Multi-Issue Opinion Diffusion under Constraints

Sirin Botan, Umberto Grandi, and Laurent Perrussel

Abstract

Most existing models of opinion diffusion on networks neglect the existence of logical constraints correlating 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. We propose a model based on individual updates on subsets of the issues at stake, overcoming the main problem of dealing with inconsistent influencing opinions. By adapting notions from the theory of boolean functions, we characterise the set of integrity constraints under which the diffusion model is closed under influence. We also identify sufficient conditions on the network topology to guarantee the termination of the diffusion process.[a]

[a] A previous version of this paper appears in the proceedings of the *4th AAMAS Workshop on Exploring Beyond the Worst Case in Computational Social Choice (EXPLORE 2017)*, with the title “Propositionwise Opinion Diffusion with Constraints”.

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 (2010) and Jackson (2008)), with applications ranging from product adoption to disaster information diffusion. 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. Two recent examples are the work of Friedkin et al. (2016) in sociological modelling—who consider how beliefs spread and change in a group—and the work by Schwind et al. (2015) in the area of belief merging.

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 jury needing to reach a decision on whether a defendant is guilty based on the validity of the evidence; or a participatory budgeting algorithm in which users decide which project to fund under a budget constraint; to 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. Let us showcase the main problems tackled by our paper with a concrete example.

Example 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. While the first three agents are rather certain of their view, 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 an 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.*

We argue that information should not always spread by looking at the entire set of issues. If the fourth agent in the example above consulted her influencers on one single issue, such as: “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.

[1] Corresponding to a simple threshold model (Granovetter, 1978).

The main problem tackled in this paper is to identify the minimal amount of information exchange—in terms of the 'scope' of the questions asked by agents to influencers—that allows an information diffusion system to work properly given a certain integrity constraint. We 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 boolean functions. We also 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.

The paper is organised as follows. In Section 2 we define our model of propositionwise opinion diffusion under constraints. Section 3 introduces and studies a useful class of integrity constraints, which is used in Section 4 to obtain our main results. Section 5 identifies networks on which the termination of the diffusion system is guaranteed, and Section 6 concludes.

Related work

Diffusion on networks has been extensively studied in the field of social network analysis, be it diffusion of diseases, information, or opinions (Jackson and Yariv, 2011; Easley and Kleinberg, 2010; Shakarian et al., 2015). Building on the classical work of Granovetter (1978), DeGroot (1974), and Lehrer and Wagner (1981), a number of models were recently introduced for the diffusion of *complex* opinions, such as knowledge bases (Schwind et al., 2015, 2016), preferences over alternatives (Ghosh and Velázquez-Quesada, 2015; Brill et al., 2016; Bredereck and Elkind, 2017), and binary evaluations over multiple issues (Grandi et al., 2015, 2017; Christoff and Grossi, 2017a). Our paper builds on the latter model, including an integrity constraint that logically correlates the issues at stake. To the best of our knowledge, the only work in opinion diffusion under constraints is the recent work of Friedkin et al. (2016), which however represents opinions as real-valued beliefs, as well as the work of Christoff and Grossi (2017b). Let us also mention the literature on boolean networks (Kaufmann, 1969), which is used for modeling biological regulatory networks (see, e.g., Shmulevich et al. (2002)), and focuses on updates on one single binary issue. To the best of our knowledge, multiple issues and constraints have never been considered in this literature.

2 The General Framework

This section presents our diffusion model for binary opinions over multiple issues correlated by an integrity constraint.

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 $\boldsymbol{B} = (B_1, \dots, B_n)$ 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 $\boldsymbol{B}$. Thus if $B = (110)$, then $B(p) = B(q) = 1$ and $B(r) = 0$.

An *integrity constraint* $\mathrm{IC} \subseteq \mathcal{D}$ defines a domain of feasible opinions. We say that B is IC-consistent when $B \in \mathrm{IC}$, or equivalently that B is a model of IC. For each agent i, we assume that $B_i \in \mathrm{IC}$, 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 $\mathrm{IC} = \{(110), (011), (101), (100), (010), (001), (000)\}$. In further sections we will 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 $\mathit{Inf}(i)_G = \{j \in \mathcal{N} \mid (j, i) \in E\}$ is the set of influencers of agent i in the network G.[2]

We model social influence as a transformation function, which takes as input a profile of IC-consistent opinions $\boldsymbol{B} = (B_1, \ldots, B_n)$, and returns a set of profiles which are each the result of some opinion update on $\boldsymbol{B}$. If clear from the context, we omit reference to G and IC.

Let $F = (F_1, \ldots F_n)$ be composed of *aggregation procedures* $F_i : \mathrm{IC}^{\mathit{Inf}(i)} \to \mathcal{D}$, one for each agent i. We assume that aggregation functions satisfy the minimal requirement of *unanimity*, i.e., whenever $B_j = B^*$ for all $j \in \mathit{Inf}(i)$ then $F_i(\boldsymbol{B}) = B^*$. 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 Endriss (2016); Grossi and Pigozzi (2014).

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 is a set of issues, let $(B\restriction_{\mathcal{I}\setminus S}, B'\restriction_S)$ be the opinion obtained from B with the opinions on the issues in S replaced by those in B'.

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

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

In this paper 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. Our opinion diffusion model is hence defined as follows.

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

$$\begin{aligned} \mathrm{PWOD}^k_F(\boldsymbol{B}) =\{&\boldsymbol{B}' \mid \exists M \subseteq \mathcal{N}, S : M \to 2^{\mathcal{I}} \textit{ with } |S(i)| \leq k, \\ &\textit{s.t. } B'_i = F\text{-UPD}(\boldsymbol{B}, i, S(i)) \textit{ for } i \in M \\ &\textit{and } B'_i = B_i \textit{ otherwise.}\} \end{aligned}$$

PWOD^k_F defines, for each consistent profile of opinions $\boldsymbol{B}$, 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. To obtain the more classical view of diffusion as a discrete time iterative process, it is sufficient to combine PWOD^k_F with an agent-scheduler—i.e., a turn-taking function—and an issue-scheduler deciding which issues are updated by each agent.

Example 2 (Pairwise preference diffusion)**.** *The framework of pairwise preference diffusion by Brill et al. (2016) can be seen as an instance of* PWOD^1_F *where F is the (strict) majority rule. To see this, consider a set A of alternatives. A linear order $\succ$ is an irreflexive, transitive and complete binary relation over A, which can be represented as a binary evaluation over a set of issues $\mathcal{I}_A = \{p_{ab} \mid (a, b) \in A \times A \text{ and } a \neq b\}$, such that $B(p_{ab}) = 1$ if and only if $a \succ b$.*[3] *The integrity constraint* $\mathrm{IC}_\succ$ *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. (2016) propose to update on one pair of alternatives at the time, which corresponds to a propositionwise update on the analogous issue.*

[2]Observe that we do not make any assumption on whether $i \in \mathit{Inf}(i)$, thereby defining the framework in full generality.

[3]Representing preferences with binary evaluations is an idea that can be traced back to the work of Wilson (1975).

As a last definition, we introduce the following concept:

Definition 2. *A profile $\boldsymbol{B}$ is a termination profile for* PWOD_F^k *and* IC *if* $\text{PWOD}_F^k(\boldsymbol{B}) = \{\boldsymbol{B}\}$.

Termination profiles are fixed points of PWOD_F^k. We stress the role of IC in determining which updates can be performed. To clarify our definitions, consider the following:

Example 3. *Consider a scenario similar to Example 1. Three agents are voting on three proposals for their city; a skyscraper, (S) an hospital (H), and a new road (R). Recall that the constraint in this setting is $(S \wedge H \to R)$. The three agents are connected in the following network, where the initial profile is* $\boldsymbol{B} = (111, 011, 101)$.

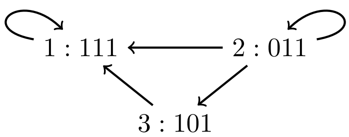

Assume that F_i is the strict majority rule for each i, accepting an issue only if a strict majority of their influencers accept it. If all agents update simultaneously under PWOD_F^2*, the resulting profile at termination (where each agent updates first on $\{S, H\}$, then $\{H, R\}$) will be* $(011, 011, 011)$. *If the agents update simultaneously under* PWOD_F^1 *and we assume the same order on issues for all agents —first on the first issue, then the second, and so on—we will reach the same termination profile—(011, 011, 011), after four rounds.*

3 Geodetic Integrity Constraints

In this section we build on notions from the theory of boolean functions (see, e.g., Crama and Hammer (2011)) to identify a useful class of integrity constraints that we will later use to characterise termination profiles of our diffusion model.

3.1 Basic Definitions

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 $H(B, B') = \Sigma_{p \in \mathcal{I}} |B(p) - B'(p)|$. Consider the following:

Definition 3. *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 $B \in$* IC*,*
- (ii) *the set of edges E_{IC}^k is defined as follows: $(B_1, B_2) \in E_{\text{IC}}^k$ iff $H(B_1, B_2) \leq k$, for any* $B_1, B_2 \in \text{IC}$.

Intuitively, the k-graph of IC connects two models if one can be reached from the other by swapping at most k issues. As it is clear from Definition 3, $\mathcal{G}_{\text{IC}}^k \subseteq \mathcal{G}_{\mathcal{D}}^k$ 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 4. *An integrity constraint* IC *is k-geodetic if and only if for all B_1 and B_2 in* IC*, at least one of the shortest paths from B_1 to B_2 in $\mathcal{G}_{\mathcal{D}}^k$ is also a path of $\mathcal{G}_{\text{IC}}^k$.*

For ease of notation, we denote 1-geodeticness with geodeticness *tout court*, borrowing the term from the equivalent definition for boolean functions Ekin et al. (1999). To illustrate our definitions, consider the following example.

Example 4. *Let there be three issues, and let* IC $= \{(000), (001), (010), (100), (011), (111)\}$*. The graph below corresponds to* $\mathcal{G}^1_{\text{IC}}$*, connecting only those models that satisfy* IC *with a continuous edge. The graph consisting of all edges (continuous and dashed) corresponds to* $\mathcal{G}^1_{\mathcal{D}}$*.*

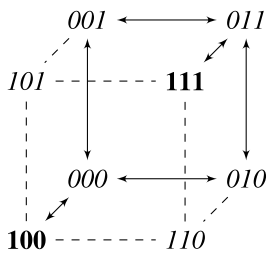

We can now observe that IC *is not geodetic: the shortest paths between* (100) *and* (111) *in* $\mathcal{G}^1_{\mathcal{D}}$ *pass through either* (110) *or* (101)*, which however are not nodes of* $\mathcal{G}^1_{\text{IC}}$*. As for 2-geodeticness (and similarly for* k*-geodeticness for higher* k*), it is easy to see that* IC *satisfies it, since there is a direct path from* (100) *to* (111) *in* $\mathcal{G}^2_{\text{IC}}$*.*

3.2 Recognising 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). To see this, let $\prec$ and $\prec'$ 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 $a \succ_i c \succ_i b$ or $b \succ_i c \succ_i a$.[4] Knowing this, it becomes straightforward to show that $\text{IC}_{\succ}$ is geodetic (for the particular encoding of preferences explained in Example 2). Similar encodings can be used to show that partial and weak orders and equivalence relations can be modelled by geodetic constraints.

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 $(p \vee q) \wedge (\neg p \vee \neg q)$, as shown in Example 5.

Example 5. *Let there be two issues, and consider the following* IC $= \{(01), (10)\}$*, i.e.,* p *XOR* q*. The graphs* $\mathcal{G}^1_{\text{IC}}$ *(boldface nodes) and* $\mathcal{G}^1_{\mathcal{D}}$ *(all nodes and edges) can be represented as follows:*

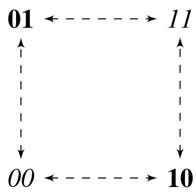

Clearly, IC *is not geodetic since the two nodes* (01) *and* (10) *are disconnected in* $\mathcal{G}^1_{\text{IC}}$ *but connected in* $\mathcal{G}^1_{\mathcal{D}}$*.*

Another interesting class is that of *budget constraints*, which specify the list of subsets of the issues $\mathcal{I}$ that 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. A number of logical characterisation of 1-geodetical integrity constraints can be found in the work of Ekin et al. (1999), including the fact that negative formulas are 1-geodetic. To the best of our knowledge, for

[4]This result is folklore, a formal proof can be found in in Elkind et al. (2009).

k-geodetic constraints no such characterisation is available. While similar results would be outside the scope of this paper, we show the following simple facts:

Fact 1. *If $|\mathcal{I}| = m$, then for all $k \geq m$ any* IC *is k-geodetic.*

This is straightforward, since all nodes of G^k_{IC} are directly connected if $k \geq m$. Observe also the following:

Fact 2. *If* IC *is k-geodetic for a set of issues $\mathcal{I}$, 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-geodeticness of a constraint, in the following:

Lemma 1. *An integrity constraint* IC *is k-geodetic iff for all models $B_1, B_2 \in$ IC, there is a path in G^k_{IC} from B_1 to B_2 of lenght smaller than $\left\lceil \frac{H(B_1,B_2)}{k} \right\rceil$.*

Proof sketch. Let B_1 and B_2 be two models of IC. The length of the shortest path from B_1 to B_2 in the hypercube $\mathcal{G}^k_{\mathcal{D}}$ is exactly $\left\lceil \frac{H(B_1,B_2)}{k} \right\rceil$, since $H(B_1, B_2)$ is the number of issues that has to be changed to move from B_1 to B_2, and the edges in $\mathcal{G}^k_{\mathcal{D}}$ change k symbols at most. As $\mathcal{G}^k_{\text{IC}} \subseteq \mathcal{G}^k_{\mathcal{D}}$, if there is a path of minimal length connecting B_1 to B_2 in $\mathcal{G}^k_{\text{IC}}$, then it is one of the shortest paths of $\mathcal{G}^k_{\mathcal{D}}$. By repeating for all B_1 and B_2 in IC we obtain the desired statement. □

4 Termination Profiles

In this section we investigate how the structure of the integrity constraint influences the set of PWOD^k_F termination profiles.

4.1 Influence-Closure of PWOD^k_F

As observed in the introduction, when limiting the influence updates to sets of k issues, the influence process may be blocked by the structure of the integrity constraint at hand. We therefore give the following definition:

Definition 5. PWOD^k_F *is influence-closed wrt. an integrity constraint* IC *if for any termination profile $\boldsymbol{B}$, and any $i \in \mathcal{N}$, we have that if $F(\boldsymbol{B}_{\mathit{Inf}(i)}) \in$ IC, then $B_i = F(\boldsymbol{B}_{\mathit{Inf}(i)})$.*

Influence closure of PWOD^k_F simply means that whenever possible, an agent will move towards, and eventually adopt the aggregate opinion of her influencers. Clearly, if $k = |\mathcal{I}|$ then PWOD^k_F is influence-closed, irrespective of the constraint, since agents update on all issues at the same time; we now give exact bounds on the integrity constraints and the degree of issue-wise communication for this to happen:

Theorem 1. PWOD^k_F *is influence-closed with respect to* IC *if and only if* IC *is k-geodetic.*

Proof. For the right to left direction, assume that IC is k-geodetic. Suppose PWOD^k_F terminates on a profile $\boldsymbol{B}$ and, by contradiction, that there exists an agent i such that $F(\boldsymbol{B}_{\mathit{Inf}(i)}) \in$ IC and $B_i \neq F_i(\boldsymbol{B}_{\mathit{Inf}(i)})$. Since IC is k-geodetic, and both B_i and $F(\boldsymbol{B}_{\mathit{Inf}(i)})$ are IC-consistent, by Definition 4 the shortest path in $\mathcal{G}^k_{\mathcal{D}}$ between them is composed of IC-consistent opinions. Let B^1 be the first model on such path after B_i, and let $p_1, \ldots, p_\ell$ be the issues on which B_i and B^1 differ. By the definition of $\mathcal{G}^k_{\mathcal{D}}$ we know that $\ell \leq k$. Moreover, since B^1 is on the shortest path between B_i and $F(\boldsymbol{B}_{\mathit{Inf}(i)})$, we can infer that $B^1 = (B_i{\restriction_{\mathcal{I}\setminus S}}, F_i(\boldsymbol{B}_{\mathit{Inf}(i)}){\restriction_S})$, where $S = \{p_1, \ldots, p_\ell\}$. If we now consider profile $\boldsymbol{B}'$, obtained by setting $I = \{i\}$ and S as defined above in Definition 1, we obtain that $\boldsymbol{B}' \in \text{PWOD}^k_F(\boldsymbol{B})$, against the assumption that $\boldsymbol{B}$ is a termination profile.

For the left to right direction, suppose IC is not k-geodetic. This implies the existence of two IC-consistent opinions B_1 and B_2 that are not connected in $\mathcal{G}^k_{\mathcal{D}}$ by any shortest path. We now construct an instance of PWOD^k_F that is not influence-closed. Let there be two agents, and the network G be such that $E = \{(1, 2)\}$. Let F be the strict majority rule. Profile $\boldsymbol{B} = (B_1, B_2)$ is a termination profile of PWOD^k_F, in which however $B_2 \neq B_1 = F(\boldsymbol{B}_{\mathit{Inf}(2)})$, showing that PWOD^k_F is not influence-closed wrt to IC. □

Theorem 1 shows that if aggregating an agent's influencers using F gives an opinion in the k-geodetic set IC, PWOD^k_F will eventually reach a state where each agent's opinion equals the aggregated opinion of their respective influencers.

4.2 Update Order Independence

While the outcome of PWOD^k_F in most cases depends on the order in which agents update on the network, we are also interested in the order in which each agent updates their opinion wrt. the issues; when $k < |\mathcal{I}|$, this order matters in determining possible termination profiles. Consider the following example:

Example 6. *Let a network and a profile of opinions be as described in the figure below, and let the integrity constraint* IC $= \mathcal{D} \setminus \{(111)\}$.

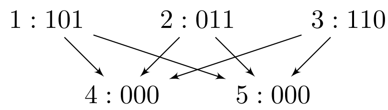

Agents 4 and 5 have the same initial opinions and set of influencers. If agent 4 updates in the order p, q, r, obtaining 110, and agent 5 in the order r, q, p, obtaining 011*, these will be their – different – opinions in the termination profile.*

As the above example shows; when agents update towards an inconsistent opinion, they might do so in radically different ways. This is however not possible when the aggregated opinion is IC-consistent for a geodetic IC.

Definition 6. *A pair $(\boldsymbol{B}^0, G)$, where $\boldsymbol{B}^0$ is a profile and G a network, has the* local IC-consistency property *if for all profiles $\boldsymbol{B}$ reachable from $\boldsymbol{B}^0$ and each $i \in \mathcal{N}$ we have that $F(\boldsymbol{B}_{\mathit{Inf}(i)})$ is* IC-*consistent.*

We also say that a profile $\boldsymbol{B}$ is i-reachable from profile $\boldsymbol{B}^0$ if there exists a sequence of PWOD^k_F updates from $\boldsymbol{B}^0$ to $\boldsymbol{B}$ with set of updating agents $M = \{i\}$. An i-termination profile is therefore a fixed point of any i-update. PWOD^k_F is *issue-order-independent* if for all $i \in \mathcal{N}$ and profile $\boldsymbol{B}$, there is a unique i-termination profile i-reachable from $\boldsymbol{B}$. We can now prove the following:

Theorem 2. *If $\boldsymbol{B}^0$ and G have the local* IC*-consistency property wrt to a k-geodetic* IC*, then* PWOD^k_F *is issue-order-independent.*

Proof sketch. By the local IC-consistency property of $\boldsymbol{B}$ 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 $\boldsymbol{B}^0$ is unique, and is such that $B_i = F(\boldsymbol{B}_{\mathit{Inf}(i)})$. □

In particular, Theorem 2 applies to trees, simple cycles, and any network in which the in-degree of the nodes is at most one, showing that PWOD^k_F is issue-order-independent on these classes of networks.

4.3 Inconsistent Profiles

Propositionwise updates were introduced to get as close as possible to an *inconsistent* aggregated opinion. We provide here a formal justification for this claim. As a measure of closeness between opinions we use the Hamming distance (see Section 3 for a definition). We show the following:

Theorem 3. *Let* IC *be k-geodetic, and let $\boldsymbol{B}$ be a termination profile of* PWOD_F^k. *Then, $\boldsymbol{B}$ is also a termination profile for* PWOD_F^K *for any $K \geq k$.*

Proof. Let $i \in \mathcal{N}$. If $F(\boldsymbol{B}_{\mathit{Inf}(i)}) \in \text{IC}$, then by Theorem 1, since $\boldsymbol{B}$ is a termination profile, we have that $B_i = F(\boldsymbol{B}_{\mathit{Inf}(i)})$ and no further update of any size is possible. Assume then that $F(\boldsymbol{B}_{\mathit{Inf}(i)}) \notin \text{IC}$ and, for the sake of contradiction, that there exists a possible update of PWOD_F^K for a specific $K \geq k$ resulting in a different $B_i' \neq B_i$. That is, there exists a set of issues S with $|S| > k$ such that $F\text{-UPD}(\boldsymbol{B}, i, S) = B_i'$. Since $B_i' \in \text{IC}$, and IC is k-geodetic, there is a path of updates of size k or smaller that reaches B_i' from B_i. This implies that S can be partitioned in smaller subsets of less than k issues reaching B_i' from B_i, against the assumption that $\boldsymbol{B}$ is a termination profile. □

As a consequence of Theorems 1 and 3, we obtain that to identify the correct level of communication between agents it is sufficient to identify the minimal k such that the integrity constraint is k-geodetic. Larger degrees of communication would be costly and useless, and smaller would not allow to reach consistent opinions.

However, the following result shows that when faced with an outcome that does not satisfy the constraint, it is possible to build examples in which individual opinions are as far as possible from the aggregated opinion of their influencers:

Theorem 4. *For any finite n and $m > 3$, there is a geodetic* IC *over m issues, a network G over n agents, and a termination profile $\boldsymbol{B}$ for* PWOD_F^1*, such that there is one agent i with $H(B_i, F(\boldsymbol{B}_{\mathit{Inf}(i)})) \in O(m)$.*

Proof sketch. Let there be $m > 3$ issues and n agents. First, let $\text{IC} = (\neg p_1 \wedge \neg p_m) \rightarrow (p_2 \wedge, \ldots, \wedge p_{m-1})$, i.e. IC allows all opinions except those which reject the first and last issue and at least one other issue. Second, let the network G be the following simple directed acyclic graph $E = \{(i, n) \mid 1 \leq i \leq n-1\}$, i.e. the first $n-1$ agents are the influencers of the n-th agent. Third, let $\boldsymbol{B}$ be such that $B_n = (0, 1, \ldots, 1, 0)$, and B_i for $i < n$ be such that $B_i(i) = 1$ and $B_i(j) = 0$ otherwise. If F is the strict majority rule, then it is easy to see that $F(\boldsymbol{B}_{\mathit{Inf}(i)}) = (0, \ldots, 0)$, which is not IC-consistent.

Clearly, $H(B_n, F(\boldsymbol{B}_{\mathit{Inf}(n)}))=m-2$. Even worse, we can observe that $\sum_{j \neq n} H(B_n, B_j)=(n-1)\times(m-2)$. Profile $\boldsymbol{B}$ is also a termination profile. Agent n, which is the only influenced agent, cannot move towards the aggregated opinion $(0, \ldots, 0)$ by any propositionwise update. It remains to be shown that IC is geodetic. Let B, B' be models of IC. If both B and B' accept the first (last) issue, since all opinions accepting the first (last) issue satisfy IC then we can move between the two with one-issue updates. If both B and B' reject both the first and the last issue, then $B = B' = (0, 1, \ldots, 1, 0)$ as this is the only IC-consistent opinion. A simple case study concludes the proof, showing that IC is geodetic.

□

4.4 Computational Complexity

As a consequence of Theorems 1 and 3, if a mechanism designer faces a situation described by an integrity constraint IC, it should allow communication on the network on up to k issues, where k is the smallest number such that IC is k-geodetic. We now investigate the computational complexity of this task.

Theorem 5. *Let* IC *be a constraint over* m *issues and* $k < m$*. Checking whether* IC *is* k*-geodetic is* co-NP*-complete.*

Proof sketch. To find a counterexample for k-geodeticness, it is sufficient to find two models B_1 and B_2 of IC that are not connected by any of the shortest paths of $\mathcal{G}^k_{\mathcal{D}}$. A co-NP algorithm guesses two opinions B_1 and B_2, checks that B_1 and B_2 are IC-consistent, and that for all subsets $S \in \mathcal{I}$ of $|S| \leq k$ we have that $(B_1\restriction_{\mathcal{I}\setminus S}, B_2\restriction_S) \not\models$ IC, showing a counterexample to the k-geodeticness of IC. Note that the number of subsets of size k is an exponential figure in k but not in m, which is the input size.

As for hardness, we exploit a result by Hegedüs and Megiddo (1996), 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 means (a) observing that 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 IC $\wedge\, p$ and IC $\wedge \neg p$ must also be k-geodetic for all $p \in \mathcal{I}$. To show (c), suppose that B_1 and B_2 are two models of IC $\wedge\, p$ that are not connected by any shortest path of $\mathcal{G}^k_{\mathcal{D}}$. Since B_1 and B_2 are also models of IC, and $\mathcal{G}^k_{\text{IC}\wedge p} \subseteq \mathcal{G}^k_{\text{IC}}$, this would imply that IC is not k-geodetic, against the assumption. ☐

The hardness result above is shown for 1-geodetic formulas by Ekin et al. (1999). By using the algorithm of Theorem 5 as an oracle, with binary search we obtain the following:

Theorem 6. *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^p_2.

Putting together the previous result with Theorems 1 and 3, we obtain the following operational result for PWOD^k_F:

Corollary 1. *Let* IC *be an integrity constraint over* m *issues and let* $k < m$*. Checking whether* k *is the minimal* $k < m$ *such that* PWOD^k_F *is influence-closed is in* Θ^p_2.

5 Termination of the Iterative Process

In this section we analyse the termination of discrete-time iterative processes that are defined by PWOD^k_F updates.

5.1 Basic Definitions

Recall our Definition 1, introducing propositionwise opinion diffusion as a transformation function that associates a set of updated profiles with every IC-consistent profile. Thus, PWOD^k_F 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 sets of issues $S(i)$, one for each updating individual. Termination states, as defined by our Definition 2, are the attractors of such a transition system.

In line with the existing literature on propositional opinion diffusion Grandi et al. (2015); Brill et al. (2016); Bredereck and Elkind (2017) and on boolean networks Kaufmann (1969), we define *asynchronous* PWOD^k_F by restricting transitions to those involving only one single agent at a time, and *synchronous* PWOD^k_F by restricting transitions to those involving all individuals. We call a transition from $\boldsymbol{B}$ to $\boldsymbol{B}'$ *effective* if $\boldsymbol{B}' \neq \boldsymbol{B}$. We say that PWOD^k_F *terminates universally* if there exists no infinite sequence of effective transitions, while it *terminates asymptotically* if from any IC-consistent profile there is a sequence of transitions that reaches a termination profile. Finally, a *consensual termination profile* is a termination profile $\boldsymbol{B}$ such that for all $i, j \in \mathcal{N}$ we have that $B_i = B_j$.

5.2 Simple Cycles

A simple cycle is a finite connected network E such that every agent has exactly one outgoing edge and exactly one incoming edge.

Theorem 7. *If G is a simple cycle and* IC *is k-geodetic, then asynchronous* PWOD_F^k *terminate asymptotically to a consensual termination profile.*

Proof sketch. Let $\boldsymbol{B}^0$ be a profile on G, and let $i^* \in \mathcal{N}$ be such that $B_{i^*}^0 \neq B_{i^*+1}^0$. Since IC is k-geodesic, by Lemma 1 there is a sequence of propositionwise updates of length $k = \lceil \frac{H(B_{i^*}^0, B_{i^*+1}^0)}{k} \rceil$ that transforms the latter opinion into the former. By having agent $i^* + 1$ updating at time $t = 0, \ldots, k$, and $S(i)$ according to the sequence of updates above, we obtain a resulting profile $\boldsymbol{B}^k$ such that $B_{i^*+1}^k = B_{i^*}^0$ and $B_j^k = B_j^0$ for all $j \neq i^* + 1$. Repeat the process for $i^* + 2$, and continue on the cycle until agent i^*, obtaining a consensual termination profile in which all opinions are $B_{i^*}^0$. □

Observe that the termination profiles reachable from the same initial profile on a cycle can depend on the sequence of updates. A characterisation of such a set is an interesting open problem, as already observed by Brill et al. (2016).

5.3 DAGs and Complete Graphs

A directed acyclic graph (DAG) is a directed graph that contains no cycle involving two or more vertices. A simple argument of propagation allows us to prove the following:

Theorem 8. *If G is a DAG, then both synchronous and asynchronous* PWOD_F^k *terminate universally.*

Proof sketch. We define potential functions h_i for each node i, as follows: $h_i(t) = H(B_i^t, F_i(\boldsymbol{B}_{\mathit{Inf}(i)}^t))$, measuring the distance between an individual's opinion and the aggregated opinion of its influencers in profile $\boldsymbol{B}^t$. Each PWOD_F^k update decreases one or more such functions, those of the updating agents, and possibly increases others, those of the agents influenced by the one updating. By ordering such potential functions based on the distance from a node to a source, which is possible given that G is a DAG, we obtain a lexicographic ordering of all functions h_i that decreases strictly with each effective transition. It is therefore impossible to build an infinite sequence of PWOD_F^k effective transitions. □

Let a complete graph be a graph $G = (\mathcal{N}, E)$ where $E = \mathcal{N} \times \mathcal{N}$. With a similar argument as the one used in the previous proof (and generalising a result by Farnoud et al. (2013)) we show that:

Theorem 9. *If G is the complete graph, then both synchronous and asynchronous* PWOD_F^k *converge universally.*

Proof sketch. On a complete graph the set of influencers $\mathit{Inf}(i) = \mathcal{N}$ for all i. Let therefore $h(t) = \sum_i H(B_i, F(\boldsymbol{B}))$ be a potential function measuring the overall distance of individual opinions from the aggregated one. Every effective transition for both PWOD_F^k decreases the value of h. □

A general result on the asymptotic convergence of PWOD_F^k is an interesting open problem. A proof similar to the one used by Brill et al. (2016) could be adapted to show that PWOD_F^k asymptotically converges on any graph, under the local IC-consistency property introduced in Definition 6, for a k-geodetic IC. Universal convergence cannot be guaranteed even on simple cycles, at least when more than two issues are present. To see this it is sufficient to consider a simple cycle with only one agent having opinion 11 and all others 00, and devise a sequence of updates that make the 11 opinion turn in the cycle whilst keeping all other opinions at 00. Termination results are well-established for boolean networks, which however consider the diffusion of a single binary issue and do not consider integrity constraints (see, for a survey, Cheng et al. (2010)).

6 Conclusion and Future Work

In this paper we defined the first formal framework of opinion diffusion with binary issues under constraints. We proposed a setting in which agents in a social influence network change their opinions by asking about the opinions of their influencer on sets of issues of bounded size. We identified the relation between the structure of the integrity constraint and the minimal size of communication sets that allows the influence to lead to changing opinions, keeping the integrity constraint satisfied. We also analysed the computational complexity of recognising k-geodetic integrity constraints and identifying the minimal k for which a constraint is k-geodetic, and investigated the termination of the associated diffusion process.

This paper raises a number of open questions, and suggests compelling directions for future research. First, observe that our model easily generalises to cases in which agents might be uncertain about, or abstain from giving an opinion on certain issues; it would be sufficient to change the aggregation procedures to accommodate such input. Second, obtaining termination results for arbitrary constraints, or characterising the set of constraints that guarantee termination on arbitrary networks, would be a major advancement. Last, strategic issues are at play, motivating a deeper investigation of the incentive structure behind influence updates, especially when a collective decision is expected after the influence process.

References

R. Bredereck and E. Elkind. Manipulating opinion diffusion in social networks. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 2017.

M. Brill, E. Elkind, U. Endriss, and U. Grandi. Pairwise diffusion of preference rankings in social networks. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.

D. Cheng, Z. Li, and H. Qi. A survey on boolean control networks: A state space approach. In *Three Decades of Progress in Control Sciences*, pages 121–139. Springer, 2010.

Z. Christoff and D. Grossi. Stability in binary opinion diffusion. In *Proceedings of the 6th International Workshop on Logic, Rationality, and Interaction (LORI)*, 2017a.

Z. Christoff and D. Grossi. Binary voting with delegable proxy: An analysis of liquid democracy. In *Proceedings of the 16th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, 2017b.

Y. Crama and P. L. Hammer. *Boolean Functions: Theory, Algorithms, and Applications*. Cambridge University Press, 2011.

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

D. Easley and J. Kleinberg. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press, New York, NY, USA, 2010. ISBN 0521195330, 9780521195331.

O. Ekin, P. L. Hammer, and A. Kogan. On connected boolean functions. *Discrete Applied Mathematics*, 96-97:337–362, 1999.

E. Elkind, P. Faliszewski, and A. Slinko. Swap bribery. In *Proceedings of International Symposium on Algorithmic Game Theory*, pages 299–310. Springer, 2009.

U. Endriss. Judgment aggregation. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 17. Cambridge University Press, 2016.

F. Farnoud, E. Yaakobi, B. Touri, O. Milenkovic, and J. Bruck. Building consensus via iterative voting. In *Proceedings of the 2013 IEEE International Symposium on Information Theory*, 2013.

N. E. Friedkin, A. V. Proskurnikov, R. Tempo, and S. Parsegov. Network science on belief system dynamics under logical constraints. *Science*, 354(6310):321–326, 2016.

S. Ghosh and F. R. Velázquez-Quesada. Agreeing to agree: Reaching unanimity via preference dynamics based on reliable agents. In *Proceedings of the 14th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2015.

U. Grandi, E. Lorini, and L. Perrussel. Propositional opinion diffusion. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2015.

U. Grandi, E. Lorini, A. Novaro, and L. Perrussel. Strategic disclosure of opinions on a social network. In *Proceedings of the 16th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2017)*, 2017.

M. Granovetter. Threshold models of collective behavior. *American Journal of Sociology*, 83(6): 1420–1443, 1978.

D. Grossi and G. Pigozzi. *Judgment Aggregation: A Primer*. Morgan and Claypool, 2014.

T. Hegedüs and N. Megiddo. On the geometric separability of boolean functions. *Discrete Applied Mathematics*, 66(3):205–218, 1996.

M. O. Jackson. *Social and Economic Networks*. Princeton University Press, 2008.

M. O. Jackson and L. Yariv. Diffusion, strategic interaction, and social structure. In J. Benhabib, A. Bisin, and M. O. Jackson, editors, *Handbook of Social Economics*, volume 1, pages 645–678. North-Holland, 2011.

S. Kaufmann. Homeostasis and differentiation in random genetic control networks. *Nature*, 224: 177–178, 1969.

K. Lehrer and C. Wagner. *Rational Consensus in Science and Society*. Springer, 1981.

N. Schwind, K. Inoue, G. Bourgne, S. Konieczny, and P. Marquis. Belief revision games. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, 2015.

N. Schwind, K. Inoue, G. Bourgne, S. Konieczny, and P. Marquis. Is promoting beliefs useful to make them accepted in networks of agents? In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.

P. Shakarian, A. Bhatnagar, A. Aleali, E. Shaabani, and R. Guo. *Diffusion in Social Networks*. Springer, 2015.

I. Shmulevich, E. R. Dougherty, S. Kim, and W. Zhang. Probabilistic boolean networks: a rule-based uncertainty model for gene regulatory networks. *Bioinformatics*, 18:261–274, 2002.

R. Wilson. On the theory of aggregation. *Journal of Economic Theory*, 10(1):89–99, 1975.

Sirin Botan
ILLC, University of Amsterdam
Amsterdam, The Netherlands
Email: sirin.botan@gmail.com

Umberto Grandi
IRIT, University of Toulouse
Toulouse, France
Email: umberto.grandi@irit.fr

Laurent Perrussel
IRIT, University of Toulouse
Toulouse, France
Email: laurent.perrussel@irit.fr