

Multi-Issue Opinion Diffusion under Constraints

Sirin Botan

ILLC, University of Amsterdam
Netherlands
s.botan@uva.nl

Umberto Grandi

IRIT, University of Toulouse
France
umberto.grandi@irit.fr

Laurent Perrussel

IRIT, University of Toulouse
France
laurent.perrussel@irit.fr

ABSTRACT

Most existing models of opinion diffusion on networks neglect the existence of logical constraints that might correlate individual opinions on multiple issues. In this paper we study the diffusion of constrained opinions on a social network as an iterated process of aggregating neighbouring opinions. Individual views are modelled as vectors of yes/no answers to a number of propositions subject to integrity constraints, and each individual updates her opinion by looking at the aggregated opinion of her influencers. To overcome the problem of updating towards inconsistent influencing opinions, we propose a model based on individual updates on subsets of the issues of limited size called propositionwise updates. By adapting notions from the theory of boolean functions, we identify classes of integrity constraints on which propositionwise updates decrease the influence gap between nodes of the network and their influencers caused by the presence of an integrity constraint. Furthermore, we provide a detailed study of the termination of the proposed diffusion processes.

KEYWORDS

Social networks, judgment aggregation, opinion transformation, boolean functions

ACM Reference Format:

Sirin Botan, Umberto Grandi, and Laurent Perrussel. 2019. Multi-Issue Opinion Diffusion under Constraints. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019*, IFAAMAS, 9 pages.

1 INTRODUCTION


The diffusion of information in a social network is the subject of a vast literature combining sociological with algorithmic considerations (see, e.g., Easley and Kleinberg [10] and Jackson [25]), whose applications range from product adoption to disaster information management. In this diverse range of applications, only a few models have considered that opinions may be structured by the presence of an integrity constraint, relating the multiple issues at stake. Three recent examples are the work of Friedkin et al. [16] in sociological modelling of beliefs spread and change in a group, the work by Schwind et al. [30] in belief merging, and the analysis by Christoff and Grossi [6] of liquid democracy under constraints.

In this paper we consider individual opinions defined on a set of binary issues. The presence of constraints permits us to define a variety of applications: a participatory budgeting algorithm in which users decide which project to fund under a budget constraint;

a jury or a committee needing to reach a decision, or the problem of artificial agents influencing each other in a distributed manner.

We take a normative perspective to opinion diffusion in a constrained domain, replying to the question of how the diffusion process should be constructed to “fit” the integrity constraint defining the problem. Our focus is on settings where an opinion diffusion may precede a collective decision-making process. Let us showcase the main problems tackled by our paper with a concrete example of such a collective decision-making problem.

Example 1.1. Consider the case of four agents deciding whether a skyscraper (S), a hospital (H), or a new road (R) should be constructed in their city. Assume the first three agents are rather certain of their view as they have already considered their influencers’ opinion; the fourth agent is influenced by the first three, and will change her opinion according to the majority.¹



The law imposes that when both a hospital and a skyscraper are built then a new road must be constructed as well, a constraint that can be represented as $(S \wedge H) \rightarrow R$. Suppose that the first agent wants only the hospital; the second, only the skyscraper; and the third would like the whole package: skyscraper, hospital and road. Thus the fourth agent is facing an aggregated opinion which says yes to the skyscraper and the hospital, but no to the road; this opinion, of course, does not satisfy the constraint, hence blocking the influence of the first three agents on the fourth regardless her possible initial opinion.

We argue that information should not always spread through the network on all issues at once, or in other words, that agents should update their opinions locally rather than globally. If the fourth agent in the example above consulted her influencers on one single issue at a time, such as asking: “should a hospital be built?”, then she would be able to update her opinion to a consistent one by changing her opinion on this single issue. We call this opinion diffusion process, *propositionwise* diffusion: opinions are updated on subsets of issues, rather than on all issues at once.

The main contribution of this paper is to propose and characterise such propositionwise opinion diffusion processes. Such a model allows us first to identify the minimal amount of information exchange—in terms of the ‘scope’ of the questions asked by agents to influencers—that is needed for an information diffusion system to work as desired given a certain integrity constraint. We

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada. © 2019 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

¹Corresponding to a simple threshold model [22].

then characterise the class of constraints that allow influence to spread for bound k on the number of issues updated, borrowing and building on notions from the theory of boolean functions. We also characterise the improvement of the diffusion process when only a subset of issues are considered. Finally, we investigate the effects of the order of the updates on the result of the diffusion process, and provide intuitive initial results on the termination of iterative processes defined by propositionwise updates.

Related work

Diffusion on networks has been extensively studied from many angles in the field of social network analysis [10, 26]. Building on the classical work of Granovetter [22], DeGroot [9], and Lehrer and Wagner [28], a number of models have been introduced for the diffusion of *complex* opinions, such as knowledge bases [30, 31], preferences over alternatives [3], and binary evaluations [20, 21]. The model of binary opinion diffusion is also related to the literature on boolean networks [27], which is used for modelling biological regulatory networks (see, e.g., Shmulevich et al. [32])—which focuses on updates on one single binary issue. Our paper builds on the model of binary evaluations by including an integrity constraint that logically correlates the issues at stake. We examine the diffusion process when certain updates are rendered impossible by a global constraint. To the best of our knowledge; the only work whose primary focus is opinion diffusion under constraints is a recent paper by Friedkin et al. [16], which however represents opinions as real-valued beliefs, and the work of Christoff and Grossi [6], which is a special case of our propositionwise model for networks in which each node has at most one influencer.

In Section 4.2 we provide a detailed survey on results concerning the termination of opinion diffusion processes. Related to this idea, a recent line of research investigated the stabilisation of diffusion processes on a unanimous opinion [15, 17], with recent work by Auletta et al. [1] showing that for any network, there exists some majority of agents who will lead the process to terminate on a unanimous network when agents hold one of two possible opinions.

While we do not examine agents’ strategies and possible manipulation of the process in this paper, various strategic concerns such as manipulation and bribery have also been addressed in the literature. For binary opinions, Bredereck and Elkind [2] show that even for a single binary issue, identifying successful manipulation of the network through bribery, deleting edges, or controlling the order of asynchronous updates are all computationally hard problems. In the related model of preference diffusion, Faliszewski et al. [14] examine bribery on undirected graphs of voter clusters which consists of agents with similar preferences, and show that bribery is fixed-parameter tractable w.r.t. the number of candidates for rules which can be expressed as an integer linear program.

Paper overview. In Section 2 we define our model of propositionwise opinion diffusion under constraints, and we define and study a useful class of integrity constraints, which is used in Section 3 to obtain our main results. Section 4 studies network and aggregation properties to guarantee the termination of propositionwise diffusion models, and Section 5 concludes.

2 PRELIMINARIES

In this section we present our diffusion model for binary opinions over multiple issues correlated by an integrity constraint, as well as novel useful definitions for classes of integrity constraints.

2.1 Individual Opinions

Let $\mathcal{I} = \{p_1, \dots, p_m\}$ be a finite set of m issues, where each issue represents a binary choice. We call $\mathcal{D} = \{0, 1\}^{\mathcal{I}}$ the *domain* associated with this set of issues. For a finite set of agents $\mathcal{N} = \{1, \dots, n\}$, we say $B_i \in \mathcal{D}$ is the *opinion* of agent $i \in \mathcal{N}$ over all issues in \mathcal{I} . A vector $\mathbf{B} = (B_1, \dots, B_n)$ of all opinions of agents in \mathcal{N} is called a *profile*. An opinion B represents an agent’s acceptance/rejection of each of the issues in \mathcal{I} . For example, if $\mathcal{I} = \{p, q, r\}$, then $B = (110)$ is the opinion accepting p and q and rejecting r . We denote with $B_i(p)$ agent i ’s judgment on $p \in \mathcal{I}$ in the profile \mathbf{B} . Thus if $B = (110)$, then $B(p) = B(q) = 1$ and $B(r) = 0$.

An *integrity constraint* $\text{IC} \subseteq \mathcal{D}$ defines a domain of feasible opinions. We say that B is IC-consistent when $B \in \text{IC}$. For each agent i , we assume that $B_i \in \text{IC}$, meaning each individual opinion must satisfy the given integrity constraint. For instance, if we have three issues, p, q and r , and each agent can only accept at most two of the three, then $\text{IC} = \{(110), (011), (101), (100), (010), (001), (000)\}$. In further sections we will often assume that integrity constraints are represented compactly by means of a formula of propositional logic, such as $(\neg p \vee \neg q \vee \neg r)$ for the previous example.

2.2 The Social Influence Process

We assume that agents are connected by a *social influence network* $G = (\mathcal{N}, E)$ where $(i, j) \in E$ means agent i influences agent j and $\text{Infl}(i)_G = \{j \in \mathcal{N} \mid (j, i) \in E\}$ is the set of influencers of agent i in the network G . We model social influence as a transformation function, which takes as input a profile of IC-consistent opinions $\mathbf{B} = (B_1, \dots, B_n)$, and returns a set of profiles which are each the result of some opinion update on \mathbf{B} , depending on which set of agents update on which set of issues. If clear from the context, we omit reference to G and IC.

Let $F = (F_1, \dots, F_n)$ be composed of *aggregation procedures* $F_i : \text{IC}^{\text{Infl}(i)} \rightarrow \mathcal{D}$, one for each agent i . We assume that aggregation functions satisfy the minimal requirement of *unanimity*, i.e., whenever $B_j = B^*$ for all $j \in \text{Infl}(i)$ then $F_i(\mathbf{B}) = B^*$. In words, whenever all influencers are unanimous, F updates according to the influencers (no negative influence is possible). Our running example for an aggregator is the issue-by-issue majority rule, but we refer to the literature on judgment aggregation for other well-studied examples of aggregation rules [13, 23].

Our first model is a straightforward adaptation of propositional opinion diffusion [21], in which agents update their opinion on all issues towards the aggregated opinion of their influencers, provided that the latter satisfies the integrity constraint.

Definition 2.1. Given network G and aggregators F , we call *propositional opinion diffusion* the following transformation function:

$$\begin{aligned} \text{POD}_F(\mathbf{B}) = \{ \mathbf{B}' \mid \exists M \subseteq \mathcal{N} \\ \text{s.t. } B'_i = F_i(\mathbf{B}_{\text{Infl}(i)}) \text{ if IC-consistent and } i \in M \\ \text{and } B'_i = B_i \text{ otherwise.} \} \end{aligned}$$

POD_F defines the set of updates that are possible from a given consistent profile, depending on the set of agents that will perform the update. As shown by Example 1.1, social influence in presence of integrity constraints is often blocked when performing propositional updates on all the issues at the same time. Therefore, we now provide a definition for a *propositionwise* model of social influence.

Once an agent i and a subset of issues $S \subseteq \mathcal{I}$ is specified, aggregation functions F can be combined with a network G to obtain an update function for agent i 's opinions on the issues in S . If B and B' are two opinions and S a set of issues, let $(B|_{\mathcal{I} \setminus S}, B'|_S)$ be the opinion obtained from B with the opinions on the issues in S replaced by those in B' . We define an F -update as follows:

$$F\text{-UPD}(\mathbf{B}, i, S) = \begin{cases} (B_i|_{\mathcal{I} \setminus S}, F_i(\mathbf{B}_{\text{Inf}(i)})|_S) & \text{if IC-consistent} \\ B_i & \text{otherwise.} \end{cases}$$

That is, agent i looks at the aggregated opinion of its influencers $F_i(\mathbf{B}_{\text{Inf}(i)})$, and copies this opinion on all issues in S only if this results in a new opinion that is consistent with IC.


We are interested in varying degrees of communication among the agents, from simply asking one-issue questions to their influencers, to more complex updates involving all the issues at stake. We therefore give the following definition:

Definition 2.2. Given network G , aggregation functions F , and $1 \leq k \leq |\mathcal{I}|$, we call *k-propositionwise opinion diffusion* the following transformation function:

$$\text{PWOD}_F^k(\mathbf{B}) = \{ \mathbf{B}' \mid \exists M \subseteq \mathcal{N}, S : M \rightarrow 2^{\mathcal{I}} \text{ with } |S(i)| \leq k, \\ \text{s.t. } B'_i = F\text{-UPD}(\mathbf{B}, i, S(i)) \text{ for } i \in M \\ \text{and } B'_i = B_i \text{ otherwise.} \}$$

PWOD_F^k defines, for each consistent profile of opinions \mathbf{B} , the set of possible updates obtained by selecting a subset of agents $M \subseteq \mathcal{N}$ and a subset of issues $S(i) \subseteq \mathcal{I}$ for $i \in M$ on which agent i 's opinion is updated. Clearly, when $k = |\mathcal{I}|$ we have that $\text{POD}_F \subset \text{PWOD}_F^k$.

Example 2.3. Let us consider the situation in Example 1.1. The set of issues is $\mathcal{I} = \{S, H, R\}$ corresponding to building a skyscraper, a hospital, and a road, and the constraint in this situation is $(S \wedge H \rightarrow R)$. The agents are $\mathcal{N} = \{1, 2, 3, 4\}$, with opinions and social connections as in the following figure:



Clearly, if we denote with \mathbf{B} the profile described above, and if F is the strict majority rule, then $\text{POD}_F(\mathbf{B}) = \{\mathbf{B}\}$. Hence, no update on all propositions at the same time is possible. However, if we consider updates on one proposition at a time, we obtain that $\text{PWOD}_F^1(\mathbf{B}) = \{(010, 100, 111, \mathbf{010}), (010, 100, 111, \mathbf{100}), \mathbf{B}\}$. Observe that $\text{PWOD}_F^2(\mathbf{B}) = \text{PWOD}_F^1(\mathbf{B})$, as we would obtain the same set of profiles by updating on pairs of issues simultaneously.

The following example stresses the generality of our definition of propositionwise opinion diffusion.

Example 2.4 (Pairwise preference diffusion). The framework of pairwise preference diffusion by Brill et al. [3] can be seen as an

instance of PWOD_F^1 where F is the (strict) majority rule. To see this, consider a set A of alternatives. A linear order $>$ is an irreflexive, transitive and complete binary relation over A , which can be represented as a binary evaluation over a set of issues $\mathcal{I}_A = \{p_{a_i a_j} \mid (a_i, a_j) \in A \times A \text{ and } i < j\}$, such that $B(p_{ab}) = 1$ if and only if $a > b$. The integrity constraint $\text{IC}_>$ therefore contains all opinions over \mathcal{I}_A corresponding to linear orders over A . To overcome Condorcet cycles, i.e., individuals facing an aggregated majority which is not transitive, Brill et al. [3] propose to update on one pair of alternatives at the time, which corresponds to a propositionwise update on the analogous issue.

2.3 Iterative processes of opinion diffusion

To obtain the more classical view of diffusion as a discrete time iterative process, it is sufficient to combine PWOD_F^k with a turn-taking function: an *agent-scheduler* and an *issue-scheduler* deciding which issues are updated by which agent.


Let us now define the usual notions helping at characterizing some diffusion processes, namely *reachability* and *termination*.

We say that a profile \mathbf{B}' is PWOD_F^k -reachable from profile \mathbf{B} if there exists a sequence of profiles $\mathbf{B}_1, \dots, \mathbf{B}_t$ such that $\mathbf{B}_1 = \mathbf{B}$, $\mathbf{B}_t = \mathbf{B}'$, and for each $1 < j \leq t$ we have that $\mathbf{B}_j \in \text{PWOD}_F^k(\mathbf{B}_{j-1})$. We also introduce the following concept:

Definition 2.5. A profile \mathbf{B} is a *termination profile* for PWOD_F^k and IC if $\text{PWOD}_F^k(\mathbf{B}) = \{\mathbf{B}\}$.

Termination profiles are fixed points of PWOD_F^k . We stress the role of IC in determining which updates can be performed.

Example 2.6. Consider a scenario similar to Example 1: 3 agents are voting on three proposals for their city; a skyscraper, (S) an hospital (H), and a new road (R), with $\text{IC} = (S \wedge H \rightarrow R)$. The three agents are now connected in the following network, where the initial profile is $\mathbf{B} = (111, 011, 101)$.



Assume that F_i is the strict majority rule for each i , accepting an issue only if a strict majority of their influencers accept it. If we let all agents update simultaneously under POD_F we reach profile $(111, 011, 011)$; agent 3 keeping her valuation in absence of a strict majority of influencers against it. An additional round of POD_F leads to profile $(011, 011, 011)$, a consensual *termination* profile.

Consider now PWOD_F^1 : assume all agents updating simultaneously, we reach the profile $(111, 011, 011)$ after two rounds, updating first on issue S , then on issue H . Two other rounds again on issue S , and on issue H leads to the termination profile $(011, 011, 011)$. Observe that this particular network configuration always leads to the same termination profile (more results in this line in Section 4.5).

2.4 Geodetic Integrity Constraints

In this section we build on notions from the theory of boolean functions (see, e.g., Crama and Hammer [8]) to identify a useful

class of integrity constraints that we will later use to characterise termination profiles of our diffusion model.

Recall that $\mathcal{D} = 2^{\mathcal{I}}$ and that $\text{IC} \subseteq \mathcal{D}$. In this section we will call an opinion $B \in \text{IC}$ a model of IC, importing the terminology from propositional logic. Given two opinions B and $B' \in \mathcal{D}$, recall that the *Hamming distance* between them is $H(B, B') = \sum_{p \in \mathcal{I}} |B(p) - B'(p)|$. Consider the following:

Definition 2.7. Let IC be an integrity constraint for issues \mathcal{I} . The k -graph of IC is given by $\mathcal{G}_{\text{IC}}^k = \langle \text{IC}, E_{\text{IC}}^k \rangle$, where:

- (i) the set of nodes is the set of $B \in \text{IC}$,
- (ii) the set of edges E_{IC}^k is defined as follows: $(B, B') \in E_{\text{IC}}^k$ iff $H(B, B') \leq k$, for any $B, B' \in \text{IC}$.


As it is clear from Definition 2.7, $\mathcal{G}_{\text{IC}}^k \subseteq \mathcal{G}_{\mathcal{D}}^k$ for all IC. We say that a path of $\mathcal{G}_{\mathcal{D}}^k$ is also a path of $\mathcal{G}_{\text{IC}}^k$ if all nodes on the path are also nodes of $\mathcal{G}_{\text{IC}}^k$. We are now ready to give the following definition:

Definition 2.8. An integrity constraint IC is k -geodetic if and only if for all B and B' in IC, at least one of the shortest paths from B to B' in $\mathcal{G}_{\mathcal{D}}^k$ is also a path of $\mathcal{G}_{\text{IC}}^k$.

For ease of notation, we denote a 1-geodetic IC as geodetic *tout court*, borrowing the term from the equivalent definition for boolean functions [11]. To illustrate our definitions, consider the following:

Example 2.9. First, consider the integrity constraint of our running example: $\text{IC} = S \wedge H \rightarrow R$ or $\text{IC} = \{(000), (001), (010), (011), (100), (101), (111)\}$. Clearly, all shortest paths between any two models of IC belong to $\mathcal{G}_{\text{IC}}^1$ and thus IC is geodetic.

Assume now that $\text{IC} = \{(000), (001), (010), (100), (011), (111)\}$. The graph below corresponds to $\mathcal{G}_{\text{IC}}^1$, connecting only those models that satisfy IC with a continuous edge. The graph consisting of all edges (continuous and dashed) corresponds to $\mathcal{G}_{\mathcal{D}}^1$.



We can now observe that IC is not geodetic: the shortest paths between (100) and (111) in $\mathcal{G}_{\mathcal{D}}^1$ pass through either (110) or (101), which however are not nodes of $\mathcal{G}_{\text{IC}}^1$.

Preferences and Geodetic Constraints. An important class of integrity constraints that are geodetic is the one commonly used to represent preferences as linear orders over a set of alternatives (see Example 2.4). To see this, let $<$ and $<'$ be two distinct linear orders over a set A of alternatives. Then, they also must differ on a pair which is adjacent in one of them, i.e., there exists a pair ab such that $B(p_{ab}) \neq B'(p_{ab})$ and there is no $c \in A$ such that $a > c > b$ or $b > c > a$.² Knowing this, it becomes straightforward to show that $\text{IC}_{>}$ is geodetic (for the particular encoding of preferences explained in Example 2.4). Similar encodings can be used to show that

²This result is folklore, a formal proof is in [12].

partial and weak orders and equivalence relations can be modelled by geodetic constraints.

Budget constraints. Another important class is that of *budget constraints*, which specify the list of subsets of the issues \mathcal{I} that do not exceed a given budget. Such formulas can be shown to be *negative* formulas, i.e., there is a DNF representation in which all propositional symbols only occur as negated. This specific representation guarantees geodeticity [11].

Syntactic restrictions. Integrity constraints are typically represented compactly by means of propositional formulas. It is easy to see that all conjunctions of literals are k -geodetic for any k , as well as simple clauses of any length. However, the conjunction of two k -geodetic formulas is not necessarily k -geodetic, as can be seen by considering an XOR formula such as $(\neg p \vee q) \wedge (p \vee \neg q)$. Clearly, this formula is not geodetic. This example also shows that known syntactic restrictions such as Horn clauses or 2CNF formulas are not relevant for determining geodeticity.

A number of logical characterisations of 1-geodetic integrity constraints can be found in the work of Ekin et al. [11]. To the best of our knowledge, for k -geodetic constraints no such characterisation is available. While similar results would be outside the scope of this paper, we show the following simple proposition, whose proof is straightforward from our definitions:

PROPOSITION 2.10. *If IC for a set of issues \mathcal{I} is k -geodetic, then it is also k -geodetic for any larger set of issues $\mathcal{I}' \supseteq \mathcal{I}$.*

We also obtain a more operational definition of k -geodeticity of a constraint, in the following:

PROPOSITION 2.11. *An integrity constraint IC is k -geodetic iff for all models $B_1, B_2 \in \text{IC}$, there is a path in $\mathcal{G}_{\text{IC}}^k$ from B_1 to B_2 of length smaller than $\left\lceil \frac{H(B_1, B_2)}{k} \right\rceil$.*

PROOF SKETCH. Let B and B' be two models of IC. The length of the shortest path from B to B' in the hypercube $\mathcal{G}_{\mathcal{D}}^k$ is exactly $\left\lceil \frac{H(B, B')}{k} \right\rceil$, since $H(B, B')$ is the number of issues that has to be changed to move from B to B' , and the edges in $\mathcal{G}_{\mathcal{D}}^k$ change k symbols at most. As $\mathcal{G}_{\text{IC}}^k \subseteq \mathcal{G}_{\mathcal{D}}^k$, if there is a path of minimal length connecting B to B' in $\mathcal{G}_{\text{IC}}^k$, then it is one of the shortest paths of $\mathcal{G}_{\mathcal{D}}^k$. By repeating for all B and B' in IC we obtain the statement. \square

3 PROPOSITIONWISE UPDATES

In this section we show that propositionwise updates of size k reduce the influence gap between an agent and its influencers, provided that the constraint under consideration is k -geodetic.

3.1 Reachability under k -Geodetic Constraints

In most examples considered so far PWOD_F^k was able to perform additional updates compared to POD_F even when $k < m$ for m issues. Consider however the following example:

Example 3.1. Let $\text{IC} = p \text{ XOR } q$ which is 2-geodetic but not 1-geodetic, and let there be two agents with $E = \{(1, 2)\}$, that is, 1 is the only influencer of 2. Assume $B_1 = (0, 1)$ and $B_2 = (1, 0)$. Whatever the unanimous F , we have that $\text{PWOD}_F^2(B)$ contains profile B'

in which $B'_1 = B'_2 = (0, 1)$, while we have that $\text{PWOD}_F^1(\mathbf{B}) = \{\mathbf{B}\}$, i.e., \mathbf{B} is a termination profile for PWOD_F^1 .

Given an integrity constraint IC it is therefore of crucial importance to identify the right “level of communication”, i.e., the value of k in the Definition 2.2, that allows propositionwise updates to reach the same profiles that are reachable by POD_F , and eventually move further. The previous example illustrates that PWOD_F^1 does not perform well when considering a 2-geodetic constraint. More generally, we have the following central result:

THEOREM 3.2. *Let IC be an integrity constraint, and let \mathbf{B}' be POD_F -reachable from an IC-consistent initial profile \mathbf{B} . Then, \mathbf{B}' is PWOD_F^k -reachable from \mathbf{B} iff IC is k -geodetic.*

PROOF. Let IC be k -geodetic. We prove that any profile that is POD_F -reachable from \mathbf{B} is also reachable by propositionwise updates of size at most k . Wlog, we can assume that \mathbf{B}' has been obtained from profile \mathbf{B} with one single agent updating from B_i to $F(\mathbf{B}_{\text{Inf}(i)})$. Since IC is k -geodetic, and both B_i and $F(\mathbf{B}_{\text{Inf}(i)})$ are IC-consistent, by Definition 2.8 there exists an IC-consistent shortest path in $\mathcal{G}_{\mathcal{D}}^k$ that connects the two opinions. Let B_i^1 be the first model on such path after B_i , and let p_1, \dots, p_ℓ be the issues on which B_i and B_i^1 differ. By the definition of $\mathcal{G}_{\mathcal{D}}^k$ we know that $\ell \leq k$. Moreover, since B_i^1 is on the shortest path between B_i and $F(\mathbf{B}_{\text{Inf}(i)})$, we can infer that $B^1 = (B_i \upharpoonright_{I \setminus S}, F_i(\mathbf{B}_{\text{Inf}(i)}) \upharpoonright_S)$, where $S = \{p_1, \dots, p_\ell\}$. Let us now set $I = \{i\}$ and S as defined above in Definition 2.2, obtaining that $(B_1, \dots, B_i^1, \dots, B_n) \in \text{PWOD}_F^k(\mathbf{B})$. To conclude the proof, repeat the same construction for all the IC-consistent opinions B_i^k on the shortest path to $F_i(\mathbf{B}_{\text{Inf}(i)})$, showing that \mathbf{B}' is PWOD_F^k -reachable from \mathbf{B} .

For the converse, assume that IC is not k -geodetic, and let us show an example of a profile \mathbf{B}' that is POD_F -reachable from \mathbf{B} but not PWOD_F^k -reachable. Let there be two agents, and let $E = \{(1, 2)\}$. Since IC is not k -geodetic, there exist two IC-consistent opinions B_1 and B_2 that are not connected in $\mathcal{G}_{\text{IC}}^k$ by any of the shortest paths of $\mathcal{G}_{\mathcal{D}}^k$. Profile $\mathbf{B} = (B_1, B_2)$ is therefore a termination profile of PWOD_F^k , in which however $B_2 \neq B_1 = F(\mathbf{B}_{\text{Inf}(2)})$ (the last equality obtained by unanimity of F). However, $\text{POD}_F(\mathbf{B})$ includes the unanimous profile (B_1, B_1) , since agent 2 can update directly on all issues, thus concluding the proof. \square

The converse of Theorem 3.2 does not hold, as shown by Example 2.3, where profile $\{(010, 100, 111, 010)\}$ is PWOD_F^1 -reachable but not POD_F -reachable from the initial profile \mathbf{B} .

Theorem 3.2 shows that if IC is k -geodetic then k is the appropriate “level of communication” to set in definition of PWOD_F^k to be able to reach at least the same profiles as those that are reachable by updating on all issues at the same time. We now prove that nothing would be gained by considering any K larger than k :

LEMMA 3.3. *Let IC be k -geodetic. If \mathbf{B}' is PWOD_F^K -reachable from a consistent profile \mathbf{B} for $K \geq k$, then it is also PWOD_F^k -reachable.*

PROOF SKETCH. Suppose \mathbf{B}' is PWOD_F^K -reachable from \mathbf{B} . A similar construction to the proof of Theorem 3.2 can be used to show that the updates on K issues from \mathbf{B} to \mathbf{B}' can be simulated by (a larger number of) smaller updates of size at most k . \square

The following is an immediate consequence of Lemma 3.3:

COROLLARY 3.4. *Let IC be k -geodetic, and let \mathbf{B} be a termination profile of PWOD_F^k . Then, \mathbf{B} is also a termination profile for PWOD_F^K for any $K \geq k$.*

Thus, when confronted with a situation of opinion diffusion under constraint IC, our results suggests to set PWOD_F^k with the minimal k such that IC is k -geodetic (for the computational complexity of determining this parameter see Section 3.3).

3.2 Minimising the Influence Gap

We now want to investigate how the presence of an integrity constraint entails some form of cost on the diffusion process. This cost should reflect the difference between an opinion diffusion conducted under constraints rather than without. In this section we omit the reference to F , which is assumed to be clear from the context. We define the following notion:

Definition 3.5. If \mathbf{B} is a profile, and G a network, the *influence-gap* of \mathbf{B} on G is defined as follows:

$$\text{GAP}(\mathbf{B}, G) = \sum_{i \in \mathcal{N} \text{ s.t. } \text{Inf}(i) \neq \emptyset} H(B_i, F(\mathbf{B}_{\text{Inf}(i)}))$$

The influence gap of a profile is therefore the sum of all the disagreements between each agent and the aggregated opinion of her influencers, for those agents that have influencers. In the absence of integrity constraints, a strictly positive influence gap implies that updates are still possible. This is not the case when integrity constraints are present. Hereafter, we only focus on the termination profiles in order to quantify the loss entailed by a constraint IC:

Definition 3.6. Let G be a network, and IC a constraint over \mathcal{I} . The *price* of IC over G is the maximal influence gap among all PWOD_F^k -termination profiles for G and \mathcal{I} .

To simplify notation, we denote price_{IC} the price of IC for POD_F , and $\text{price}_{\text{IC}}^k$ the price of IC for PWOD_F^k for $k < m$. Clearly, in the absence of constraints (i.e., $\text{IC} = \mathcal{D}$), then $\text{price}_{\text{IC}} = 0$. An immediate consequence of Theorem 3.2 and Example 2.3 is the following:

PROPOSITION 3.7. *If IC is k -geodetic, then $\text{price}_{\text{IC}}^k \leq \text{price}_{\text{IC}}$, while the converse does not hold.*

We now show tight bounds for the price of an integrity constraint under POD_F and PWOD_F^k , showing that for the latter this price is lower. We begin with the following:

PROPOSITION 3.8. *Let $\text{IC} \neq \mathcal{D}$ and \bar{n} be the number of agents having at least one influencer in G . The following is a tight bound:*

$$\text{price}_{\text{IC}} \leq \left(\max_{B \in \text{IC}, B' \notin \text{IC}} H(B, B') \right) \times \bar{n}$$


PROOF. The upper bound is easy to obtain. Let B_i be the opinion of any agent with at least one influencer in a termination profile \mathbf{B} . $F(\mathbf{B}_{\text{Inf}(i)})$ cannot be IC-consistent, for the profile would not be a termination profile. Hence, $F(\mathbf{B}_{\text{Inf}(i)}) \notin \text{IC}$ and the above bound applies. For tightness, consider the following situation. Let $\text{IC} = p_1 \vee \dots \vee p_m$, and let there be $m+1$ individuals such that $B_i(p_j) = 1$ iff $i = j$ and $1 \leq i \leq m$, and $B_{m+1} = (1, \dots, 1)$. Let F be the strict majority rule, and let G be such that agents $1, \dots, m$ all influence agent

$m + 1$. Clearly, \mathbf{B} is an IC consistent profile, $\text{POD}_F(\mathbf{B}) = \{\mathbf{B}\}$, and $\text{GAP}(\mathbf{B}, G) = m$, since agent $m + 1$ is the only agent having at least one influencer, and is confronted with an inconsistent aggregated opinion of $(0, \dots, 0)$. \square

Let us now move to $\text{price}_{\text{IC}}^k$ and show that it is significantly lower. We first need an additional definition:

Definition 3.9. An IC-inconsistent-path from B to B' is a path of $\mathcal{G}_{\mathcal{D}}^k$ such that all opinions B'' between B and B' are not in IC. Two opinions B and B' are *totally IC-disconnected* if all shortest paths $\mathcal{G}_{\mathcal{D}}^k$ from B to B' are IC-inconsistent paths.

Example 3.10. Let there be three issues, and let $\text{IC} = \{(011), (111), (101), (100)\}$. IC is 1-geodetic and $\mathcal{G}_{\text{IC}}^1$ is represented as follows:



Opinions (011) and (000) are totally disconnected, since paths $\langle(011), (010), (000)\rangle$ and $\langle(011), (001), (000)\rangle$ are both inconsistent-paths.

We are now ready to prove the following:

PROPOSITION 3.11. Let $\text{IC} \neq \mathcal{D}$ be k -geodetic. Let \bar{n} be the number of agents having at least one influencer. The following is a tight bound:

$$\text{price}_{\text{IC}}^k \leq \left(\max_{\substack{B \in \text{IC}, B' \notin \text{IC} \\ B, B' \text{ totally IC-disconnected}}} H(B, B') \right) \times \bar{n}$$

PROOF. To prove the upper bound, let \mathbf{B} be a termination profile, and let i be such that $\text{Inf}(i) \neq \emptyset$. Since \mathbf{B} is a termination profile, if B_i differs from $F(\mathbf{B}_{\text{Inf}(i)})$, then $F(\mathbf{B}_{\text{Inf}(i)}) \notin \text{IC}$, for otherwise by k -geodeticity there would be possible updates from B_i to $F(\mathbf{B}_{\text{Inf}(i)})$. Moreover, B_i and $F(\mathbf{B}_{\text{Inf}(i)})$ must be totally IC-disconnected. To see this, assume that there is a $B' \in \text{IC}$ on one of the shortest paths between the two models. By k -geodeticity again, there is a shortest path of IC models from B_i to B' , which translates into a sequence of PWOD_F^k updates from B_i towards $F(\mathbf{B}_{\text{Inf}(i)})$, against the assumption that \mathbf{B} is a termination profile. By observing that $B_i \in \text{IC}$ while $F(\mathbf{B}_{\text{Inf}(i)})$ is not, we obtain the desired bound.

For tightness, consider the following case. Let IC be as in Example 3.10, and let there be four individuals such that $E = \{(1, 4), (2, 4), (3, 4)\}$. If we take F as the unanimous rule, which accepts an issue only if all the agents accept it, and profile $\mathbf{B} = (100, 101, 011, 011)$, then $F(100, 101, 011) = 000$. Thus, $\text{GAP}(\mathbf{B}, G) = H(011, 000) = 2$, corresponding to the formula in the statement. \square

Let us go back to Example 3.10. The maximal distance between an IC-model and non-model is 3, take for example 000 and 111. Instead, the maximal distance between an IC model that is totally disconnected from a non-model equals to 2, as can be seen by considering model 011 and non-model 000.

3.3 Computational Complexity

As observed in Section 3.1, when defining opinion diffusion processes in presence of an integrity constraint, the best option is to allow for propositionwise updates on up to k issues, where k is the smallest number such that IC is k -geodetic. We now investigate the computational complexity of finding such a threshold.

THEOREM 3.12. Let IC be a constraint over m issues and $k < m$. Checking whether IC is k -geodetic is co-NP-complete.

PROOF. To find a counterexample for k -geodeticity, it is sufficient to find two models B and B' of IC that are not connected by any of the shortest paths of $\mathcal{G}_{\mathcal{D}}^k$. A co-NP algorithm guesses two opinions B and B' , checks that they are IC-consistent, and that for all subsets $S \in \mathcal{I}$ of $|S| \leq k$ we have that $(B|_{\mathcal{I} \setminus S}, B'|_S) \notin \text{IC}$, showing a counterexample to the k -geodeticity of IC whose correctness can be checked in polynomial time.

As for hardness, we exploit a result by Hegedüs and Megiddo [24], stating that the membership problem for classes of boolean functions that satisfy the *projection property* is co-NP-hard. To show that the class of k -geodetic IC has the projection property we have to show that (a) the constant function \top is k -geodetic, (b) that for any k there is always a non- k -geodetic function, and (c) that if IC is k -geodetic then both $\text{IC} \wedge p$ and $\text{IC} \wedge \neg p$ must also be k -geodetic for all $p \in \mathcal{I}$. (a) is an immediate consequence of the definition of k -geodetic constraints. As for (b), we need to show that for any k there always exists a non-geodetic IC. Let $k < m$, and consider the full graph $\mathcal{G}_{\mathcal{D}}^k$. Let B and B' be two assignments at distance exactly $m - 1$. Let IC be composed of B, B' , and all other assignments except for those on any shortest path of $\mathcal{G}_{\mathcal{D}}^k$ between B and B' . Clearly, B and B' are not connected in $\mathcal{G}_{\text{IC}}^k$ by any of the shortest path of $\mathcal{G}_{\mathcal{D}}^k$, and IC is not empty. To show (c), suppose that B and B' are two models of $\text{IC} \wedge p$ that are not connected by any shortest path of $\mathcal{G}_{\mathcal{D}}^k$. Since B and B' are also models of IC, we have that $\mathcal{G}_{\text{IC} \wedge p}^k \subseteq \mathcal{G}_{\text{IC}}^k$ and therefore IC is not k -geodetic, against the assumption. \square

For 1-geodetic constraints, the hardness result above has already been shown by Ekin et al. [11].

Combining a simple binary search with the co-NP-complete problem shown in Theorem 3.12, we obtain the following:

THEOREM 3.13. Let IC be an integrity constraint over m issues and let $k < m$. Checking whether k is the minimal $k < m$ such that IC is k -geodetic is in Θ_2^P .

4 TERMINATION OF ITERATIVE DIFFUSION

In this section we analyse the termination of discrete-time iterative processes defined by PWOD_F^k updates, generalising results from the literature and opening interesting directions for future work.

4.1 Basic Definitions

Recall our Definitions 2.1 and 2.2, introducing propositionwise opinion diffusion as a transformation function that associates a set of updated profiles with every IC-consistent profile. Thus, PWOD_F^k induces a state transition system in which states are all profiles of IC-consistent opinions, and each transition is induced by the choice of a set of updating individuals M and a set of issues $S(i)$ for each

updating individual. Termination states, as defined by Definition 2.5, are the attractors of the transition system.

In line with the most recent literature on propositional opinion diffusion [2, 3, 21] and on boolean networks [27], we define *asynchronous* opinion diffusion processes by restricting transitions to those involving only one single agent at a time, and *synchronous* ones by restricting transitions to those involving all individuals. The two processes could be equivalently defined by introducing an agent-scheduler, indicating at each point in time the set of updating agents: all the agents for a synchronous scheduler, and sets of cardinality one for the asynchronous scheduler. A specific instance of asynchronous scheduler is the one that follows a predetermined order on \mathcal{N} in the updates, as studied by Goles and Tchuente [19].

We call a transition from \mathbf{B} to \mathbf{B}' *effective* if $\mathbf{B}' \neq \mathbf{B}$. We say that an opinion diffusion process *terminates universally* if there exists no infinite sequence of effective transitions starting at any IC-consistent profile, while it *terminates asymptotically* if from any IC-consistent profile there exists a sequence of transitions that reaches a termination profile. Note that for the case of synchronous POD_F there is only one sequence of effective updates for each initial profile, hence for this iterative process the two notions coincide.

4.2 Previous work

We summarise here results from related work that are close to our setting, using whenever possible the terminology introduced above.

Synchronous processes are the most studied. For one single binary issue POD_F and PWOD_F coincide, and the work of Goles and Olivos [18] showed that such processes either terminate or produce infinite sequences of effective transitions with period 2, under the assumption that F is a (generalised) threshold rule and the graph is undirected. For directed graphs and arbitrary aggregation procedures, Christoff and Grossi [7] characterised the set of profiles that lead to termination on a given graph, generalising results by Grandi et al. [21] who studied sufficient conditions on the network graph to guarantee universal termination. As seen in Example 2.4, preference diffusion can be viewed as an instance of PWOD_F^1 , and the work of Brill et al. [3] showed the asymptotic termination of the synchronous update process on arbitrary graphs, under a restrictive condition on the initial profile.

For asynchronous diffusion processes, Brederick and Elkind [2] showed that for one issue, majority POD_F asymptotically terminates on any undirected graph, and identify two sequences of transitions leading to two well-defined termination profiles either maximising the number of 0 or the number of 1 in the graph.

The work of Christoff and Grossi [6] is to the best of our knowledge the only one focusing on arbitrary integrity constraints on binary issues, albeit on specific networks called delegation graphs, where each node has at most one influencer. Finally, well-established termination results for boolean networks only consider the case of a single binary issue (see, e.g., Cheng et al. [4]).

4.3 Universal Termination

Let a complete graph be a graph $G = (\mathcal{N}, E)$ where $E = \mathcal{N} \times \mathcal{N}$, and let us define the following property of aggregation procedures:

Ballot-Monotonicity: for all profiles $\mathbf{B} = (B_1, \dots, B_n)$, if $F(\mathbf{B}) = B^*$ then for any $1 \leq i \leq n$ we have that $F(\mathbf{B}_{-i}, B^*) = B^*$.

Ballot monotonicity avoids aggregators modelling situations of “negative influence”. Generalising a result by Grandi et al. [21], originally stated for synchronous processes, we can show that:

THEOREM 4.1. *Let G be the complete graph. Synchronous POD_F terminates universally, and asynchronous POD_F terminates universally if F is ballot-monotonic.*

PROOF. Let \mathbf{B}^0 be an arbitrary initial profile. On a complete graph $\text{Inf}(i) = \mathcal{N}$ for all i , therefore every individual updates towards the same aggregated opinion $F(\mathbf{B}^0)$. The case of synchronous POD_F is a straightforward adaptation of the analogous result by Grandi et al. [21] and its proof is omitted. For the case of asynchronous POD_F , consider the influence gap of Definition 3.5. By ballot-monotonicity of F , the influence gap is a potential function for this iterative process: after any number of updates t , the aggregated opinion will not change, i.e., $F(\mathbf{B}^t) = F(\mathbf{B}^0)$, since by ballot-monotonicity $F(\mathbf{B}_{-i}, F(\mathbf{B})) = F(\mathbf{B})$. Therefore, the influence gap is strictly decreasing at each effective update. \square

For the case of PWOD_F^k we need to introduce a stronger property for the aggregation function, which is known in the literature on judgment aggregation as monotonicity (see, e.g., Endriss [13]):

Monotonicity: for any $j \in \mathcal{I}$ and any profiles \mathbf{B}, \mathbf{B}' , if $B_i(j)=1$ entails $B'_i(j)=1$ for all $i \in \mathcal{N}$, and for some $s \in \mathcal{N}$ we have that $B_s(j)=0$ and $B'_s(j)=1$, then $F(\mathbf{B})(j)=1$ entails $F(\mathbf{B}')(j)=1$

Clearly, monotonicity is a stronger property and implies ballot-monotonicity. We show the following:

THEOREM 4.2. *If G is the complete graph and F is monotonic, then both synchronous and asynchronous PWOD_F^k terminate universally.*

PROOF SKETCH. The same proof works for synchronous and asynchronous PWOD_F^k . As in the previous proof, we show that $F(\mathbf{B}^t) = F(\mathbf{B}^0)$ for any sequence of t updates, and therefore that the influence gap $\text{GAP}(\mathbf{B}^t, G)$ is a potential function. By induction, assume that at time $t-1$ a set of agents M —either a singleton for asynchronous processes or equal to \mathcal{N} —is updating on issues defined by the selection function $S : M \rightarrow 2^{\mathcal{I}}$. Assume that $F(\mathbf{B}^t)$ is IC-consistent: since every update reinforces the agents agreements with $F(\mathbf{B}^t)$ on those issues in the image of S , by monotonicity we conclude that $F(\mathbf{B}^{t+1}) = \text{PWOD}_F^k(\mathbf{B}^t) = F(\mathbf{B}^t)$. If $F(\mathbf{B}^t)$ is not IC-consistent, then some of these updates are blocked, as they would generate an inconsistent result. Still, this implies that some of the issues are copied towards $F(\mathbf{B}^t)$, and therefore by monotonicity that $F(\mathbf{B}^{t+1}) = F(\mathbf{B}^t)$. \square

Both assumptions of ballot-monotonicity and monotonicity are necessary in the respective theorems.

A directed acyclic graph (DAG) is a directed graph that contains no cycle involving two or more vertices. By constructing a suitable potential function we can prove the following theorem, whose proof is omitted in the interest of space:

THEOREM 4.3. *If G is a DAG and F is ballot-monotonic (respectively, monotonic), then both synchronous and asynchronous POD_F (respectively, PWOD_F^k) terminate universally.*

Universal termination cannot be guaranteed even on simple cycles, as can easily be shown on a cycle of arbitrary length with one

of the agents having opinion 1 and all others 0. In conclusion, POD_F and PWOD_F^k are comparable in terms of universal termination, with the latter requiring a slightly stronger property of monotonicity to avoid situations of negative influence.

4.4 Asynchronous Asymptotic Termination

The following condition adapts a property from Brill et al. [3], requiring all influence updates to be based on IC-consistent opinions.

Definition 4.4. A pair (\mathbf{B}^0, G) , where \mathbf{B}^0 is a profile and G a network, has the *local IC-consistency property* if for all profiles \mathbf{B} that is reachable from \mathbf{B}^0 and each $i \in \mathcal{N}$ we have that $F(\mathbf{B}_{\text{Inff}(i)})$ is IC-consistent.

Depending on the diffusion process considered, the above definition needs to be specified considering profiles that are reachable via POD_F or PWOD_F^k . Albeit restrictive, observe that this property holds for any function F which is collectively rational for the given integrity constraint, including all distance-based functions [29], as well as for simple cycles, trees, and any network where each node has at most one influencer. We prove the following:

THEOREM 4.5. *If \mathbf{B}^0 is an IC-consistent profile such that (\mathbf{B}^0, G) satisfies the local IC-consistency property, then asynchronous POD_F terminates asymptotically.*

PROOF SKETCH. The proof is based on an original construction by Chierichetti et al. [5]. It is also used by Brill et al. [3] and Bredereck and Elkind [2] and is thus only sketched. Following a fixed ordering of the issues, perform two rounds of asynchronous updates for each issue: a first round in which all individuals who disagree with their influencers and have opinion 0 update their opinion to 1, and a second round in which individuals who disagree with their influencers and have opinion 1 update their opinion to 0. In the resulting profile no further POD_F -update is possible. For suppose not. Wlog we can assume that such an update will revert issue i_1 from 0 to 1. But such an update should have taken place in the first round, since by the local IC-consistency property the aggregated opinion of influencers is always IC-consistent, and thus all possible updates are covered by the above procedure. \square

We can now use our Theorem 3.2 to show that:

COROLLARY 4.6. *If \mathbf{B} is an IC-consistent profile such that (\mathbf{B}, G) satisfies the local IC-consistency property and IC is k -geodetic, then asynchronous PWOD_F^k terminates asymptotically.*


PROOF SKETCH. If from every (\mathbf{B}, G) that satisfies the local IC-consistency we can reach a termination profile by asynchronous POD_F updates, then by Theorem 3.2 we can reach the same profile by means of PWOD_F^k -updates as well. \square

Theorem 4.5 generalises Proposition 1 by Bredereck and Elkind [2], which is stated for one single issue and the majority rule. Corollary 4.6 instead generalises Theorem 10 by Brill et al. [3].

4.5 Update Order Independence

While the outcome of asynchronous diffusion processes typically depends on the order of agents updating, for PWOD_F^k for $k < m$ we also need to investigate the order of updates w.r.t. the issues.

Example 4.7. Let a network and a profile of opinions be as in the figure below and let $\text{IC} = \mathcal{D} \setminus \{(111)\}$.



Agents 4 and 5 have the same initial opinions and set of influencers. If agent 4 updates in the order p, q, r , obtaining 110, and agent 5 in the order r, q, p , obtaining 011, these will be their (different) opinions in the termination profile.

As the above example shows, when agents update towards an inconsistent opinion, they might do so in radically different ways. Let us say that a profile \mathbf{B} is i -reachable from profile \mathbf{B}^0 if there exists a sequence of PWOD_F^k updates from \mathbf{B}^0 to \mathbf{B} with set of updating agents $M = \{i\}$. An i -termination profile is therefore a fixed point of any i -update. We also say that PWOD_F^k is *issue-order-independent* if for all $i \in \mathcal{N}$ and profile \mathbf{B} , there is a unique i -termination profile that is i -reachable from \mathbf{B} . We prove the following:

THEOREM 4.8. *If \mathbf{B}^0 and G have the local IC-consistency property w.r.t. to a k -geodetic IC, then PWOD_F^k is issue-order-independent.*

PROOF SKETCH. By the local IC-consistency property of \mathbf{B} and G , every influence update of an agent i is based on an IC-consistent opinion. If IC is k -geodetic, every influence update between two models must be part of an IC-consistent shortest path connecting them. To see this, observe that a k -geodetic IC either contains all models of a shortest path between two models, or does not contain any. Therefore, no matter the update order, the i -termination profile i -reachable from \mathbf{B}^0 is unique, and is such that $B_i = F(\mathbf{B}_{\text{Inff}(i)})$. \square

5 CONCLUSIONS AND FUTURE WORK

In this paper we defined a formal framework for opinion diffusion with binary issues under constraints. We proposed a setting in which agents in a social influence network update their opinions towards the aggregated opinion of their influencers, either on all issues at the same time, or on sets of issues of bounded size. We showed that if the integrity constraint satisfies a property called k -geodeticity, then the influence gap created by the constraint can be reduced by considering updates on sets of propositions of size at most k . We also investigated the termination of the associated diffusion processes, generalising several results from the literature.

We raise a number of open questions, and suggests compelling directions for future research. First, our model easily generalises to cases in which agents might be uncertain about, or abstain from giving an opinion on certain issues. Second, obtaining termination results in absence of the local IC-consistency profile, or characterising the set of constraints that guarantee termination on arbitrary networks, would be a major advancement. Last, strategic issues might be at play, motivating a deeper investigation of the incentive structure behind influence updates, especially when a collective decision is expected after the influence process.

Acknowledgments

The authors acknowledge the support of the Labex CIMI project "Social Choice on Social Networks" (ANR-11-LABX-0040-CIMI) and the ANR JCJC project SCONE (ANR 18-CE23-0009-01).

REFERENCES

- [1] Vincenzo Auletta, Diodato Ferraioli, and Gianluigi Greco. 2018. Reasoning about Consensus when Opinions Diffuse through Majority Dynamics. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*.
- [2] Robert Bredebeck and Edith Elkind. 2017. Manipulating Opinion Diffusion in Social Networks. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*.
- [3] Markus Brill, Edith Elkind, Ulle Endriss, and Umberto Grandi. 2016. Pairwise Diffusion of Preference Rankings in Social Networks. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*.
- [4] Daizhan Cheng, Zhiqiang Li, and Hongsheng Qi. 2010. A Survey on Boolean Control Networks: A State Space Approach. In *Three Decades of Progress in Control Sciences*. Springer, 121–139.
- [5] Flavio Chierichetti, Jon Kleinberg, and Sigal Oren. 2013. On Discrete Preferences and Coordination. In *Proceedings of the 14th ACM Conference on Electronic Commerce (EC-2013)*.
- [6] Zoé Christoff and Davide Grossi. 2017. Binary Aggregation with Delegable proxy: An Analysis of Liquid Democracy. In *Proceedings of the 16th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*.
- [7] Zoé Christoff and Davide Grossi. 2017. Stability in Binary Opinion Diffusion. In *Proceedings of the 6th International Conference on Logic, Rationality and Interaction (LORI)*.
- [8] Yves Crama and Peter L. Hammer. 2011. *Boolean Functions: Theory, Algorithms, and Applications*. Cambridge University Press.
- [9] Morris H. DeGroot. 1974. Reaching a consensus. *Journal of the American Statistical Association* 69, 345 (1974), 118–121.
- [10] David Easley and Jon Kleinberg. 2010. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press, New York, NY, USA.
- [11] Oya Ekin, Peter L. Hammer, and Alexander Kogan. 1999. On Connected Boolean Functions. *Discrete Applied Mathematics* 96-97 (1999), 337–362.
- [12] Edith Elkind, Piotr Faliszewski, and Arkadii Slinko. 2009. Swap bribery. In *Proceedings of International Symposium on Algorithmic Game Theory (SAGT)*. Springer, 299–310.
- [13] Ulle Endriss. 2016. Judgment Aggregation. In *Handbook of Computational Social Choice*, F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia (Eds.). Cambridge University Press, Chapter 17.
- [14] Piotr Faliszewski, Rica Gonen, Martin Koutecký, and Nimrod Talmon. 2018. Opinion Diffusion and Campaigning on Society Graphs. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*.
- [15] Michal Feldman, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. 2014. Reaching Consensus via Non-Bayesian Asynchronous Learning in Social Networks. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM)*.
- [16] Noah E. Friedkin, Anton V. Proskurnikov, Roberto Tempo, and Sergey Parsegov. 2016. Network Science on Belief System Dynamics under Logical Constraints. *Science* 354, 6310 (2016), 321–326.
- [17] Sujata Ghosh and Fernando R. Velázquez-Quesada. 2015. Agreeing to Agree: Reaching Unanimity via Preference Dynamics Based on Reliable Agents. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.
- [18] Eric Goles and J. Olivos. 1980. Periodic behaviour of generalized threshold functions. *Discrete Mathematics* 30, 2 (1980), 187 – 189.
- [19] Eric Goles and Maurice Tchuente. 1983. Iterative behaviour of generalized majority functions. *Mathematical Social Sciences* 4, 3 (1983), 197 – 204.
- [20] Umberto Grandi, Emiliano Lorini, Arianna Novaro, and Laurent Perrussel. 2017. Strategic Disclosure of Opinions on a Social Network. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.
- [21] Umberto Grandi, Emiliano Lorini, and Laurent Perrussel. 2015. Propositional Opinion Diffusion. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.
- [22] Mark Granovetter. 1978. Threshold models of collective behavior. *Amer. J. Sociology* 83, 6 (1978), 1420–1443.
- [23] Davide Grossi and Gabriella Pigozzi. 2014. *Judgment Aggregation: A Primer*. Morgan and Claypool.
- [24] Tibor Hegedüs and Nimrod Megiddo. 1996. On the Geometric Separability of Boolean Functions. *Discrete Applied Mathematics* 66, 3 (1996), 205–218.
- [25] Matthew O. Jackson. 2008. *Social and Economic Networks*. Princeton University Press.
- [26] Matthew O. Jackson and Leeat Yariv. 2011. Diffusion, Strategic Interaction, and Social Structure. In *Handbook of Social Economics*, J. Benhabib, A. Bisin, and M. O. Jackson (Eds.). Vol. 1. North-Holland, 645–678.
- [27] Stuart Kaufmann. 1969. Homeostasis and Differentiation in Random Genetic Control Networks. *Nature* 224 (1969), 177–178.
- [28] Keith Lehrer and Carl Wagner. 1981. *Rational Consensus in Science and Society*. Springer.
- [29] Michael K Miller and Daniel Osherson. 2009. Methods for distance-based judgment aggregation. *Social Choice and Welfare* 32, 4 (2009), 575–601.
- [30] Nicolas Schwind, Katsumi Inoue, Gauvain Bourgne, Sébastien Konieczny, and Pierre Marquis. 2015. Belief revision games. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*.
- [31] Nicolas Schwind, Katsumi Inoue, Gauvain Bourgne, Sébastien Konieczny, and Pierre Marquis. 2016. Is promoting beliefs useful to make them accepted in networks of agents? In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*.
- [32] Ilya Shmulevich, Edward R. Dougherty, Seungchan Kim, and Wei Zhang. 2002. Probabilistic Boolean networks: a rule-based uncertainty model for gene regulatory networks. *Bioinformatics* 18 (2002), 261–274.