

Maximizing the Spread of an Opinion in Few Steps: Opinion Diffusion in Non-Binary Networks

Robert Brederbeck*, Lilian Jacobs and Leon Kellerhals

Technische Universität Berlin, Chair of Algorithmics and Computational Complexity
{robert.bredereck, leon.kellerhals}@tu-berlin.de, mail2lilian.jacobs@gmail.com

Abstract

We consider the setting of asynchronous opinion diffusion with majority threshold: given a social network with each agent assigned to one opinion, an agent will update its opinion if more than half of its neighbors agree on a different opinion. The stabilized final outcome highly depends on the sequence in which agents update their opinion. We are interested in *optimistic sequences*—sequences that maximize the spread of a chosen opinion. We complement known results for two opinions where optimistic sequences can be computed in time and length linear in the number of agents. We analyze upper and lower bounds on the length of optimistic sequences, showing quadratic bounds in the general and linear bounds in the acyclic case. Moreover, we show that in networks with more than two opinions determining a spread-maximizing sequence becomes intractable; surprisingly, already with three opinions the intractability results hold in highly restricted cases, e.g., when each agent has at most three neighbors, when looking for a short sequence, or when we aim for approximate solutions.

1 Introduction

Analyzing the diffusion of opinions in social networks is an important topic with obvious applications to marketing, disease spreading, and news propagation (we refer to a recent survey by Grandi [2017] for a comprehensive overview on the general topic of opinion diffusion in social networks).

As up-to-date social network data is becoming more and more readily available, as urgent is attaining an understanding for diffusion processes, as well as having algorithms that help to predict how opinions change over time. In particular, being interested in a specific opinion, answering the following questions is critical, and the main focus of our paper.

(1) What is the maximum possible spread of the opinion in the network? (2) How many update steps are needed to reach a stable state with maximum spread of some opinion? (3) What is the computational complexity of computing the maximum spread of the opinion in the network?

Exhaustive work was done with respect to these three questions with main focus on the case of two opinions (in the following called black and white) and when agents update their opinions according to the majority of their neighbors. Goles and Olivos [1980] showed that a sequence of synchronous updates always converges to a stable state or cycles with period two. Moreover, their analysis provides an upper bound of $\mathcal{O}(n^2)$ on the number of updates with n being the number of agents. Frischknecht *et al.* [2013] showed that this quadratic bound is essentially tight, both for the synchronous model and for the asynchronous model. Note that the number of update steps for the synchronous update only depends on the network and the initial opinions. For the lower bound on the number of asynchronous update steps, one assumes that the sequence of updates is selected adversarially. In the synchronous model it is NP-hard to decide if there exists an initial opinion assignment for some given network such that convergence (to the two-periodic stable state) takes more than k rounds [Kaaser *et al.*, 2016]. Kralovic [2001] studies a variant where changes from black to white are irreversible and where the social network is acyclic. He develops an algorithm that tests whether a given configuration on an arbitrary tree evolves into an all-black state; the algorithm runs in time $\mathcal{O}(t \log t)$, where t is the number of black nodes.

Probably closest to our work is that of Brederbeck and Elkind [2017], in which all three questions are answered for the binary case. In particular, they show that one can compute an asynchronous update sequence that maximizes the number of black nodes in $\mathcal{O}(n)$ time. The sequence computed by their algorithm is of length at most $2n$. Moreover, to maximize the number of black nodes, every agent needs to change its opinion at most twice.

Opinions have been modeled binary [Goles and Olivos, 1980], ternary [Auletta *et al.*, 2019], as arbitrary finite set with a distance metric between the opinions [Chierichetti *et al.*, 2013], via continuous values from the interval $[0, 1]$ [Clifford and Sudbury, 1973], by rankings of alternatives [Brill *et al.*, 2016], or even via collections of propositional formulas [Grandi *et al.*, 2015; Slavkovik and Jamroga, 2016]. With respect to our three questions, however, little is known beyond the case of binary opinions. The result of Goles and Olivos [1980] for synchronous updates was subsequently extended to more general settings including an arbitrary number of opinions as well as weights on opinions and rela-

*Contact Author

tions between agents [Goles and Tchuente, 1983]. Auletta *et al.* [2019] have considered the case of three opinions and showed that computing a sequence that maximizes the spread a specific opinion becomes more involved with a third opinion. They show, first, that a greedy algorithm as used in the binary case may fail dramatically, and second, that finding an update sequence that maximizes the number of black nodes is NP-hard.

In this paper we contribute to our three main questions with the following results. First, we extend the results of Frischknecht *et al.* [2013] and show that an update sequence is of length at most quadratic in the number of agents even if there are more than two opinions. Further we show that the lower bound of $\Omega(n^2)$ on the number of updates holds even for spread-maximizing update sequences. Moreover we show that in acyclic influence networks, the length of a spread-maximizing update sequence may be linear in the number of agents, and there may exist an agent whose opinion is changed a linear number of times. Towards the third question, we show multiple hardness results: First of all we show that there is (presumably) no efficient algorithm for finding a spread-maximizing update sequence on influence networks with more than two opinions, even if we require the sequence length to be short, or if we want to find an approximate spread-maximizing update sequence, or if every agent has at most three neighbors.

2 Model and Notation

An (undirected) k -opinion influence network (InfNet) is a pair (G, \circ) where $G = (V(G), E(G))$ is a simple, undirected graph and $\circ: V(G) \rightarrow \{0, \dots, k-1\}$ is an *opinion mapping* that describes the initial opinions of all nodes. We set $n = |V(G)|$ and $m = |E(G)|$. Without loss of generality we only consider InfNets (G, \circ) where G is connected.¹ For simplicity of illustration, if $k = 3$, then we say node v is *black* if $\circ(v) = b = 0$, *white* if $\circ(v) = w = 1$ and *gray* if $\circ(v) = g = 2$. For a non-black opinion $x \in \{w, g\}$ we denote by $\bar{x} \in \{w, g\} \setminus \{x\}$ the non-black opinion opposing x .

For a node $v \in V(G)$, we denote by $N(v) := \{u \in V(G) \mid \{u, v\} \in E(G)\}$ the (*open*) *neighborhood* of v , that is, the set of all nodes that are connected to v by an edge. The degree of a node is the number of its neighbors. The maximum degree of G is the highest degree of a node in G .

During an *update step*, a node v in an influence network G changes its opinion if more than half of the nodes in its neighborhood $N(v)$ have the same opinion which differs from the opinion of v . Updates are performed asynchronously, that is, the updates are performed one after another. We describe the order of updates by a sequence σ of nodes; the i -th node considers changing its opinion at step i . We denote by $\circ[G, \sigma, z]$ the opinion mapping resulting from \circ after z updates following the sequence σ . We use $\circ[G, \sigma]$ as shortcut for the final outcome $\circ[G, \sigma, |\sigma|]$. A node v is *stable* if not more than half of the nodes in the neighborhood agree on a different opinion, and unstable otherwise. An opinion mapping \circ is stable in a

¹If the InfNet consists of more than one connected component, we can find an optimistic sequence for each InfNet separately and concatenate them in an arbitrary order.

network G if every node is stable. A sequence is stable if the corresponding final outcome is stable. A node v is *permanently stable* if there exists no sequence such that the opinion of v is changed. We say that an update sequence is *optimistic* if it is stable and maximizes the number of black nodes in the final outcome.

3 Lengths of Optimistic Sequences

This section deals with bounds on the lengths of optimistic update sequences. First, we provide upper bounds on the length of stable update sequences (which, of course, also bound the length of optimistic sequences). Second, we show lower bounds on the length of optimistic sequences. We provide a tight linear bound for acyclic influence networks and a tight quadratic bound for general influence networks.

3.1 Upper Bounds

As a warm-up we generalize two sequence length bounds provided by Frischknecht *et al.* [2013]. We first show that no sequence can be longer than the number of edges, essentially following the proof line of Frischknecht *et al.* [2013].

Proposition 1. *A k -opinion InfNet on m edges reaches stability after at most m updates.*

Frischknecht *et al.* proved the existence of 2-opinion InfNets that do not admit stable update sequences shorter than $\Omega(n)$ (clearly this holds for $k > 2$ as well). We provide an algorithm to match this lower bound for any k .

Proposition 2. *For a k -opinion InfNet one can compute a stable update sequence of length at most n in $\mathcal{O}(kn)$ time.*

Proof. We start with a greedy sequence σ : For $i = 0, 1, \dots, (k-1)$, in that order, update all nodes that can be updated to opinion i . Observe that nodes being updated to i either could have updated to i at any prior point in σ (this is always true for the first node to update to i), or can update to i only after some of their neighbors updated to i . We call these nodes fixed, resp. loose. Now we do the following for $i = (k-1), \dots, 0$, for every loose node v being updated to i , for every opinion $j < i$: If v is updated to j , then remove this update, and remove all updates of nodes v' that cannot be updated to j due to this removal. Note that v' indeed is stable in this point of the sequence. After completing this for i , the nodes that are updated to $j \geq i$ are updated at most once and are all stable at the end of the sequence. As their update can only depend on updates of nodes to j , they are not affected by the following iterations. After having iterated through all i , every node was updated at most once; thus the sequence length is at most n . The claimed running time follows trivially. \square

Note that Proposition 2 does not show a linear upper bound for *optimistic* update sequences. Indeed, in the next section, we show that to maximize the number of black nodes, one may indeed need to update a quadratic number of nodes.




Figure 1: The construction of the acyclic influence network for which every optimistic update sequence requires $\Omega(n)$ updates, and in which the node x is updated $\Omega(n)$ times. Thick lines represent multiple edges to multiple subtrees and the edge label gives the multiplicity.

3.2 Lower Bounds

We now focus on *optimistic* update sequences and show lower bounds on the required number of updates on general and acyclic influence networks. For both results we only require three opinions, so throughout this section we assume that the opinion mapping maps to the opinions black, white and gray.

We start showing a tight linear lower bound on the minimum length of optimistic sequences for acyclic influence networks. While just a linear lower bound for acyclic influence networks easy to see, our construction has the interesting property that it contains one node that needs to update its opinion a $\Omega(n)$ times.

Theorem 1. *There exists a family of acyclic 3-opinion InfNets on $\mathcal{O}(n)$ nodes for which every optimistic update sequence updates at least one node $\Omega(n)$ times.*

Proof sketch. Let $c \in \mathbb{N}$. We construct an acyclic InfNet with $\mathcal{O}(c)$ nodes. First, we introduce a *root* node x with $\circ(x) = w$. Attached to x are $c + 1$ absorbing gadgets with opinion g , c absorbing gadgets with opinion w and two emitting gadgets.

An absorbing gadget with opinion z is constructed as follows. The root has opinion z , and two children: one path on two nodes with opinion \bar{z} and one child with opinion z . The latter also has two children: one path on two nodes with opinion \bar{z} and one child with opinion b . Attached to the latter is a path on two nodes with opinion z .

An emitting gadget is constructed as follows. The root has opinion g , and $2c$ children with opinion g , each of which have two paths on two nodes with opinion w attached to them, as well as $2c + 2$ children with opinion w , each of which have two paths on two nodes with opinion g attached to them. Figure 1 shows the construction of the influence network.

Observe first that the black opinion cannot spread within this network. Next, note that in order to make the black node within an absorbing gadget with opinion z stable, the opinion of the two ancestors with opinion z must be changed to \bar{z} . In order to do so, node x must acquire opinion \bar{z} . Note though that only one absorbing gadget with opinion z can be updated at a time, since otherwise, x cannot update to z and the black nodes of all absorbing gadgets with opinion \bar{z} cannot be made stable. Thus, x must change its opinion to white c times, and to gray c times. The task of the emitting gadgets is to change the opinion of their respective root node such that the majority of neighbors of x can be shifted from opinion z to \bar{z} . In order to change the opinion of the root of an emitting gadget from z to \bar{z} , two of its children with opinion z must update their opinion (recall that these children have two children on their own with opinion \bar{z}). \square

Note that, since in acyclic InfNets we have $n = m - 1$, the lower bound in Theorem 1 is tight up to a constant.

We now present a family of influence networks for which the length of any shortest optimistic update sequence is quadratic in the number of nodes. One component of our network is the two-opinion network (called *snake graph* in the following) introduced by Frischknecht *et al.* [2013] (Lemma 2) to show that there exists a (not necessarily optimistic) update sequence whose length is quadratic in the number of its nodes. Our goal is to enforce that every optimistic update sequence must have $\Omega(n^2)$ updates. To this end, we add *synchronization gadgets* to the snake graph. We first present the construction of the two parts.

Snake graph. Let $n \in \mathbb{N}$. We construct the following graph G on $6n$ nodes v_0, \dots, v_{6n-1} : If $i \leq 2n - 1$, then we connect v_i to v_j for all j whose parity is different from the parity of i . For all even $i \leq 2n - 2$ we connect v_i to the nodes $\{v_0, v_2, v_4, \dots, v_{6n-2i-2}\}$. For all odd $i \leq 2n - 1$ we connect v_i to the nodes $\{v_1, v_3, v_5, \dots, v_{6n-2i-3}\}$. Finally, the nodes v_i with even $2n \leq i \leq 6n - 2$ form a clique, as do the nodes v_i with odd $2n + 1 \leq i \leq 6n - 1$. Initially, $\circ(v_i) = g$ if i is even, otherwise $\circ(v_i) = w$.

For a node v we define the *potential* $P(v)$ as

$$P(v) = \max_{x \in \{b, w, g\} \setminus \{\circ(v)\}} |\{u \in N(v) \mid \circ(u) = \circ(x)\}| - |\{u \in N(v) \mid \circ(u) = \circ(v)\}|.$$

In our InfNet we initially have $P(v_i) = i + 1$ if i is even, and $P(v_i) = i + 2$ otherwise. Since the nodes v_0, \dots, v_{2n-1} induce a complete graph we can observe the following.

Observation 1. *Let $0 \leq i \leq 2n - 1$. If v_i changes its opinion to $x \in \{w, g\}$, then for all $0 \leq j \leq 2n - 1$, $j \neq i$, $P(v_j)$ increases by two if $\circ(v_j) \neq x$ and decreases by two otherwise.*

Synchronization gadgets. The adversary sequence presented by Frischknecht *et al.* performs $2n^2 + n$ updates on v_0, \dots, v_{2n-1} in a snake-like order. We now present the construction of the *synchronization gadgets* that ensure that every optimistic update sequence has at least $\Omega(n^2)$ updates.

For $0 \leq i \leq n - 1$ we introduce a *white synchronization gadget* SG_i^w and a *gray synchronization gadget* SG_i^g .

Let $x \in \{g, w\}$. Then, for $0 \leq i \leq n - 1$ the synchronization gadget SG_i^x is constructed as follows: We introduce two main nodes a_i^x and b_i^x with $\circ(a_i^x) = x$ and $\circ(b_i^x) = \bar{x}$. Both a_i^x and b_i^x have edges to the nodes in $N_i = c_0 \cup c_{i-1} \cup c_i$ (thus $N_0 = c_0$, $N_1 = c_0 \cup c_1$, and for $i \geq 2$, $|N_i| = 6$). We add $|N_i| + 4$ paths on two nodes with opinion \bar{x} , and connect one endpoint of each of the paths to a_i^x . Analogously, we add $|N_i|$ paths on two nodes with opinion \bar{x} , and connect one endpoint for each of the paths to b_i^x . Next, we add a path on three nodes with one black endpoint and the other two nodes




Figure 2: The synchronization gadgets $SG_i(w)$ and $SG_i(g)$ for $0 \leq i \leq 2$. Numbers on edges depict the number of copies of the branch attached to the nodes a_i and b_i , i.e., there are six paths of length two attached to a_1^w .

having opinion x , and add an edge between the black endpoint and a_i^x . We do the same for b_i^w , but the two non-black nodes have opinion \bar{x} . Lastly, we add two edges that start in b_i^x . If $x = g$ then we add edges to b_i^w and to a_i^w . If $x = w$ then we add edges to b_{i-1}^g and to a_{i-1}^g . For the special case of b_0^w , we add a path on two nodes with opinion g and a path on two nodes with opinion w , and connect one endpoint of each to b_0^w . Finally, we add one path on two nodes with opinion g and connect one endpoint to b_{n-1}^g . Figure 2 shows the first six synchronization gadgets for $i \leq 2$.

Sequence length. The general idea of the synchronization gadgets is the following. Observe first that the black opinion cannot spread within our network. In order to save the two black nodes of a synchronization gadget SG_i^x for $x \in \{w, g\}$ and $0 \leq i \leq n-1$, the node a_i^x must change its opinion to \bar{x} , and the node b_i^x must change its opinion to x . If b_i^x is updated to x , then we say that the synchronization gadget SG_i^x is *fulfilled*. But this is only possible if all neighbors of b_i^x in the snake graph have opinion x , and if all synchronization gadgets are updated in the correct order.

We first show that there actually exists an update sequence that preserves all black nodes in our constructed InfNet.

Lemma 1. *There is an optimistic update sequence that keeps all black nodes stable.*

Proof sketch. We denote by $\bar{\sigma}$ the reverse of the sequence σ . The sequence is the concatenation

$$\sigma_0^w \circ (b_0^w, a_0^w) \circ \sigma_0^g \circ (b_0^g, a_0^g) \circ \sigma_1^w \circ (b_1^w, a_1^w) \circ \sigma_1^g \circ (b_1^g, a_1^g) \circ \dots \circ \sigma_{n-1}^w \circ (b_{n-1}^w, a_{n-1}^w) \circ \sigma_{n-1}^g \circ (b_{n-1}^g, a_{n-1}^g),$$

where $\sigma_0^w = (v_0)$,

$$\sigma_i^g = (v_{2i-1}) \circ \bar{\sigma}_i^w, \text{ for } 0 \leq i \leq n-1, \text{ and}$$

$$\sigma_i^w = (v_{2i}) \circ \bar{\sigma}_{i-1}^g, \text{ for } 0 \leq i \leq n-1. \quad \square$$

We now show that the black nodes can only be rescued completely only if all synchronization gadgets are fulfilled in the correct order. For that we introduce some auxiliary notation: Let $Y = \{a_i^x, b_i^x \mid 1 \leq i \leq n-1, x \in \{w, g\}\}$. Next, for $1 \leq i \leq n-1$ and $x \in \{w, g\}$ we denote by $N_Y(b_i^x) = N(b_i^x) \cap Y$ the set of “a” and “b” nodes that are neighbors of b_i^x .

Lemma 2. *Every optimistic update sequence updates a_i^g and b_i^g to white, resp. gray, after a_i^w and b_i^w are updated to gray, resp. white, for $0 \leq i \leq n-1$. Every optimistic update sequence updates $a_{i'}^w$ and $b_{i'}^w$ to gray, resp. white, after $a_{i'-1}^g$ and $b_{i'-1}^g$ are updated to white, resp. gray, for $0 < i' \leq n-1$.*

Proof sketch. We prove this lemma by proving the following two statements for a fixed j , $2 \leq j \leq n-1$:

(1) If a_j^w and b_j^w must be updated after a_{j-1}^g and a_{j-1}^g , then a_{j-1}^g and b_{j-1}^g must be updated after a_{j-1}^w and b_{j-1}^w .

(2) If a_j^g and b_j^g must be updated after a_j^w and b_j^w , then a_j^w and b_j^w must be updated after a_{j-1}^g and b_{j-1}^g .

We show how to prove (1) only; the correctness of statement (2) can be proven analogously. We first postulate first that, if b_{j-1}^w is updated before b_{j-1}^g , then a_{j-1}^w must be updated before b_{j-1}^g , and a_{j-1}^g must be updated after b_{j-1}^g : In order to set $\circ(b_{j-1}^g) = g$ we require that all but one of the nodes in $N_{j-1} \cup N_Y(b_{j-1}^g)$ are gray (this holds since $|N_Y(b_{j-1}^g)| = 4$, and b_{j-1}^g has $|N_{j-1}|$ stable white neighbors, and a black neighbor). Note that, by assumption, $\circ(b_{j-1}^w) = w$; thus we must have $\circ(u) = g$ for all $u \in N_{j-1} \cup N_Y(b_{j-1}^g) \setminus \{b_{j-1}^w\}$, and specifically the nodes a_{j-1}^g and a_{j-1}^w must be gray (we know $\circ(b_{j-1}^w) = g$ by assumption of statement (1)). So a_{j-1}^g must have already been updated, and a_{j-1}^w can only be updated after b_{j-1}^g .

So to prove the correctness of (1), it suffices to show that b_{j-1}^g is updated after b_{j-1}^w . Assume towards a contradiction that this is not the case. In order to set $\circ(b_{j-1}^w) = w$ we require that all but one of the nodes in $N_{j-1} \cup N_Y(b_{j-1}^w)$ are white. But $\circ(b_{j-1}^g) = g$; thus the rest of the nodes in $N_{j-1} \cup N_Y(b_{j-1}^w)$ must be white; specifically, $\circ(a_{j-1}^w) = \circ(a_{j-2}^g) = \circ(b_{j-2}^g) = w$. But then, after the update of b_{j-1}^w , its neighbor b_{j-2}^g must be updated to gray. Again, for this we require that all but one of the nodes in $N_{j-2} \cup N_Y(b_{j-2}^g)$ are gray. Observe that the two nodes in the latter set, a_{j-2}^g and b_{j-1}^w are both white. Clearly, the former of the two is stable, and from this it follows that the latter is stable as well. Hence, b_{j-2}^g is stable and cannot be updated to gray; thus the update sequence is not optimistic—a contradiction.

We use the two statements as steps of an induction, the base case of which is that a_{j-1}^g and b_{j-1}^g must be updated after a_{j-1}^w and b_{j-1}^w . The correctness of the base case is similar to the correctness of statement (1) with $j = n$; note that b_{n-1}^g does not have a neighbor b_n^w , but it does have a stable gray neighbor in its place. The last step of the induction (showing that a_0^w and b_0^w are updated before all other nodes in Y) uses statement (2) with $j = 1$; note that b_0^w does not have neighbors a_{-1}^g and b_{-1}^g , but it does have a stable gray neighbor and a stable white neighbor in their place. \square

Let us introduce the notion of *phases*. For $0 \leq i \leq n-1$, phase i ends with the update of b_i^w , that is, when SG_i^w is fulfilled. Phase 0 starts before the first update of the update sequence. For $1 \leq i \leq n-1$, phase i starts after the last update of phase $i-1$.

Lemma 3. *In phase i there are $\Omega(i)$ many updates within the snake graph.*

Proof sketch. Throughout this proof we assume that node a_i^x was updated immediately after b_i^x . Then, within the snake gadget, the potential of the nodes v_0, \dots, v_{2n-1} is not affected by the opinions of the nodes in the snake gadgets. Note though that if a_i^x is updated at another point in the sequence, this changes the number of required updates in phase i only by a constant.

Let $n_x = |\{v_i \mid 0 \leq i \leq 2n-1 \wedge \circ(v_i) = x\}|$ be the number of nodes within v_0, \dots, v_{2n-1} with opinion x . We take a close look at v_{2i-1} just after phase $i-1$ ended. Note that at the end of phase $i-1$, the nodes $v_0, v_1, v_{2i-2}, v_{2i-1}, v_{2i}, v_{2i+1}$ must have opinion white. Let $x = P(v_{2i-1}) = 2i + n_g - n_w + 1$. Note that $n_w - n_g = 2i + 1 - x$. So

$$\begin{aligned} P(v_{2i-2}) &= -(2i - 2 + n_w - n_g + 1) \\ &= -(2i - 2 + 2i + 1 - x + 1) = -4i + x. \end{aligned}$$

We consider three cases and show that in each case we require $\Omega(i)$ updates during phase i .

(1) *Suppose $x \leq 1$.* Then $P(v_{2i-2}) \leq -4i + x$, and in order to update v_{2i-2} to gray (which is required to fulfill $\text{SG}_i(g)$), at least $2i = \Omega(i)$ nodes must be updated to gray to make v_{2i-2} unstable, that is, to make the potential of v_{2i-2} positive.

(2) *Suppose $x = 4i - c > 1$, where $0 \leq c \leq 4i$ is an odd integer number.* It suffices to consider odd c since the degree of all nodes in the snake graph is odd; thus their potential is odd as well. Observe that $P(v_{2i-2}) \leq -4i + x = -c$, and in order to update v_{2i-2} to gray, at least $(c+1)/2$ nodes must be updated to gray in order to make v_{2i-2} unstable.

Note that in order to fulfill SG_i^g we must have updated v_{2i-2} as well as at least $(c+1)/2$ other nodes (including v_{2i-1}) to gray. Suppose that we did this; more specifically, suppose that between the end of phase $i-1$ and fulfilling SG_i^g we updated $y \geq 0$ nodes to white and $z \geq (c+1)/2 + y + 1$ nodes to gray. Then n_g increases by $z - y$ and n_w increases by $y - z$. Thus $n_w - n_g$ increases by $2(y - z) \leq 2y - 2((c+1)/2 + y + 1) = -c - 1 + 2$. Then

$$\begin{aligned} P(v_{2i-1}) &= -2i + n_w - n_g + 1 \\ &= -2i + 1 + (2i + 1 - x + 2(y - z)) \\ &< -4i + c + 2 - c - 1 + 2 = -4i + 3, \end{aligned}$$

and in order to fulfill SG_{i+1}^w and thus to end phase i , we require at least $2i - 1 = \Omega(i)$ updates from gray to white.

(3) *Suppose $x = 4i + c$ where $c \geq 0$ is an odd integer number.* Observe that $P(v_{2i-2}) \leq -4i + x = c$. Then, in order to fulfill SG_i^g , the six nodes controlled by said synchronization gadget must be updated from white to gray. Suppose that between the end of phase $i-1$ and fulfilling SG_i^g we updated $y \geq 0$ nodes to white and $z \geq y + 6$ nodes to gray. Then n_g increases by $z - y$ and n_w increases by $y - z$.

Thus $n_w - n_g$ increases by $2(y - z) \leq -12$. Then

$$\begin{aligned} P(v_{2i-1}) &= -2i + n_w - n_g + 1 \\ &= -2i + 1 + (2i + 1 - x + 2(y - z)) \\ &< -4i - c + 2 - 12 = -4i - c - 10, \end{aligned}$$

and in order to fulfill $\text{SG}_{i+1}(w)$ and thus to end phase i , we require at least $(4i + c + 10 + 1)/2 = 2i + 5 + (c+1)/2 = \Omega(i)$ updates from gray to white. \square

Thus, overall, an optimistic update sequence must be of length at least $\sum_{i=0}^{n-1} \Omega(i) = \Omega(n^2)$.

Theorem 2. *There exists a family of 3-opinion InfNets on n nodes for which every optimistic update sequence updates at least one node $\Omega(n)$ times, and is of length $\Omega(n^2)$.*

4 Computing Optimistic Sequences

In this section, we analyze the computational complexity of computing a optimistic update sequence. Formally, we consider the following decision problem.

k-OPTIMISTIC UPDATE SEQUENCE (*k*-OPT-SEQ)

Given an InfNet (G, \circ) , with k opinions, and a positive integer q , is there a stable update sequence σ such that the final number of black nodes is at least q , i.e., $|\{v \in V(G) \mid \circ[G, \sigma](v) = 0\}| \geq q$?

Auletta *et al.* [2019] have shown NP-hardness for a very related problem. The main difference is that in their diffusion model, an agent might even update its opinion when only a weak majority of neighbors agrees on a different opinion. Moreover, they assume that every agent updates its opinion only a constant number of times (and this constant is given as a parameter to their problem name). We remark that it seems possible to adapt their hardness proof for our model so that general NP-hardness of 3-OPT-SEQ is not surprising.

What might, however, be unexpected is that we can show intractability of 3-OPT-SEQ for highly restricted cases. In the next theorem, we show this problem remains intractable even if the InfNet has few black nodes, few initially unstable nodes, and short optimistic update sequences, while aiming only for an approximate solution.

Theorem 3. *3-OPT-SEQ is NP-hard and W[2]-hard parameterized by sequence length $|\sigma|$. Both holds even if the number of black nodes in the given InfNet is two and the number of initially unstable nodes is three.*

Proof. We show W[2]-hardness using a parameterized reduction from SET COVER, which is defined as follows: Given a finite set $U = \{u_1, \dots, u_t\}$, a family $\mathcal{S} = \{S_1, \dots, S_r\}$ of subsets of U , and a nonnegative integer h , SET COVER asks whether there is a *set cover* of size h , that is, whether there is a size- h set $\mathcal{S}' \subseteq \mathcal{S}$ such that $\bigcup_{S \in \mathcal{S}'} S = U$. Let (U, \mathcal{S}, h) be a SET COVER instance. We construct an equivalent 3-OPT-SEQ instance as follows (see Figure 3 for an illustration).

For each element $u_i \in U$, create an *element node* v_i with opinion $\circ(v_i) = g$. For each subset $S_j \in \mathcal{S}$, create one *subset node* w_j with opinion $\circ(w_j) = w$. For each $i \in [t], j \in [r]$ with $u_i \in S_j$ add the edge $\{v_i, w_j\}$.

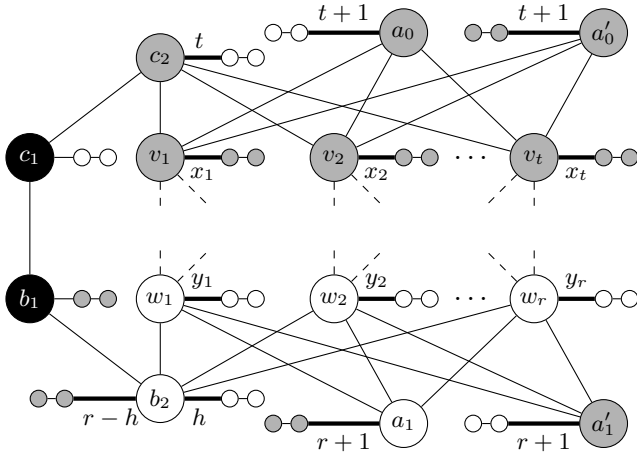


Figure 3: Illustration of the reduction from SET COVER to OPTIMISTIC UPDATE SEQUENCE to show $W[2]$ -hardness wrt. sequence length $|\sigma|$. We set $x_i = \#(u_i) - 2$ for $i \in [t]$ and $y_j = |S_j| - 1$ for $j \in [r]$. Thick lines represent multiple edges to multiple stable leaves and the edge label gives the multiplicity. Dashed lines represent possible edges (depending on the SET COVER input instance).

We create two *covering nodes* c_1 with $\circ(c_1) = b$ and c_2 with $\circ(c_2) = g$ as well as two *budget nodes* b_1 with $\circ(b_1) = b$ and b_2 with $\circ(b_2) = w$. Finally, we create four *activator nodes* a_0, a'_0, a_1 , and a'_1 with $\circ(a_0) = g$ and $\circ(a'_0) = g$, $\circ(a_1) = w$, and $\circ(a'_1) = g$.

We connect the covering node c_2 with every element node v_i , the budget nodes b_2 as well as activator nodes a_1 and a'_1 , with every subset node w_j , as well as c_1 with c_2 , b_1 with b_2 , and c_1 with b_1 . The activator nodes a_0 and a'_0 are both connected to every element node. The subset node w_j is connected to element node v_i if and only if $u_i \in S_j$.

Each node has a certain number of *stable leaf nodes*. Technically, a stable leaf node ℓ with opinion $\circ(\ell) = x$ attached to some node v in fact has a second neighbor ℓ' (besides neighbor v) with the same opinion $\circ(\ell') = x$. In effect, both ℓ and ℓ' are permanently stable.

To describe the leaf nodes, let $\#(u_i) := |\{S_j \in \mathcal{S} \mid u_i \in S_j\}|$ be the number of occurrences of element u_i . Node c_1 has one stable leaf with opinion w and node b_1 has one stable leaf with opinion g . Node c_2 has t stable leaves with opinion w , node a_0 has $t + 1$ stable leaves with opinion w , node a'_0 has $t + 1$ stable leaves with opinion g , node a_1 has $r + 1$ stable leaves with opinion g , and node a'_1 has $r + 1$ stable leaves with opinion w . Every element node v_i has $\#(u_i) - 2$ stable leaves with opinion g and every subset node w_j has $|S_j| - 1$ stable leaves with opinion w . Finally, node b_2 has $r - h$ stable leaves with opinion g and h leaves with opinion w .

Setting $q := 2$ finishes the construction which can clearly be performed in polynomial time.

Let us sketch the *idea and basic observations* of the correctness. Observe that no node can change its opinion to black. So we need to ensure that c_1 and b_1 remain black. For this we must ensure that c_2 and b_2 will also keep their opinions. However, node c_2 can only keep its opinion if all element nodes keep their opinion, which in turn is only pos-

sible if every element node has at least one neighboring subset node that changes its opinion to g (corresponding to a set cover). Moreover, if more than h subset nodes change their opinions, then b_2 must also change. Initially, only a_0, a_1 and a'_1 are unstable. At some point, a_0 must change its opinion to w so that, every element node, to keep opinion g , needs at least one neighboring subset node that changes its opinion to g . Nodes a_1 and a'_1 will swap their opinions. If we update a_1 before a'_1 , we can change the opinion of arbitrary subset nodes to g between the two updates, while we require for our graph to update at most h subset nodes such that every element node is neighboring one of the updated ones. Further, to keep the chosen subset nodes stable, at least one of the neighboring leaves has to update at any later point eventually as we require these subset nodes to be of opinion gray at the end of the update sequence, also requiring them to be stable. \square

Chen *et al.* [2006] show that, unless the Exponential Time Hypothesis (ETH) breaks, there is no $f(h)r^{o(h)}$ -time algorithm for SET COVER. Consequently we obtain:

Corollary 1. *There is no $f(|\sigma|)n^{o(|\sigma|)}$ -time algorithm for 3-OPT-SEQ unless ETH breaks.*

We can actually show that the reduction behind Theorem 3 still works when setting $q = 1$. This excludes approximation algorithms even if they may use $f(|\sigma|) \cdot n^{o(|\sigma|)}$ time.

Corollary 2. *Let (G, \circ) be an InfNet. Finding a stable update sequence σ with $|\{v \in V(G) \mid \circ[G, \sigma](v) = 0\}| \geq 1$ is NP-hard and $W[2]$ -hard with respect to $|\sigma|$.*

Finally, using a reduction from MAX 2-SAT we show that 3-OPT-SEQ remain hard even if every node has degree three.

Theorem 4. *3-OPT-SEQ is NP-hard, even if the maximum degree of G is at most three.*

5 Conclusion

This paper essentially shows that *few steps are not enough for optimistic update sequences*, first in the sense that computing (even approximate) optimistic update sequences is highly intractable, and second in the sense that in general networks optimistic sequences may have to be of quadratic length and that even in acyclic networks single nodes may have to update their opinions many times. The last result is of particular interest since it answers an open question implied by Auletta *et al.* [2019]: They formulated a polynomial-time algorithm for finding a spread-maximizing update sequence on acyclic three-opinion influence networks in which agents may change their opinion a constant number of times. Thus, our result shows that this algorithm cannot be used on acyclic networks in general, and whether one can find optimistic sequences on acyclic networks efficiently remains open.

Acknowledgements

Lilian Jacobs was supported by the DFG project ‘‘FPTinP’’ (NI 369/16).

References

- [Auletta *et al.*, 2019] Vincenzo Auletta, Diodato Ferraioli, Valeria Fionda, and Gianluigi Greco. Maximizing the spread of an opinion when tertium datur est. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS '19)*, pages 1207–1215. International Foundation for Autonomous Agents and Multiagent Systems, 2019.
- [Bredereck and Elkind, 2017] Robert Bredereck and Edith Elkind. Manipulating opinion diffusion in social networks. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence, (IJCAI '17)*, pages 894–900. ijcai.org, 2017.
- [Brill *et al.*, 2016] Markus Brill, Edith Elkind, Ulle Endriss, and Umberto Grandi. Pairwise diffusion of preference rankings in social networks. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI'16)*, pages 130–136, 2016.
- [Chen *et al.*, 2006] Jianer Chen, Xiuzhen Huang, Iyad A. Kanj, and Ge Xia. Strong computational lower bounds via parameterized complexity. *Journal of Computer and System Sciences*, 72(8):1346–1367, 2006.
- [Chierichetti *et al.*, 2013] Flavio Chierichetti, Jon M. Kleinberg, and Sigal Oren. On discrete preferences and coordination. In *Proceedings of the 14th ACM Conference on Electronic Commerce (EC'13)*, pages 233–250. ACM, 2013.
- [Clifford and Sudbury, 1973] Peter Clifford and Aidan Sudbury. A model for spatial conflict. *Biometrika*, 60(3):581–588, 1973.
- [Frischknecht *et al.*, 2013] Silvio Frischknecht, Barbara Keller, and Roger Wattenhofer. Convergence in (social) influence networks. In *Proceedings of the 27th International Symposium on Distributed Computing (DISC '13)*, pages 433–446. Springer, 2013.
- [Goles and Olivos, 1980] Eric Goles and J. Olivos. Periodic behaviour of generalized threshold functions. *Discrete Mathematics*, 30(2):187–189, 1980.
- [Goles and Tchuente, 1983] Eric Goles and Maurice Tchuente. Iterative behaviour of generalized majority functions. *Mathematical Social Sciences*, 4(3):197–204, 1983.
- [Grandi *et al.*, 2015] Umberto Grandi, Emiliano Lorini, and Laurent Perrussel. Propositional opinion diffusion. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems (AAMAS '15)*, pages 989–997. IFAAMAS, 2015.
- [Grandi, 2017] Umberto Grandi. Social choice and social networks. In U. Endriss, editor, *Trends in Computational Social Choice*. AI Access, 2017.
- [Kaaser *et al.*, 2016] Dominik Kaaser, Frederik Mallmann-Trenn, and Emanuele Natale. On the voting time of the deterministic majority process. In *Proceedings of the 41st International Symposium on Mathematical Foundations of Computer Science (MFCS '16)*, volume 58 of *LIPICs*, pages 55:1–55:15. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2016.
- [Kralovic, 2001] Rastislav Kralovic. On majority voting games in trees. In *Proceedings of the 28th International Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM '01)*, volume 2234 of *LNCS*, pages 282–291. Springer, 2001.
- [Slavkovik and Jamroga, 2016] Marija Slavkovik and Wojciech Jamroga. Iterative judgment aggregation. In *Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI'16)*, pages 1528–1536, 2016.