

On the Complexity of Opinion Consensus under Majority Dynamics^{*}

Vincenzo Auletta¹, Diodato Ferraioli¹, and Gianluigi Greco²

¹ University of Salerno
{auletta,dferraioli}@unisa.it

² University of Calabria
gianluigi.greco@unical.it

Abstract. Opinion diffusion is studied on social graphs where agents hold opinions and where social pressure leads them to conform to the opinion manifested by the majority of their neighbors. Within this setting, questions related to which extent a minority/majority can spread the opinion it supports to the other agents are considered. It is shown that if there are only two available opinions, no matter of the underlying social graph $G = (N, E)$, there is always a group formed by a half of the agents that can annihilate the opposite opinion. A polynomial-time algorithm to compute a group of agents enjoying these properties is also devised and analyzed. The result marks the boundary of tractability, since the influence power of minorities is shown to depend on certain features of the underlying graphs, which are NP-hard to be identified. Finally, for more than two opinions we show that even the simpler problem of deciding whether there exists a sequence of updates leading to consensus is NP-hard.

Keywords: opinion diffusion · stability · consensus

1 Introduction

Consider the following prototypical scenario. The members of a department are organizing a social dinner, and they have to decide whether to go to a restaurant or to a pizzeria. Initially, each of them holds an opinion on her ideal choice. At a certain point, they will exchange their viewpoints and each of them will be affected by a social pressure leading to adapt her opinion to the one manifested by the majority of her friends. So, we ask: Would be they capable to reach a *consensus* for some/all profiles of their initial opinions? Can a *minority* have enough social “power” to influence all other agents? Is it any easier for a *majority* to guide convergence towards consensus? Can they reach an equilibrium different from a consensus?

Our goal is to analyze the above kinds of questions under the lens of algorithm design and computational complexity, by focusing on a setting where social

^{*} Copyright © 2019 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

relationships are encoded as the edges of a social graph $G = (N, E)$ whose nodes correspond to the agents. In particular, by starting from an initial configuration where agents hold some innate opinions, we consider a model of opinion diffusion over the underlying graph where each agent is stable if, and only if, her current opinion agrees with the opinion held by a (non-strict) majority of her neighbors. Hence, at any time step of the dynamics, agents that are not stable can change their opinion *asynchronously*, thereby leading to a non-deterministic evolution where the final configuration may depend on the specific order in which updates have been performed.

In fact, the study of opinion diffusion constitutes nowadays an active area of research in the computer science literature (e.g., [15, 1, 3]). However, few results are known for the following fundamental CONSENSUS problem: Given a rational number α with $0 \leq \alpha \leq 1$, is there any set $S \subseteq N$ of agents with $|S| \leq \lceil \alpha|N| \rceil$ that is able to influence all the other agents? It is known to be NP-hard [9]: However, we missed so far a finer-grained analysis of the complexity of CONSENSUS which charts its frontier of tractability w.r.t. the ranges of the possible values of α , the classes of social graphs being considered, the number of available alternatives.

In the paper¹, we fill this gap and we provide the following contributions. First, we analyze the CONSENSUS problem on arbitrary social graphs, but focusing on scenarios with only two available opinions, and for which the fraction α of the agents that already agree on the opinion to be propagated is such that $\alpha \geq \frac{1}{2}$. We show that, in this case, CONSENSUS is tractable and in particular a majority of $\lceil |N|/2 \rceil$ agents always exists (and can be efficiently computed) which is capable of annihilating the opposite opinion. Second, we show that the value $\alpha = \frac{1}{2}$ defines a sharp boundary for the CONSENSUS problem when there are only two available opinions. Indeed, we show that a minority (i.e., $\alpha < \frac{1}{2}$) that can spread its opinion to all the agents exists only in certain graphs. In fact, a result of this kind holds for the problem of assessing the existence of a minority that can become a majority [2, 4, 5]. However, in that case the graphs enjoying the desired property admit a computationally simple characterization, while in our case a characterization of these graphs is NP-hard. Third, we consider the case that there are more than two available opinions. We show that the scenario radically changes with respect to the binary case. Indeed, we show that it is NP-hard even to verify if there is a sequence of updates leading from a given set of initial opinions to consensus.

2 The Model

Let $G = (N, E)$ be a *social* graph, that is, an undirected connected graph encoding the interactions of a set N of agents. Moreover, a set O of opinions are

¹ Some of the results appeared in the 27th International Joint Conference on Artificial Intelligence [8] (Distinguished Paper) and in the 18th International Conference on Autonomous Agents and MultiAgent Systems [7].




Fig. 1. Greedy dynamics do not maximize opinion when there are three available opinions.

available to the agents. Accordingly, we model a *configuration* for G as a function $c : N \rightarrow O$; its intended meaning is that agent $x \in N$ with $c(x) = o$ holds opinion o . For any set $S \subseteq N$ of agents, we define $S_{o/c}$ (shortly S_o , if the configuration c is clearly understood) as the set of all agents in S with opinion o . For each agent $x \in N$, the set $\{y \mid \{x, y\} \in E\}$ of her neighbors is denoted by $\delta(x)$. Agent $x \in N$ is *stable* in c if her opinion agrees with the opinion held by a (non-strict) majority of her neighbors. A configuration c is *stable* if all agents in N are stable. We consider an *asynchronous* and *non-deterministic* model of opinion diffusion where, at each time instant, some arbitrarily chosen agents that are not stable change their opinion. More formally, a *dynamics* for G is modeled throughout the paper as a sequence of configurations c_0, \dots, c_k such that c_{h+1} , for each $h \in \{0, \dots, k-1\}$, is obtained from c_h by changing the opinions of a non-empty subset of agents that are not stable in c_h .

A kind of dynamics that will play a prominent role in this paper is the *greedy dynamics* for a given opinion $op \in O$. It works as follows: As long as there are agents that are not stable and for which op is a majority in the neighborhood, change their opinion to op ; otherwise change the opinion of some other non-stable nodes. Greedy dynamics turn out to be crucial in the case that only two opinions are available, i.e. $O = \{0, 1\}$. Indeed, in this case the following property holds [9]: Given a configuration c for a graph G , the greedy dynamics converges to a stable configuration $\max_{op}(c)$ and, for each dynamics $c = c_0, \dots, c_k$ such that c_k is a stable configuration, it holds that $|N_{op/\max_{op}(c)}| \geq |N_{op/c_k}|$.

In the light of the above result, when we consider that only two opinions are available, we can focus our analysis, w.l.o.g., on greedy dynamics. Unfortunately, this result does not extend to the case in which there are at least three available opinions. Consider indeed that $O = \{white, black, gray\}$ and the graph is as in Figure 1.(a). Let c be the initial configuration where each agent holds the opinion corresponding to the coloring of the nodes in the figure. Consider the (non-greedy) dynamics, illustrated in Figure 1.(b): First, c adopts opinion *gray*; then, all agents adopt opinion *white* in the following order: $b, d, c, h_1, \dots, h_{n-1}, h_n$. Eventually, an opinion profile is reached where all agents hold opinion *white*.

On the other hand, observe that a greedy dynamics must start with agent d changing her opinion to *white*, since in the profile \mathbf{c} this is the only agent interested in adopting opinion *white*. Let \mathbf{c}' be the configuration obtained after this change, which is graphically illustrated in Figure 1.(c). From this configuration, no other agent is interested in changing her opinion to *white*. Rather, agent a will eventually adopt opinion *black* (possibly after that some of the agents in $\{h_1, \dots, h_n\}$ changed her opinion to *black* or *gray*). Hence, we have that the maximum number of whites in a stable configuration reachable from \mathbf{c} via the dynamics $\mathbf{c} \rightsquigarrow \max_{\text{white}}(\mathbf{c})$ is only 3.

3 Our Results

Define $\forall\text{op}$ as the configuration where all agents hold opinion op , and consider the CONSENSUS problem: Given an undirected graph $G = (N, E)$, a rational number α such that $0 < \alpha < 1$, compute a configuration \mathbf{c} for G such that (i) $|N_{\text{op}/\mathbf{c}}| \leq \lceil \alpha |N| \rceil$, (ii) $\max_{\text{op}}(\mathbf{c}) = \forall\text{op}$, or check that there is no configuration enjoying (i), (ii). We start by considering that $O = \{0, 1\}$ and we study the complexity of the CONSENSUS problem by assuming, w.l.o.g., that $\text{op} = 1$.

From Majority to Consensus. We start our study by showing that, for each undirected graph G , whenever the fraction α covers at least a majority of the agents, in particular, even if $\alpha = \frac{1}{2}$, a configuration \mathbf{c} with $|N_{1/\mathbf{c}}| \leq \lceil \alpha |N| \rceil$ and $\max_1(\mathbf{c}) = \forall 1$ always exist and can be computed in polynomial-time.

To establish the results, we explore the space of the *binary partitions* \mathcal{P} of the agent set N , that is, of the pairs $\mathcal{P} = (A, B)$ where A and B are non-empty sets such that $A \cup B = N$. Given a partition $\mathcal{P} = (A, B)$ of N , we write $X \in \mathcal{P}$ to denote that a set $X \subseteq N$ belongs to $\{A, B\}$; moreover, we define $\bar{X} = N \setminus X$. For any agent $x \in N$, let \mathcal{P}_x (resp., $\bar{\mathcal{P}}_x$) denote the set of \mathcal{P} to which x belongs (resp., does not belong), and let us define the *utility* of x in \mathcal{P} as the value $u(x, \mathcal{P}) = |\delta(x) \cap \mathcal{P}_x| - |\delta(x) \cap \bar{\mathcal{P}}_x|$. Moreover, for each “side” $X \in \mathcal{P}$, we denote by $G[X]$ the subgraph of G induced by X , and we define $\text{Zc}(X, \mathcal{P})$ as the set of all connected components of $G[X]$ such that $u(y, \mathcal{P}) = 0$ for each component $C \in \text{Zc}(X, \mathcal{P})$ and each agent y of C . Elements in $\text{Zc}(X, \mathcal{P})$ are called *zero components*. Finally, for any two disjoint sets of agents A' and B' , not necessarily forming a partition, let $E(A', B')$ be the set of edges $e \in E$ such that e has one endpoint in A' and the other in B' .

A partition \mathcal{P} is called *nice* if it has a nice side $X \in \mathcal{P}$ such that: (i) $u(x, \mathcal{P}) + u(y, \mathcal{P}) \leq -2|E(\{x\}, \{y\})|$, for each pair of agents $x \in X$ and $y \in \bar{X}$; (ii) either there is an agent $x^* \in X$ with $u(x^*, \mathcal{P}) > 0$, or $u(x, \mathcal{P}) \leq 0$ holds for each $x \in \bar{X}$ and $\text{Zc}(\bar{X}, \mathcal{P}) = \emptyset$. We have that consensus can be reached from a configuration that can be easily computed when a nice partition is given at hand.

Lemma 1. *Let \mathcal{P} be a nice partition, X its nice side, and $\bar{\mathbf{c}}(\mathcal{P})$ be the configuration that assigns opinion 1 only to agents in X . Then, $\max_1(\bar{\mathbf{c}}(\mathcal{P})) = \forall 1$.*

The question is now whether we can efficiently single out a nice partition. To answer positively this question, we individuate the obstructions to a partition

P to be nice. These are pairs of agents $\{x, y\}$, which we call *critical*, lying on distinct sides of \mathcal{P} , and such that $\mathcal{P}' = (\overline{\mathcal{P}}_x \cup \{x\} \setminus \{y\}, \mathcal{P}_x \cup \{y\} \setminus \{x\})$, either has less edges across its two sides than P , or has less zero components than P , or the maximum distance between two zero components, each on a different side, is smaller than in P . The following result links critical pairs and nice partitions.

Lemma 2. *Assume that no critical pair exists in a partition \mathcal{P} . Then, \mathcal{P} is nice.*

We can now state the main results of this section. Indeed, it turns out that a nice partition can be computed by iteratively swapping agents belonging to some critical pairs, as long as one exists.

Theorem 1. *Given any graph $G = (N, E)$, a configuration c for G can be computed in polynomial time such that $|N_{1/c}| \leq \lceil \frac{1}{2}|N| \rceil$, $\max_1(c) = \forall 1$.*

From Minority to Consensus. So far, we have analyzed the CONSENSUS problem by focusing on instances where the opinion 1 to be propagated to all the agents is initially already supported by some majority. In the following, we complete the picture of this analysis by considering the CONSENSUS problem restricted first to instances such that $\alpha < \frac{1}{2}$. Our results will show that Theorem 1 essentially charts the frontier of tractability for the CONSENSUS problem.

Inspired by similar results in earlier literature [16, 11], we show that CONSENSUS is NP-hard for $\alpha < \frac{1}{2}$ by exhibiting a reduction from the well-known VERTEX COVER problem [14].

Theorem 2. *On the class of instances where α is such that $0 < \alpha < \frac{1}{2}$, CONSENSUS is NP-hard.*

Multiple Opinions. Consider now the case that $|O| > 2$. As discussed above, we cannot assume that the dynamics leading to consensus, if there is any, is the greedy dynamics. Hence, before that the CONSENSUS problem would be addressable, we need to consider the simpler COMPLETE-SPREAD problem: to decide if there is a sequence of updates leading from the given configuration c to a consensus. We can show that even this problem is intractable, through a reduction from the NP-hard problem ONE-IN-THREE POSITIVE 3-SAT [14], that consists in deciding whether there is a truth assignment σ such that, for each clause c_j , precisely one variable in c_j evaluates to true.

Theorem 3. *If $|O| > 2$, then the COMPLETE-SPREAD problem is NP-complete.*

4 Conclusion

We addressed a number of questions related to whether consensus can be achieved in settings where opinions of the agents are affected by social influence.

Our results pave the way for further investigations. For instance, it would be interesting to analyze the extent at which consensus can be reached and kept if the social graph is dynamic, in order to model the evolving relationships among agents. Whereas some results of this kind are known for different dynamics, see, e.g., [13, 17, 10, 12, 6], the behavior of the majority dynamics in this setting is nowadays still obscure.

References

1. Acar, E., Greco, G., Manna, M.: Group reasoning in social environments. In: Proc. of AAMAS'17. pp. 1296–1304 (2017)
2. Auletta, V., Caragiannis, I., Ferraioli, D., Galdi, C., Persiano, G.: Minority becomes majority in social networks. In: Proc. of WINE'15. pp. 74–88 (2015)
3. Auletta, V., Caragiannis, I., Ferraioli, D., Galdi, C., Persiano, G.: Generalized discrete preference games. In: Proc. of IJCAI'16. pp. 53–59 (2016)
4. Auletta, V., Caragiannis, I., Ferraioli, D., Galdi, C., Persiano, G.: Information retention in heterogeneous majority dynamics. In: Proc. of WINE'17. pp. 30–43 (2017)
5. Auletta, V., Caragiannis, I., Ferraioli, D., Galdi, C., Persiano, G.: Robustness in discrete preference games. In: Proc. of AAMAS '17. pp. 1314–1322 (2017)
6. Auletta, V., Fanelli, A., Ferraioli, D.: Consensus in opinion formation processes in fully evolving environments. In: Proc. of AAAI'19 (2019)
7. Auletta, V., Ferraioli, D., Fionda, V., Greco, G.: Maximizing the spread of an opinion when tertium datur est. In: Proc. of AAMAS'19 (2019)
8. Auletta, V., Ferraioli, D., Greco, G.: Reasoning about consensus when opinions diffuse through majority dynamics. In: Proc. of IJCAI'18. pp. 49–55 (2018)
9. Bredereck, R., Elkind, E.: Manipulating opinion diffusion in social networks. In: Proc. of IJCAI'17. pp. 894–900 (2017)
10. Carvalho, A., Larson, K.: A consensual linear opinion pool. In: IJCAI. pp. 2518–2524 (2013)
11. Chen, N.: On the approximability of influence in social networks. In: Proc. of SODA'08. pp. 1029–1037 (2008)
12. Ferraioli, D., Ventre, C.: Social pressure in opinion games. In: Proc. of IJCAI'17. pp. 3661–3667 (2017)
13. Fortunato, S.: On the consensus threshold for the opinion dynamics of krause–hegselmann. *International Journal of Modern Physics C* **16**(02), 259–270 (2005)
14. Garey, M.R., Johnson, D.S.: *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co. (1979)
15. Grandi, U., Lorini, E., Novaro, A., Perrussel, L.: Strategic disclosure of opinions on a social network. In: Proc. of AAMAS'17. pp. 1196–1204 (2017)
16. Kempe, D., Kleinberg, J., Tardos, E.: Influential nodes in a diffusion model for social networks. In: Proc. of ICALP'05. pp. 1127–1138 (2005)
17. Lorenz, J., Urbig, D.: About the power to enforce and prevent consensus by manipulating communication rules. *Advances in Complex Systems* **10**(02), 251–269 (2007)

Acknowledgments. Vincenzo Auletta and Diodato Ferraioli were partially supported by “GNCS-INdAM” and by the Italian MIUR PRIN 2017 Project ALGADIMAR “Algorithms, Games, and Digital Markets”. Gianluigi Greco was partially supported by Regione Calabria under POR project “Explora Process”.