beginning of phase i the opinion $o(i)$ according to (iii) and all nodes have changed their opinion in phase i according to (ii), all nodes with $id \leq i$ must have the opposite opinion at the end of phase i , namely R if i is even or B otherwise. Therefore (iii) holds as well.

We now have proven that in phase i , i nodes change their opinion. As the adversary starts $n/3$ phases, the total number of steps is $1/2 \cdot n/3 \cdot (n/3 - 1) \in \Omega(n^2)$. \square

Directly from Lemma 1 and Lemma 2, we get the following theorem.

Theorem 4. *A worst case sequential IN reaches a fixed state after $\Theta(n^2)$ steps.*

We have seen, that with an adapted graph and an adversary an IN takes up to $\Theta(n^2)$ steps until it stabilizes. But how bad can it get, if the process is benevolent instead?

Theorem 5. *An IN with a benevolent sequential process reaches a fixed state after $\Theta(n)$ steps.*

Proof. A benevolent process needs $\Omega(n)$ steps to reach a stable state. This can be seen by considering the complete graph K_n with initially $\lfloor n/2 \rfloor - 1$ red nodes and $\lceil n/2 \rceil + 1$ blue nodes. Independently of the chosen sequence this IN needs exactly $\lfloor n/2 \rfloor - 1$ steps to stabilize because the only achievable stable state is all nodes being blue. To prove that the number of steps is bounded by $\mathcal{O}(n)$ we define the following two sets: The set of all red nodes which want to change: $C_R^i = \{v \mid o(v) = R \wedge P(v) > 0\}$ and the set of all blue nodes which want to change: $C_B = \{v \mid o(v) = B \wedge P(v) > 0\}$. A benevolent process chooses nodes in two phases. In the first phase it chooses nodes from C_B until the set is empty. During this phase, it may happen that additional nodes join C_B (e.g. a leaf of a node $v \in C_B$, after v made a step). However, no node which left C_B will rejoin, as those nodes turned red and can not turn blue again in this phase. In the second phase, the benevolent process chooses nodes from C_R until this set is empty. The set C_B will stay empty during the second phase since nodes turning blue can only reinforce blue nodes in their opinion. Both phases take at most n steps, therefore proving our upper bound. \square

References

- [AAB⁺11] Afek, Y., Alon, N., Barad, O., Hornstein, E., Barkai, N., Bar-Joseph, Z.: A Biological Solution to a Fundamental Distributed Computing Problem. *Science*, 183–185 (2011)
- [AG92] Adler, F.R., Gordon, D.M.: Information Collection and Spread by Networks of Patrolling Ants. *The American Naturalist*, 373–400 (1992)
- [AG10] Asuncion, A.U., Goodrich, M.T.: Turning Privacy Leaks into Floods: Surreptitious Discovery of Social Network Friendships and Other Sensitive Binary Attribute Vectors. In: *Workshop on Privacy in the Electronic Society (WPES)*, pp. 21–30 (2010)
- [ALP12] Avin, C., Lotker, Z., Pigolet, Y.A.: On The Elite of Social Networks. In: *Personal Communication* (2012)
- [BP98] Brin, S., Page, L.: The Anatomy of a Large-Scale Hypertextual Web Search Engine. In: *Seventh International World-Wide Web Conference, WWW* (1998)
- [BTZ⁺09] Bickson, D., Tock, Y., Zymnis, A., Boyd, S.P., Dolev, D.: Distributed Large Scale Network Utility Maximization. In: *International Conference on Symposium on Information Theory, ISIT* (2009)
- [CYZ10] Chen, W., Yuan, Y., Zhang, L.: Scalable Influence Maximization in Social Networks under the Linear Threshold Model. In: *Industrial Conference on Data Mining (ICDM)*, pp. 88–97 (2010)

- [DFF11] Doerr, B., Fouz, M., Friedrich, T.: Social Networks Spread Rumors in Sublogarithmic Time. In: Symposium on Theory of Computing (STOC), pp. 21–30 (2011)
- [FKPS10] Floréen, P., Kaski, P., Polishchuk, V., Suomela, J.: Brief Announcement: Distributed Almost Stable Marriage. In: Principles of Distributed Computing (PODC), pp. 281–282 (2010)
- [FR91] Fruchterman, T.M.J., Reingold, E.M.: Graph Drawing by Force-Directed Placement. In: Software: Practice and Experience, pp. 1129–1164 (1991)
- [Gar70] Gardner, M.: Mathematical Games: The Fantastic Combinations of John Conway’s new Solitaire Game “Life”. *Scientific American*, 120–123 (1970)
- [GO80] Goles, E., Olivos, J.: Periodic Behaviour of Generalized Threshold Functions. *Discrete Mathematics*, 187–189 (1980)
- [GS62] Gale, D., Shapley, L.S.: College Admissions and the Stability of Marriage. *The American Mathematical Monthly*, 9–15 (1962)
- [GT83] Goles, E., Tchuente, M.: Iterative Behaviour of Generalized Majority Functions. In: *Mathematical Social Sciences*, pp. 197–204 (1983)
- [Hei46] Heider, F.: Attitudes and Cognitive Organization. *Journal of Psychology*, 107–112 (1946)
- [Hoe11] Hoefler, M.: Local Matching Dynamics in Social Networks. In: Aceto, L., Henzinger, M., Sgall, J. (eds.) *ICALP 2011, Part II. LNCS*, vol. 6756, pp. 113–124. Springer, Heidelberg (2011)
- [Kel58] Kelman, H.C.: Compliance, Identification, and Internalization: Three Processes of Attitude Change. *Journal of Conflict Resolution*, 51–60 (1958)
- [KK89] Kamada, T., Kawai, S.: An Algorithm for Drawing General Undirected Graphs. *Information Processing Letters*, 7–15 (1989)
- [KKT05] Kempe, D., Kleinberg, J.M., Tardos, É.: Influential Nodes in a Diffusion Model for Social Networks. In: Caires, L., Italiano, G.F., Monteiro, L., Palamidessi, C., Yung, M. (eds.) *ICALP 2005. LNCS*, vol. 3580, pp. 1127–1138. Springer, Heidelberg (2005)
- [Kle99] Kleinberg, J.M.: Hubs, Authorities, and Communities. *ACM Computing Surveys*, CSUR (1999)
- [Koh89] Kohonen, T.: *Self-Organization and Associative Memory*. Springer (1989)
- [KOW08] Kostka, J., Oswald, Y.A., Wattenhofer, R.: Word of mouth: Rumor dissemination in social networks. In: Shvartsman, A.A., Felber, P. (eds.) *SIROCCO 2008. LNCS*, vol. 5058, pp. 185–196. Springer, Heidelberg (2008)
- [KPS10] Kipnis, A., Patt-Shamir, B.: On the Complexity of Distributed Stable Matching with Small Messages. *Distributed Computing*, 151–161 (2010)
- [KSSV00] Karp, R., Schindelhauer, C., Shenker, S., Vocking, B.: Randomized Rumor Spreading. In: *Foundations of Computer Science, FOCS (2000)*
- [KW01] Kaufmann, M., Wagner, D.: *Drawing Graphs: methods and models*. Springer (2001)
- [LHK10] Leskovec, J., Huttenlocher, D.P., Kleinberg, J.M.: Signed Networks in Social Media. In: *Conference on Human Factors in Computing Systems (CHI)*, pp. 1361–1370 (2010)
- [LSK06] Leskovec, J., Singh, A., Kleinberg, J.M.: Patterns of Influence in a Recommendation Network. In: Ng, W.-K., Kitsuregawa, M., Li, J., Chang, K. (eds.) *PAKDD 2006. LNCS (LNAI)*, vol. 3918, pp. 380–389. Springer, Heidelberg (2006)
- [MMG⁺07] Mislove, A., Marcon, M., Gummadi, K.P., Druschel, P., Bhattacherjee, B.: Measurement and Analysis of Online Social Networks. In: *7th ACM SIGCOMM Conference on Internet Measurement*, pp. 29–42 (2007)
- [Neu66] Von Neumann, J.: *Theory of Self-Reproducing Automata*. University of Illinois Press (1966)

- [NS92] Nagel, K., Schreckenberg, M.: A Cellular Automaton Model for Freeway Traffic. *Journal de Physique I*, 2221–2229 (1992)
- [Pea82] Pearl, J.: Reverend Bayes on Inference Engines: A Distributed Hierarchical Approach. In: *Second National Conference on Artificial Intelligence (AAAI)*, pp. 133–136 (1982)
- [PS83] Poljak, S., Sra, M.: On Periodical Behaviour in Societies with Symmetric Influences. *Combinatorica*, 119–121 (1983)
- [RT98] Rolls, E.T., Treves, A.: *Neural Networks and Brain Function*. Oxford University Press, USA (1998)
- [SS10] Sauerwald, T., Sudholt, D.: A Self-Stabilizing Algorithm for Cut Problems in Synchronous Networks. *Theoretical Computer Science*, 1599–1612 (2010)
- [SS11] Sauerwald, T., Stauffer, A.: Rumor Spreading and Vertex Expansion on Regular Graphs. In: *Symposium on Discrete Algorithms (SODA)*, pp. 462–475 (2011)
- [Wat02] Watts, D.J.: A Simple Model of Global Cascades on Random Networks. *Proceedings of the National Academy of Sciences*, 5766–5771 (2002)
- [Win08] Winkler, P.: *Puzzled: Delightful Graph Theory*. *Communications of the ACM*, 104 (2008)
- [Wol02] Wolfram, S.: *A New Kind of Science*. Wolfram Media (2002)