Preprints of the 21st IFAC World Congress (Virtual)
Berlin, Germany, July 12-17, 2020

A control approach to address ethical issues on social (robotic) networks ⋆

Feres A. Salem * **Ubirajara F. Moreno** *
Françoise Lamnabhi-Lagarrigue **

* *Federal University of Santa Catarina, Automation and Systems Department, Florianopolis, Brazil (e-mails: ubirajara.f.moreno@ufsc.br, feres.salem@posgrad.ufsc.br).*
** *L2S, UMR CNRS 8506, CentraleSupelec, University Paris-Sud and Paris-Saclay, 91192, Gif-sur-Yvette (e-mail: lamnabhi@l2s.centralesupelec.fr)*

Abstract: From the first Internet-based social networking applications designed to get people in contact and make friends to social networks made up of over 2 billion users, the combination of communication networks, portable devices, and AI has changed the way people interact and make decisions. The extent of this influence could be observed not only in marketing and social behavior but also in referendums and elections, leading to distortion of democratic manifestations and representations. Considering this, the consequences of the misuse of Internet-based social networks could have a substantial impact on society, and it is important to define ethical guidelines and policies for developers, rulers, operators, and social actors. Considering this aim, the objective of this paper is to show that an approach based on Systems & Control could be effective in evaluating the impact of such policies in order to meet ethical issues. The starting point from this analysis are the models of information spreading in dynamic social networks, and these models are adapted and updated to encompass the complex behavior from users, as well as some regulatory policies. The analyses were based on simulations of these models on small and large scale networks.

Keywords: Systems & Control; Social Networks, Ethics.

1. INTRODUCTION

The ability to establish complex social networks is a distinct skill from humanity; for many researchers, it is a milestone for human evolution. From small-group organization to the foundation of empires with complex lines of communication, from the diffusion of information in a tribe to global communication, the evolution of social networks could not be separated from technological transformations.

The study of the human social interaction as a field could be situated in the nineteenth century, and the modeling from social networks by applying graph theory and algebraic models finds its roots in the seminal works on Sociometry from Jacob Moreno (Freeman, 2004; Proskurnikov and Tempo, 2017). In the early twentieth century, researches in the influence of media on social behavior began. Firstly, the effect of the press was analyzed; later, with radio and television development, many relevant studies were conducted, in the 1950s and 1960s, analyzing the effect of media on social behavior and the dissemination of information.

More recently, with the development of applications on the web 2.0, combined with portable devices and AI algorithms, the first Online Social Networks (**OSNs**) applications designed to get people in contact and make friends, turned into social networks made up of over 2 billion users. This combination of communication networks, portable devices, and AI has irreversibly changed the way people interact and make decisions (Van Dijck, 2013).

This impact from **OSNs** brought the attention of many researchers, especially in the community of Systems & Control, due to possibilities of applying well-known tools of representation, such as graph theory, and the utilization of dynamical systems to model the diffusion of information and the achievement of consensus. In some measure, these problems are very similar to one studied in cooperative mobile robotics and other applications (Proskurnikov and Tempo, 2017; Bullo, 2019).

Besides the developments in control and modeling of social networks, many philosophers, sociologists, and political scientists are addressing these new phenomena that arouse in social behavior by the massive utilization of social networks (Han, 2017). Not only to understand the newly emerging collective behaviors but to foresee possible consequences and ethical dilemmas in the usage and manipulation of social networks (Vallor, 2016; Howard et al., 2018).

Despite the vast possibilities and innovations generated by **OSNs**, many of them being an essential tool for our everyday life; like any other disruptive technology, society must be aware and discuss potential undesirable side ef-

⋆ This work was supported by CAPES-PRINT, Brazil.

fects. In the last years, many disturbing effects of the abuse or misuse from **OSNs** have been reported in media. Manipulation of public opinion in elections by massive robotic nodes, mob attacking innocent people, non-precedent diffusion of fake news, bullying among teenagers, profiling of users overruling privacy, and many others.

The first works considering ethical questions on **OSNs** were more related to the effects of the overexposure and privacy breaches. However, with the use of AI to make profiles from users on social networks, more attention has been given to the possible manipulations of will, desires, and free choices. More recently, the use of apps to share messages, as a tool for disseminating fake news and influencing elections, was at the center of the debate of the misuse of **OSNs**.

Following previous works on the modeling of opinion diffusion in social networks and trying to consider ethical questions, the subject of this paper is to analyze the effect of information dissemination. It focuses on the use of robotic nodes to disseminate information and manipulate opinion consensus by applying some models presented in the literature. More specifically, two problems are addressed: worsening the humans self-appraisal and spreading malicious information. The objective is to show that a Systems & Control approach could be useful to implement and analyze policies that encompass ethical commitments.

The paper is organized as follows: in section 2, some concepts on social networks and ethics are presented with some recent reports on the use of robotic nodes in **OSNs**. In section 3, some models of opinion diffusion are presented; the application of these models, to analyze policies of utilization in simulation scenarios, is carried out in section 4. In section 5, concluding remarks and new perspectives of study are introduced.

2. ETHICAL ISSUES IN ONLINE SOCIAL NETWORKS

2.1 Social Networks

Since the beginning of the modern studies of sociology, in the nineteenth century, society was viewed as an interconnection of social actors. In the 1930s, a significant contribution was the work from Jacob Levy Moreno, whose theory of society focused on the networks of interpersonal relations that join individuals (Freeman, 2004). From the earliest works, ethical issues have played a central role in the development of theories on the behavior of society, whether in the assumption of the behaviors of specific social groups or in proposing tools to control and influence public opinion (Bernays, 2005; Lippmann, 2017).

In addition to the evolution of studies in social networks, the development of new mass media has impacted the way individuals interact socially, increasing the reach of the media to a global scale. These new *media* also gave rise to new tools to interfere and influence society. In addition to the increase in the range and speed of information dissemination, new technologies, such as press, radio, television, have defined new social relationships. According to influential studies developed in the 1960s, the medium defines new relationships, being itself a message, as Marshall McLuhan stated: "*the medium is the message*" (McLuhan, 1994).

With the evolution of the media and its uses, regulations have been developed to prevent distortion and abuse, in marketing, public opinion manipulation, and respect for privacy. In the case of influence on electoral processes, various constraints and practices have been regulated around the world (EODS, 2016). The principles underlying these initiatives are the guarantee of the right of expression of the various social players involved in the processes, avoiding asymmetries, curbing the abuse of economic power, preventing the dissemination of false information, and avoiding the manipulation of public opinion.

2.2 Online Social Networks

According to some references, the first internet-based social networks were launched at the end of the '90s, allowing people to create profiles and list friends and classmates (Boyd and Ellison, 2007). After the proliferating of many platforms, the instant messaging services increase the usage of **OSNs**, but the great step was given with the applications that enable sharing of content, turning these platforms in large corporations, and the applications into techno-cultural devices (Van Dijck, 2013).

With the use of **OSNs** new manifestations emerged, such as the decentralization of dissemination of cultural production, as well as news and opinions through blogs, and the use of social networking platforms, among others. These manifestations soon attracted the attention of scholars, trying to analyze possible ethical and social impacts, especially from social networking platforms. Many philosophers of technology, such as Albert Borgmann and Hubert Dreyfus, pointed out, very earlier, some possible risks to human relations from the overexposure in social networking (Vallor, 2016). However, due to the continued development of new applications, the widespread popularization of mobile phones, and the use of AI, the previous analyses were overwhelmed by new mass behaviors on a global scale.

Nowadays, by using shared information provided by users, applications are designed to support many activities, such as finding a job, renting a house, or even connecting citizens to run a neighborhood. One of the most critical aspects of **OSNs** is the free exchange of information by users in small or large groups. These new uses have changed many aspects of everyday life due to new social and political phenomena. According to many studies, the utilization of **OSN** changed the political space. The democratization of access to media, information production or consumption, and the increase of the power of mobilization were some of the major benefits.

However, the constitution of large groups to exchange ideas and information, gave rise to a phenomenon of polarization and a "*tendency to reward virality over veracity may harm information quality and democratic discourse*" (Neudert and Marchal, 2019; Han, 2017).

Due to lack of legislation, the distributed structure, the tendency to virality and the characteristics of fast and deep propagation of *fake news*, due the novelty and emotional content (Vosoughi et al., 2018), **OSNs** have been massively used, across the world, not only as a democratic tool for the dissemination and debate of ideas and proposals, but also for manipulation purposes (Tardáguila et al., 2018; Bradshaw and Howard, 2017).

This manipulation was based on the diffusion of personalized messages grounded on tools for profiling users but also with the use of robotic nodes, *social media bots*, in messaging applications (Howard et al., 2018; Wang et al., 2018).

The ethical issues of these misuse of **OSNs** are related not only to privacy and security but with trust and social manipulation, undermining democratic processes, urging the development of policies and regulation tools (Turculeţ, 2014; Howard et al., 2018). Another important aspect is the impact of the network misinformation in health decisions, by the spreading of *fake science*, affecting the public health system (Shao et al., 2018).

The proposal of policies to prevent the misuse and manipulation of large groups in social platforms, together with the studies about social behavior, constitutes a major concern for academia, legislators, and social actors.

The study of policies to avoid the misuse of private information has received considerable attention, and the adoption of new legislation has been debated and approved. On the other hand, the detection of fake news web-based services and algorithms to fact-checking are becoming essential tools (Zhang and Ghorbani, 2019). Also, some measures could be taken to limit the number of members of groups in instant messaging applications and the number of times that users can forward messages.

In this article, the main concern is to evaluate the use of *social media bots* and their influence on two crucial aspects: the self-appraisal of nodes representing human users and the consensus of network nodes on an issue. In this last aspect, the use of reliable information is compared to the use of robotic nodes that disseminate information without certification. Thus, it is intended to evaluate the implementation of measures to reduce *fake news* spreading, as well as to evaluate their impact on the presence of malicious robotic nodes.

The method chosen for analyzing these influences and the effect of possible countermeasures is the representation of social networks by dynamic models of information diffusion combined with gossip-based consensus algorithms (Proskurnikov and Tempo, 2017; Jia et al., 2015; Salem et al., 2019). The objective is to verify to which extent some policies could be implemented in order to tackle ethical issues in **OSNs**.

3. MODELING OF SOCIAL NETWORKS

Social networks are important case studies in the theory of complex networks and multi-agent systems because they differ from many natural and human-made complex networks in the sense that they do not present a cooperative behavior. The **OSNs** also work as tools for mapping the wiring diagram behind our social system, which is a necessary step to describe its complex behaviors (Barabási et al., 2016). The mathematical models for representing this kind of networks need to be simple enough to be examined, but able to capture the complex behavior of real social groups (Parsegov et al., 2017).

Proskurnikov and Tempo (2017, 2018) have done incredible work by highlighting the most classical models of social dynamics and also bringing the more recent models developed. Much of the theoretical knowledge needed to understand these models are compiled by Bullo (2019) in his book.

The results obtained in the present work are related to the process of information exchange through a social network and how it affects the humans involved in these processes. For this purpose, it is exposed here one model to describe how the flow of a finite set of information about a particular issue occurs in a network.

Also, it is shown a model for the evolution of self-appraisal in individuals based on the theory established by Friedkin et al. (2016) that describes how appraisals from others influence it. That is to say that self-esteem, self-efficacy, and self-confidence are a ubiquitous social construction.

3.1 Gossip Model

The model presented here is grounded on the asynchronous gossip-based consensus algorithm (Boyd et al., 2006), which is widely explored in the literature. In which is considered a pair-wise communication, such that, at each iteration k, an active agent $i(k)$ is randomly selected following a uniform distribution.

Then, the active agent has a probability p_{ij} of interacting with agent j. Those probabilities of interaction between agents are arranged in an $n \times n$ row-stochastic matrix,

$$P = [p_{ij}]$$

where n is the number of agents in the network.

So, whenever an interaction occurs, the active agent will update its information as

$$y_i(k+1) = (1-\nu_i)y_i(k) + \nu_i y_j(k) \ , \quad i = i(k) \tag{1}$$

where ν_i is a constant representing the trust of agent i in its neighbors.

The other agents (including j) persist their information to the next iteration, i.e.,

$$y_l(k+1) = y_l(k) \ , \quad \forall l \neq i(k) \tag{2}$$

In the synchronous version of this algorithm, no active agent is selected. Instead, at each step k, all the agents update their information following the probability distribution imposed by the matrix P.

3.2 DeGroot Model

This model describes the evolution of the opinions in a group of individuals trying to reach consensus and was first presented by DeGroot (1974). It assumes that individuals update their opinions as convex combinations of their own and those displayed by others (Jia et al., 2015). So, the opinions update as

$$z(k+1) = Wz(k) \tag{3}$$

where $z(k) \in \mathbb{R}^n$ is a vector holding the opinions of the n agents; W is a row-stochastic matrix that describes the influence weights.

So, each edge, $j \xrightarrow{w_{ij}} i$, represents the weight that individual i accords to the opinion from individual j or how much influence j has over i.

3.3 Handling Discrete Information

The two models mentioned previously are widely used in the literature to describe information-exchange processes. However, both of them consider the information as real numbers, which is not the case for many situations where the group needs to decide over a finite set of possibilities.

The first step to define the proposed model is to establish the connection between gossip and DeGroot models. This link is made by defining the matrix of probabilities of interaction equal to the matrix of influence weights, i.e.,

$$P \triangleq W$$

That connection between the models is the same as saying that the probability of agent i to interact with agent j is proportional to the influence that j has over i. Since the matrix of probabilities now represents the trust of the agents to one another, a new assumption over the gossip model can be made.

In (1), consider that $\nu_i = 1$, $\forall i$. Then, it simplifies to

$$x_i(k+1) = x_j(k) \tag{4}$$

where $x_i(k)$ is the information from agent i. This way, whenever an interaction occurs, the active agent does not compute a weighted average of its information and the one displayed by its neighbor.

Instead, the active agent will replace its opinion entirely, which guarantees that the information will be only exchanged, and no new information is created within the network. This way, our model is said to handle discrete information.

A fundamental aspect of this model is that the connections and its weights are the only factors influencing the information-exchange process. In other words, at this point, the content of the information is irrelevant.

3.4 Taking Reliability into Account

In order to add some notion of reliability to the information, it will be considered as a 2-tuple

$$X_i(k) = \langle m, \tau_m \rangle$$

where m is the content of the information from agent i; and τ_m is a numeric measure of the reliability of this information.

Next, a sample dependent matrix is defined to couple the influence of one agent over others with the reliability of the information held by the agent, such that,

$$\mathrm{A}(k) = [\alpha_{ij}(k)], \quad \alpha_{ij}(k) = w_{ij} X_j^{\langle 2 \rangle}(k) \ , \quad \forall i, j$$

where w_{ij} are the influence weights described before and $X_j^{\langle 2 \rangle}(k)$ is the second element of the tuple, i.e., the reliability coefficient of the information within agent j at the sample k.

Since the reliability coefficients, τ_m, multiply the influence weights, ranges can be defined for them, such that

$$\tau_m \in \mathbb{R}_+ \rightarrow \begin{cases} 0 \le \tau_m < 1 & \text{, information is not reliable} \\ \tau_m = 1 & \text{, the reliability is irrelevant} \\ \tau_m > 1 & \text{, information is reliable} \end{cases}$$

The matrix $\mathrm{A}(k)$ is not row-stochastic and, in order to be used as the matrix of probabilities of interaction, it needs to be normalized as follows

$$P = [p_{ij}], \text{ where } p_{ij} \triangleq \frac{\alpha_{ij}}{\sum_{j=0}^{n} \alpha_{ij}}$$

3.5 Evolution of Self-Appraisal

Jia et al. (2015) state that in an influence network with $n \ge 2$ agents, the self-weights (the elements of the diagonal from matrix W, i.e., $\chi(s) = w_{ii}(s) \ \ \forall i$) are updated like $\chi(s+1) = F(\chi(s))$, where $F : \Delta_n \rightarrow \Delta_n$ (Δ_n being the n-simplex) is a continuous map defined by

$$F(\chi) = \begin{cases} \mathbf{e}_i & , x = \mathbf{e}_i \ \forall i \\ \left(\dfrac{\sigma_1}{1-\chi_1}, \dots, \dfrac{\sigma_n}{1-\chi_n} \right)^T \Big/ \displaystyle\sum_{i=1}^{n} \dfrac{\sigma_i}{1-\chi_i} & , \text{otherwise} \end{cases} \tag{5}$$

where $\mathbf{e}_i$ is the ith basis vector (all elements equal to 0 except for the ith that is equal to 1), $\sigma^\top = [\, \sigma_1 \dots \sigma_n \,]$ is the vector of centrality scores (Bonacich, 1972) for each agent and it is defined as the dominant left eigenvector of the relative interaction matrix $C = [c_{ij}] \in \mathbb{R}^{n \times n}$ that is row-stochastic, zero-diagonal, and irreducible.

4. SIMULATIONS AND DISCUSSIONS

In this section, two cases are analyzed. Namely, the influence of *bots* in the human self-appraisal and the effect of bots compared with reliable information in the opinion evolution. It was considered that the bots infiltrate in social networks with malicious objectives; for this reason, they are called *malicious nodes* or *malicious agents* in some parts of the text.

4.1 Effect of bots in the human self-appraisal

The first case to be analyzed is the effect on the humans' self-appraisal caused by the presence of malicious agents infiltrated in a group that is debating over a series of issues. It is considered an influence network with ten nodes, being the nodes labeled from 1 to 8 set as humans and the nodes 9 and 10 as robots.

Here, two differences in behavior between humans and robots are considered. Robots do not change their self-weights from one issue to another as humans do. Also, robots are malicious in the sense that they will care less about opinions coming from humans than those from other robotic nodes.

The humans' self-weights are updated from one issue to another, according to the model in (5). The issue's end was considered as the moment when the group had reached consensus over a finite set of possible opinions. More details about how this simulation was carried can be found in the paper from Salem et al. (2019).

The graphs in Figure 1 illustrate the evolution of the weights in that influence network. In these, the size of the

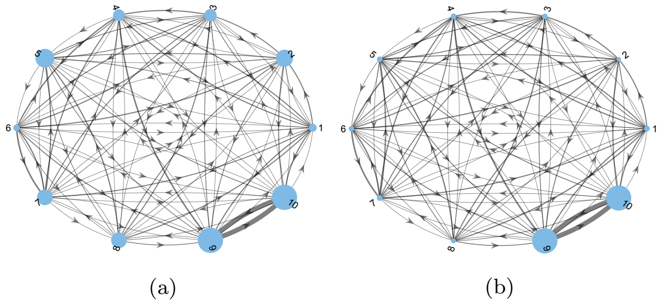

Figure 1. Evolution of the influence network with two bots across the sequence of issues: (a) initial conditions, at the first issue; (b) static state, reached at the fourth issue

nodes represents their current self-weights, i.e., how much they value their own opinion. Also, the thickness of the edges represents the interpersonal weights; for instance, a very thick edge from node i to node j means that opinions from agent i is very important to agent j.

From Figure 1a to 1b, it is noticeable that the humans' self-appraisal decayed greatly after four issues. This effect is due to humans' perception that their opinion does not affect much the opinion of the group. Also, worsening humans' self-appraisal makes them more easily influenced, producing a final opinion of the group that is hugely affected by the opinion of nodes with greater influence and self-appraisal. In the presented scenario, robots would influence more the opinion of the group than humans.

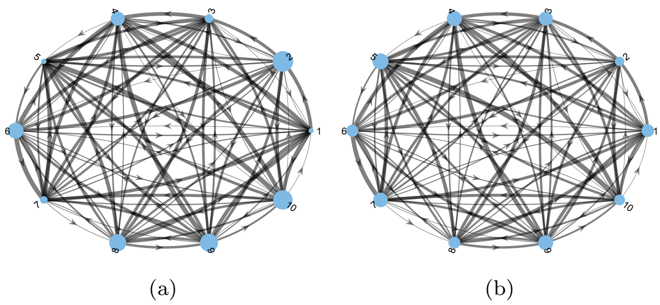

Figure 2. Evolution of the influence network without bots across the sequence of issues: (a) initial conditions, at the first issue; (b) static state, reached at the fifth issue

A scenario without the presence of bots is presented in Figure 2 to highlight the effect on humans' self-appraisal. In this case, the influence weights are more homogeneously distributed for all agents instead of being concentrated in robotic nodes.

4.2 Effect of bots and Reliable information

The second experiment to be presented here also consists of computational simulations but aiming to investigate the influence of malicious nodes in the spreading of a desired information over the network. Also, it is proposed the implementation of a stamp to mark if specific information is reliable, and the effects of this policy are analyzed.

The network configurations considered for those simulations were grounded in five fully-connected groups of twenty agents and an isolated node that only transmit information for some specific nodes in the groups. This isolated node could represent a mass media vehicle or a newspaper. However, for the present simulation, it will be considered as a central generator of malicious content, and the agents linked to it are robots infiltrated in the groups, accountable for sharing the information from the central node.

Three network configurations will be observed, differing from one another by the number of malicious agents per group. Also, in each configuration, there are ten random connections between non-malicious agents from different groups.

There are three types of information flowing through the network, represented by the colors blue, green, and red. When a node had not received any information yet, it will be represented by the color white. Figure 3 shows the initial states for each topology, and it is noticeable that each group starts with one node with each non-malicious information plus the malicious nodes.

Hence, at the first moment, it was carried 1000 simulations for each network configuration aiming to reduce the effect of the stochasticity from the results and obtain more meaningful conclusions from it. The information will be exchanged between agents following the model from the previous section and considering that all information had the same reliability coefficient, i.e., $\tau_m = 1 \;\; \forall m$.

Furthermore, a limit of three changes of information was imposed for each agent, modeling that once it made its mind about an issue, it will not change anymore. The results collected from this first set of simulations are summarized in Table 1; it is worth mentioning that the averages were rounded to the nearest integer.

Table 1. Results averaged over 1000 simulations for each scenario, considering the blue information as malicious and no information with reliability stamp

Number of Malicious Nodes	Blue Information (Bots)	Green Information	Red Information
1	48	26	26
2	69	16	16
3	80	11	10

From these results, it is noticeable that the spreading of the malicious information increases proportionally to the number of malicious agents per group.

Later, it was considered a reliability stamp for the information, meaning that a certain information can be marked as reliable. If an information is marked, then the reliability coefficient associated is set to $\tau_m = 2$.

Once more, 1000 simulations were carried for each scenario but marking the information represented by the color red as reliable. Namely, the reliability coefficients were set as $\tau_{blue} = \tau_{green} = 1$ and $\tau_{red} = 2$. In Table 2, the result of this second set of simulations is summarized.

The effect noticed was that the reliability stamp was able to reduce the spread of malicious information in the network. Still, the number of malicious agents per group influences the spreading. Another remark about

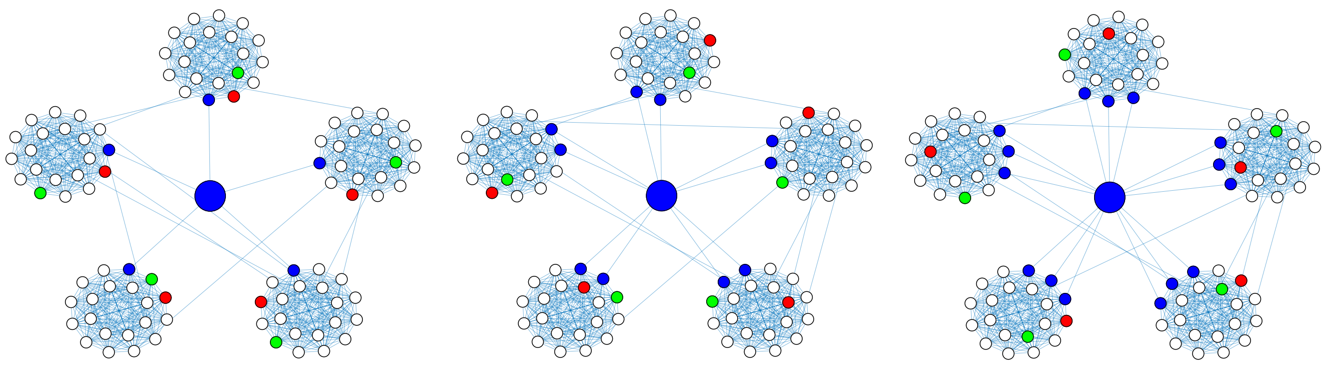

(a) (b) (c)

Figure 3. Initial conditions for each scenario: (a) one malicious node per group; (b) two malicious nodes per group; (c) three malicious nodes per group

Table 2. Results averaged over 1000 simulations for each scenario, considering the blue information as malicious and the red information as marked by the reliability stamp

Number of Malicious Nodes	Blue Information (Bots)	Green Information	Red Information (Marked)
1	21	9	70
2	38	7	55
3	52	6	42

these results is that for more than two malicious agents per group, the spreading of the malicious information surpasses that of the marked information.

It may be necessary to emphasize the fact that the bots infiltrated in the groups do not update their information with information from other members of the group. Hence, the malicious nodes are stubborn in the blue information throughout the simulations.

Figures 4-6 are presented to illustrate the difference made by the application of the reliability stamp; they picture one of the simulations carried in each scenario. It is important to emphasize that the graphs presented in Figures 4-6 do not match the result from Tables 1 and 2 since these are averaged over all the simulations carried.

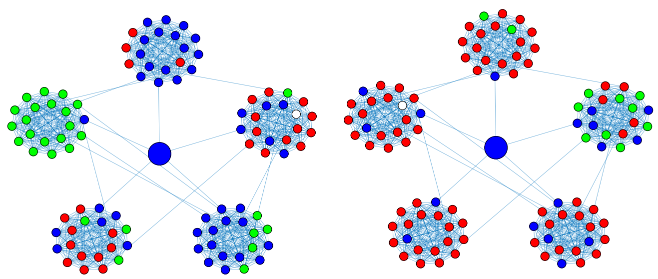

(a) (b)

Figure 4. Example of final states in the scenario with one malicious node per group: (a) without reliability stamp; (b) with reliability stamp

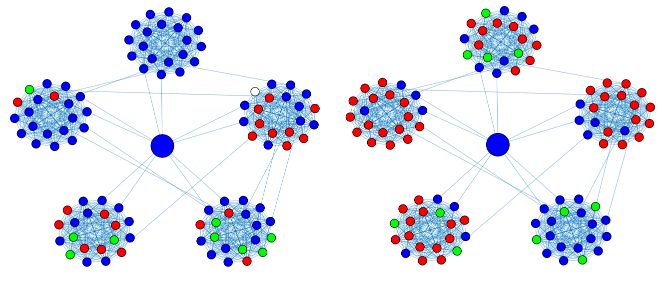

(a) (b)

Figure 5. Example of final states in the scenario with two malicious node per group: (a) without reliability stamp; (b) with reliability stamp

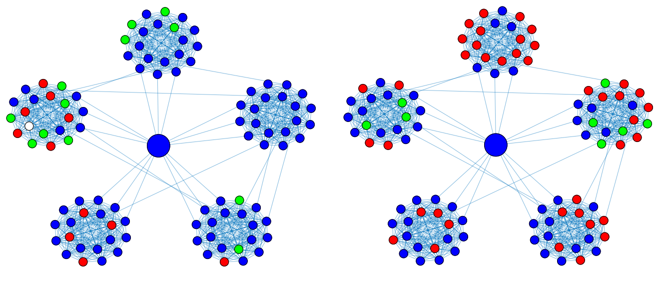

(a) (b)

Figure 6. Example of final states in the scenario with three malicious node per group: (a) without reliability stamp; (b) with reliability stamp

5. CONCLUSION

The implementation of control measures and utilization policies are important to safeguard ethical principles in the use of **OSNs**.In this paper, dynamic models that represent the diffusion of information and opinions in networks were applied to evaluate the influence of *bots* in two aspects: the humans' self-appraisal and the effect in the opinion evolution when compared with reliable information.

In the simulated examples, it was verified that the action of bots decisively influences the self-confidence of agents modeled as humans, as well as in the information diffusion and the opinion consolidation in the analyzed networks.

It was also concluded that the adoption of a mechanism to attribute some degree of reliability to the information sent to the network could counteract the influence of malicious nodes, *bots*. However, this impact could be neutralized with the increase in the numbers of bots.

Finally, by the results obtained in this work, adding measures to detect false news and implement information certification will only be effective by jointly implementing policies to detect and exclude bots from instant messaging application networks.

ACKNOWLEDGEMENTS

Financial support from Brazilian agencies Conselho Nacional de Desenvolvimento Científico e Tecnológico (**CNPq**) and Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (**CAPES/PRINT**): *Transformação Digital: Indústria e Serviços 4.0 , projeto Robótica: Competência e Pesquisa em Paralelismo e Reconfiguração* are gratefully acknowledged.

REFERENCES

Barabási, A.L. et al. (2016). *Network Science*. Cambridge University Press.

Bernays, E.L. (2005). *Propaganda*. Ig publishing.

Bonacich, P. (1972). Technique for analyzing overlapping memberships. *Sociological methodology*, 4, 176–185.

Boyd, D.M. and Ellison, N.B. (2007). Social network sites: Definition, history, and scholarship. *Journal of Computer-Mediated Communication*, 13(1), 210–230.

Boyd, S., Ghosh, A., Prabhakar, B., and Shah, D. (2006). Randomized gossip algorithms. *IEEE/ACM Transactions on Networking (TON)*, 14(SI), 2508–2530.

Bradshaw, S. and Howard, P. (2017). Troops, trolls and troublemakers: A global inventory of organized social media manipulation.

Bullo, F. (2019). *Lectures on Network Systems*. Kindle Direct Publishing, 1.3 edition. With contributions by J. Cortés, F. Dörfler, and S. Martínez.

DeGroot, M.H. (1974). Reaching a consensus. *Journal of the American Statistical Association*, 69(345), 118–121.

EODS (2016). *Compendium of international standards for elections*. Publication Office of the European Union. Fourth Edition.

Freeman, L. (2004). *The Development of Social Network Analysis: A Study in the Sociology of Science*. Empirical Press.

Friedkin, N.E., Jia, P., and Bullo, F. (2016). A theory of the evolution of social power: Natural trajectories of interpersonal influence systems along issue sequences. *Sociological Science*, 3, 444–472.

Han, B.C. (2017). *In the swarm: Digital prospects*, volume 3. MIT Press.

Howard, P. et al. (2018). How political campaigns weaponize social media bots. *IEEE Spectrum*.

Jia, P., MirTabatabaei, A., Friedkin, N.E., and Bullo, F. (2015). Opinion dynamics and the evolution of social power in influence networks. *SIAM review*, 57(3), 367–397.

Lippmann, W. (2017). *Public opinion*. Routledge.

McLuhan, M. (1994). *Understanding media : the extensions of man*. The MIT Press, New York. Reprint Edition.

Neudert, L.M. and Marchal, N. (2019). *Polarisation and the use of technology in political campaigns and communication*. European Parliamentary Research Service. STOA - Panel for the Future of Science and Technology.

Parsegov, S.E., Proskurnikov, A.V., Tempo, R., and Friedkin, N.E. (2017). Novel multidimensional models of opinion dynamics in social networks. *IEEE Transactions on Automatic Control*, 62(5), 2270–2285.

Proskurnikov, A.V. and Tempo, R. (2017). A tutorial on modeling and analysis of dynamic social networks. part i. *Annual Reviews in Control*, 43, 65–79.

Proskurnikov, A.V. and Tempo, R. (2018). A tutorial on modeling and analysis of dynamic social networks. part ii. *Annual Reviews in Control*, 45, 166–190.

Salem, F.A., Tchilian, R.d.S., and Moreno, U.F. (2019). Evolution of discrete opinions on human-swarm interaction in influence networks. *SBAI, Simpósio Brasileiro de Automação Inteligente*.

Shao, C., Ciampaglia, G.L., Varol, O., Yang, K.C., Flammini, A., and Menczer, F. (2018). The spread of low-credibility content by social bots. *Nature communications*, 9(1), 4787.

Tardáguila, C., Benevenuto, F., and Ortellado, P. (2018). Fake news is poisoning brazilian politics. whatsapp can stop it. *The New York Times*, 17(10).

Turculeţ, M. (2014). Ethical issues concerning online social networks. *Procedia - Social and Behavioral Sciences*, 149, 967 – 972. LUMEN 2014 - From Theory to Inquiry in Social Sciences, Iasi, Romania, 10-12 April 2014.

Vallor, S. (2016). Social networking and ethics. In E.N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2016 edition.

Van Dijck, J. (2013). *The culture of connectivity: A critical history of social media*. Oxford University Press.

Vosoughi, S., Roy, D., and Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380), 1146–1151.

Wang, P., Angarita, R., and Renna, I. (2018). Is this the era of misinformation yet: combining social bots and fake news to deceive the masses. In *Companion Proceedings of the The Web Conference 2018*, 1557–1561. International World Wide Web Conferences Steering Committee.

Zhang, X. and Ghorbani, A.A. (2019). An overview of online fake news: Characterization, detection, and discussion. *Information Processing & Management*, 102025.