

Space-time budget allocation policy design for viral marketing [☆]

I.C. Morărescu^{a,*}, V.S. Varma^a, L. Buşoniu^b, S. Lasaulce^a

^a*Université de Lorraine, CNRS, CRAN, F-54000 Nancy, France.*

^b*Technical University of Cluj–Napoca, Romania*

Abstract

We address formally the problem of opinion dynamics when the agents of a social network (e.g., consumers) are not only influenced by their neighbors but also by an external influential entity referred to as a marketer. The influential entity tries to sway the overall opinion as close as possible to a desired opinion by using a specific influence budget. We assume that the exogenous influences of the entity happen during discrete-time advertising campaigns; consequently, the overall closed-loop opinion dynamics becomes a linear-impulsive (hybrid) one. The main technical issue addressed is finding how the marketer should allocate its budget over time (through marketing campaigns) and over space (among the agents) such that the agents' opinion be as close as possible to the desired opinion. Our main results show that the marketer has to prioritize certain agents over others based on their initial condition, their influence power in the social graph and the size of the cluster they belong to. The corresponding space-time allocation problem is formulated and solved for several special cases of practical interest. Valuable insights can be extracted from our analysis. For instance, for most cases, we prove that the marketer has an interest in investing most of its budget at the beginning of the process and that budget should be shared among agents according to the famous water-filling allocation rule. Numerical examples illustrate the analysis.

Keywords: Social networks, hybrid systems.

1. Introduction

Opinion dynamics in social networks has become a problem of increasing research interest during the last decades. This can be explained by the multiplication of digital social networks that allow a faster and more persistent influence of opinions. In this context, governmental institution and private companies use marketing over social networks as a key tool for promoting their products or ideas. However, to the best of our knowledge, there is no formal analysis pointing out the improvements that can be achieved by using the network topology in the design of the marketing strategy. Indeed, most of the existing studies focus on the analysis of models without control, *i.e.* they study the convergence, dynamical patterns or asymptotic configurations of the open-loop opinion dynamics. Various mathematical models [1, 2, 3, 4, 5, 6] have been proposed to capture different features of these complex dynamics. Empirical models based on in vitro and in vivo experiments have also been developed [7, 8, 9].

[☆]This work has been partially funded by the CNRS PEPS project YPSOC

*Corresponding author

Email addresses: constantin.morarescu@univ-lorraine.fr (I.C. Morărescu),
vineeth.satheeskumar-varma@univ-lorraine.fr (V.S. Varma), lucian@busoniu.net (L. Buşoniu),
samson.lasaulce@univ-lorraine.fr (S. Lasaulce)

One controversial problem is related to emergence of consensus in social networks. Social studies pointed out that, in general, opinions tend to converge one toward another during interactions. Therefore, is not surprising that consensus received a particular attention in opinion dynamics literature [10, 11]. While some mathematical models naturally lead to consensus [1, 2], others lead to network clustering [4, 5, 12]. In order to enforce consensus, some recent studies propose the control of one or a few agents, see [13, 14]. Besides these methods of controlling opinion dynamics towards consensus, we also find recent attempts to control the discrete-time dynamics of opinions such that as many agents as possible reach a certain set after a finite number of influences [15]. Another relatively new line of research is based on the change of opinions through the change of susceptibility and resistance parameters [16, 17, 18]. Basically, each individual is characterized by certain parameters that make it more or less easy to influence. In [18] (and some reference therein) the authors zoom in the model and see how the persuasion can be realized by acting on the susceptibility of individuals. Instead, we are looking directly at the outcome of the persuasion strategy and use this information in the long term evolution of opinion dynamics.

Viral marketing refers to the practice where a seller attempts to artificially create word-of-mouth advertising among potential customers, and the effectiveness of this trend has been well established by social scientists and economists [19, 20]. In [21], the authors consider multiple influential entities competing to control the opinion of consumers under a game theoretical setting. However, this work assumes an undirected graph and a voter model for opinion dynamics resulting in strategies that are independent of the node centrality. On the other hand, [22] considers a similar competition with opinion dynamics over a directed graph and no budget constraints.

In this paper, we consider a different problem that requires minimizing the distance between opinions and the desired value using a given control/marketing budget. Moreover, we assume that the maximal marketing influence cannot instantaneously shift the opinion of one individual to the desired value. Basically, we consider a continuous time opinion dynamics and we want to design a marketing strategy that minimizes the distance between opinions and the desired value after a given finite number of discrete-time campaigns under budget constraints. The main motivation for this choice is the time scale of relevant events. The campaign refers to sales before some events and their duration is much smaller than the duration of the spreading of opinions related to the advertised products. There exist many practical situations where the use of an hybrid OD model seems completely natural. For instance, during a presidential campaign it is common to measure the opinions of the electors through polls just before, and just after, a time-localized event such as a big political meeting or a TV debate (see e.g., [23]). Despite its natural relevance, to the best of the authors' knowledge, no hybrid controlled OD model has been proposed to study the opinion dynamics in social networks under an external influence. To solve this control design problem we write the overall closed-loop dynamics as a linear-impulsive system and we show that the optimal strategy is to influence as much as possible the most central/popular individuals (see [24] for a formal definition of centrality) of the network as far as the graph modeling the social network is weakly connected (*i.e.* it contains at least a directed spanning tree). We also point out that the budget allocation has to take into account the size of clusters (maximal subsets of weakly connected agents, see [12]) when the graph is disconnected.

To the best of our knowledge, our work is different from all the existing results on opinion dynamics control. Unlike the few previous works on the control of opinions in social networks, we do not control the state of the influencing entity. Instead, we consider that value as fixed and we control the influence weight that the marketer has on different individuals of the social network. By doing so, we emphasize the advantages of targeted marketing with respect to broadcasting (uniform) strategies when budget constraints have to be taken into account. Moreover, we show that, although the individual control action $u_i(t_k)$ at time t_k can be chosen in the interval $[0, \bar{u}]$, the optimal choice is discrete: either 0 or \bar{u} .

The rest of the paper is organized as follows. Section 2 formulates the opinion dynamics control problem under consideration. A useful preliminary result for solving a specific optimization problem with constraints is given in Section 3. To motivate our analysis, we emphasize in Section 4 the improvements that can be obtained by targeted advertising with respect to a uniform/broadcasting control. Section 5 contains the results related to the optimal control strategy. We first analyze the case when the campaign budget is given a priori and must be optimally partitioned among the network agents. Secondly, we look at the case when the campaign budget is unknown but the campaigns are distanced in time. Thirdly, we consider the case of large networks that can be approximated as a union of clusters/sub-

networks. All these three cases point out that the optimal control contains only 0 or \bar{u} actions. These results motivate us to study in Section 6 the space-time distribution of the budget under the assumption that all the components of $u(t_k)$ are either 0 or \bar{u} . We conclude that the budget has to be allotted according to the influence power of each agent which in turn depends on the initial condition, centrality and the size of the clusters in which it lies. Numerical examples and concluding remarks end the paper.

2. Problem statement

We consider an entity (for example, a company) that is interested in attracting consumers to some product (electrical cars, healthy food, etc). Consumers belong to a social network and we refer to any consumer as an agent. For the sake of simplicity, we consider a fixed social network over the set of vertices $\mathcal{V} = \{1, 2, \dots, N\}$ of N agents. In other words, we identify each agent with its index in the set \mathcal{V} . To agent $i \in \mathcal{V}$ we assign a normalized scalar opinion $x_i(t) \in [0, 1]$ that evolves under the influence of neighbors' opinions and external entity persuasion/advertising action. We use $x(t) = (x_1(t), x_2(t), \dots, x_N(t))^T$ to denote the state of the network at any time t , where $x(t) \in \mathcal{X}$ and $\mathcal{X} = [0, 1]^N$.

In order to obtain a larger market share with a minimum investment, the external entity applies an action vector on marketing campaigns at discrete time instants. The set of campaigns time instants is finite: $\mathcal{T} = \{t_0, t_1, \dots, t_M\}$. The number of campaigns M is considered to be finite but arbitrarily large because we are interested in the finite (arbitrarily large) time behavior of the network. A given action therefore corresponds to a given marketing campaign aiming at influencing the consumer's opinion. Between two consecutive campaigns, the consumer's opinion is only influenced by the other consumers of the networks. We assume that $t_k - t_{k-1} = \delta_k \in [\delta_{\min}, \delta_{\max}]$ where $0 < \delta_{\min} < \delta_{\max}$ are two fixed real numbers.

Throughout the paper we consider $d \in \{0, 1\}$ be the desired opinion that the external entity would like to be adopted for all the consumers. We also consider $\forall i \in \mathcal{V}$ the following dynamics:

$$\begin{cases} \dot{x}_i(t) = \sum_{j=1}^N a_{ij}[x_j(t) - x_i(t)], & t \in [t_k, t_{k+1}) \\ x_i(t_k) = u_i(t_k)d + [1 - u_i(t_k)]x_i(t_k^-) \end{cases}, \quad \forall k \in \mathbb{N}, \quad (1)$$

where $u_i(t_k) \in [0, \bar{u}]$, $\forall i \in \mathcal{V}$, where $\bar{u} \in (0, 1)$ is a saturation on each component of the control, and

$$\sum_{k=0}^M \sum_{i=1}^N u_i(t_k) \leq B \quad (2)$$

where B represents the total budget of the external entity for the marketing campaigns.

It is worth pointing out that external influences are modeled through a sequence of impulsive dynamics (second equation of (1)). This corresponds either to the case when the duration of the campaign is much shorter than the time between two consecutive campaigns or to the case in which the real dynamics during the campaign is neglected and only the resulting state is used as an entry for the next inter-campaign period. It can also be noticed that the state-jump resulting from external influence is both related to the budget allocated to Agent i at time t_k i.e., $u_i(t_k)$ and the value of the state before advertising $x_i(t_k^-)$. While the former proportionality is intuitive the later expresses an increasing resistance of individuals while approaching the advertised state.

It is also important to highlight that we assume a uniform behavior of the agents with respect to external influence. In real social networks, some agents (central ones for instance) may be harder to influence. This means that for a given value of the external influence their state jump will be smaller than the jump of other agents under the same external influence. This can be done by adding a scaling factor in the second equation of (1).

Dynamics (1) can be rewritten using the collective variable $X(t) = [d, x(t)^T]^T$ as:

$$\begin{cases} \dot{X}(t) = -\mathcal{L}X(t) \\ X(t_k) = \mathcal{P}X(t_k^-) \end{cases}, \quad (3)$$

where

$$\mathcal{L} = \begin{pmatrix} 0 & \mathbf{0}_{1,N} \\ \mathbf{0}_{N,1} & L \end{pmatrix}, \mathcal{P} = \begin{pmatrix} 1 & \mathbf{0}_{1,N} \\ u(t_k) & I_N - \text{diag}(u(t_k)) \end{pmatrix}$$

with $\text{diag}(u(t_k)) \in \mathbb{R}^{N \times N}$ being the diagonal matrix having the components of $u(t_k)$ on the diagonal. Here, L is the Laplacian matrix associated to the graph formed by the adjacency matrix elements $a_{i,j}$, i.e., $L_{ij} = -a_{ij}$ for $i \neq j$ and $L_{ii} = \sum_{i \neq j} a_{ij}$.

Definition 1. The (vector) centrality of Agent i is the i^{th} component of the left eigenvector v of L associated with the eigenvalue 0 and satisfying $v^T \mathbf{1}_N = 1$.

Remark 1. It is worth noticing that:

- \mathcal{L} is a Laplacian matrix corresponding to a network of $N + 1$ agents. The first agent represents the external entity and is not connected to any other agent while the rest of the agents represents the consumers and interact through the social network defined by the influence weights a_{ij} .
- \mathcal{P} is a row stochastic matrix that can be interpreted as a Perron matrix associated with the directed graph having the external entity as a central node. This node is not influenced by the network and keeps its state constant. On the other hand it possibly influences (notice that components of $u(t_k)$ can be 0) all the other nodes. Consequently, without budget constraints, the network can reach, at least asymptotically, the state $d\mathbf{1}_N$.

Several space-time control strategies can be implemented under budget constraints. For instance, we can spend the same budget for each agent i.e. $u_i(t_k) = u_j(t_k)$, $\forall (i, j) \in \mathcal{V}^2$, we can also allocate the entire budget for specific agents of the network. Moreover, the budget can be spent either on few or many campaigns. Our objective is to design a space-time control strategy that minimizes the following cost function

$$J^\infty = \sum_{i=1}^N \lim_{t \rightarrow \infty} |x_i(t) - d| \quad (4)$$

This can be interpreted as follows. If the entity (a company for example) is interested in convincing the public to buy some product or change their habits (practice sports or quit smoking for instance), it will try to move the asymptotic consensus value of the network as close as possible from the desired value, i.e. minimize J^∞ . In some other cases, like an election campaign which targets to get the opinions close to d within a finite time T , we will minimize $J^T = \sum_{i=1}^N |x_i(T) - d|$. Therefore, in the absence of additional campaigns, the agents will exchange their opinion through the network and asymptotically converge to certain local or global agreements. It is worth noting that, after the last campaign, the system state converges exponentially fast. Indeed, in the absence of campaigns, the system dynamics is just $\dot{x}(t) = -Lx(t)$ which has the consensus manifold as a global uniformly exponentially stable attractor. This means that $x_i(T)$ is a good approximation of $\lim_{t \rightarrow \infty} x_i(t)$ when T is sufficiently large.

3. Preliminaries

Before starting the analysis of the multi-agent system in the presence of external influence, we state a key lemma which will help us to find the optimal solutions for the considered scenarios for the budget allocation problem.

Lemma 1. Given an optimization problem (OP) under the following standard form

$$\begin{aligned} & \underset{y \in \mathbb{R}^N}{\text{minimize}} && C(y) \\ & \text{subject to} && y_i - \bar{y} \leq 0, \forall i \in \{1, \dots, N\} \\ & && -y_i \leq 0, \forall i \in \{1, \dots, N\} \\ & && \sum_{i=1}^N y_i - B \leq 0 \end{aligned} \quad (5)$$

where $N \in \mathbb{N}, N \geq 1, \bar{y} < 1, B \geq 0$ and $C(y)$ is a decreasing convex function in y_i such that one of the following two conditions hold.

Case 1: $\forall i \in \{1, \dots, N\}, \exists g(y) \geq 0$ such that

$$\frac{\partial C(y)}{\partial y_i} = -c_i g(y)$$

for some $c_i \in \mathbb{R}$.

Case 2: $\frac{\partial C(y)}{\partial y_i} = -\frac{1}{1-y_i}$ for all $i \in \{1, \dots, N\}$.

Then an optimal solution y^* to this OP is given by water-filling as follows

$$y_{\circ(i)}^* = \begin{cases} \bar{y} & \text{if } i \leq \lfloor \frac{B}{\bar{y}} \rfloor \\ B - \bar{y} \lfloor \frac{B}{\bar{y}} \rfloor & \text{if } i = \lfloor \frac{B}{\bar{y}} \rfloor + 1 \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

where $\circ : \{1, \dots, N\} \mapsto \{1, \dots, N\}$ represents an ordering function which can be any bijection for Case 2 and, one satisfying $c_{\circ(1)} \geq c_{\circ(2)} \geq \dots \geq c_{\circ(N)}$ for Case 1.

Proof: Note that all the constraint functions of the considered OP are affine, which corresponds to sufficient conditions for applying KKT conditions. Since the OP is convex, KKT conditions are necessary and sufficient for optimality. By denoting the Lagrangian by

$$\ell(y, \lambda, \lambda', \widehat{\lambda}) = C(y) + \sum_{i=1}^N \lambda_i (y_i - \bar{y}) - \sum_{i=1}^N \lambda'_i y_i + \widehat{\lambda} \left(\sum_{i=1}^N y_i - B \right), \quad (7)$$

the first necessary and sufficient condition for optimality writes:

$$-\nabla C(y^*) = \sum_{i=1}^N \lambda_i^* \nabla (y_i^* - \bar{y}) - \sum_{i=1}^N (\lambda_i^*)' \nabla y_i^* + \widehat{\lambda}^* \nabla \left(\sum_{i=1}^N y_i^* - B \right). \quad (8)$$

The primal feasibility conditions write

$$0 \leq y_i^* \leq \bar{y} \quad \forall i \in \{1, \dots, N\}$$

and

$$\sum_{i=1}^N y_i^* \leq B. \quad (9)$$

All the KKT multipliers must satisfy the dual feasibility conditions: $\lambda_i^* \geq 0, (\lambda_i^*)' \geq 0, \widehat{\lambda}^* \geq 0$ for all $i \in \{1, \dots, N\}$. At last, the complementary slackness conditions are given by

$$\begin{aligned} \lambda_i^* (y_i^* - \bar{y}) &= 0, \\ (\lambda_i^*)' y_i^* &= 0, \\ \widehat{\lambda}^* \left[\left(\sum_{i=1}^N y_i^* \right) - B \right] &= 0. \end{aligned}$$

Case 1: Let us assume that $\frac{\partial C(y)}{\partial y_i} = c_i g(y)$. Then, we can simplify (8) for component i to get

$$c_i g(y^*) = \widehat{\lambda} + \lambda_i^* - (\lambda_i^*)' \quad (10)$$

which must hold for all $i \in \{1, \dots, N\}$. Since $g(y^*)$ is identical for all $i \in \{1, \dots, N\}$ and is non-negative, we must have $\lambda_i^*, (\lambda_i^*)'$ and $\widehat{\lambda}$ chose so that the above equation holds. Take y^* from (6). Set $\lambda_j^* = (\lambda_j^*)' = 0$ for $j = \left(\circ \left[\frac{B}{\bar{y}} \right] \right)$ as it

is the only component with a non-saturated solution. For any i such that $o(i) < j$, we have $c_i \geq c_j$ and this can be satisfied by setting $y_i^* = \bar{u}$ and having $\lambda_i^* > 0$ and $\lambda_i'^* = 0$. On the other hand, for any i such that $o(i) > j$, we set $y_i^* = 0$ and the KKT conditions are satisfied if $\lambda_i^* = 0$ and $(\lambda_i^*)' > 0$. The solution from (6) can also be verified to satisfy (9) and therefore, we have it satisfying all the KKT conditions.

Case 2: When $\frac{\partial C(y)}{\partial y_i} = \frac{-1}{1-y_i}$, the role of c_i is replaced by $\frac{-1}{1-y_i}$ and the agent with $y_i^* = 0$ has the lowest absolute value on the left side of the Lagrangian (*i.e.* 1), and the agent with the saturation $y_i^* = \bar{y}$ has the highest absolute value (*i.e.*, $\frac{1}{1-\bar{u}}$). Due to symmetry, any agents can be chosen to have the min or max saturation. ■

Lemma 1 provides a water-filling type optimal allocation policy for the problem considered. The solution in Case 1 is to select the best agent in terms of the coefficient c_n , allocate the maximum possible y_n to it, then allocate the remaining budget to the next best agent and so on. This implies sorting the agents based on c_n , which is done using the function o , and saturating the y_n for the first $\lfloor \frac{B}{\bar{y}} \rfloor$ agents, assigning the remaining budget to the next agent and 0 to the rest.

4. Performance analysis of the considered benchmark strategy

As a reference strategy, we consider the broadcasting-based marketing. For every campaign, it consists in allocating the available campaign budget uniformly among all the consumers. However, the campaign budget is allowed to vary over time that is, from campaign to campaign, under the total budget constraint. In order to highlight the potential of designing more advanced strategies, we show here that for some particular network topologies it is possible to quantify analytically the potential gain brought by implementing target marketing (*i.e.* using space-time strategies) over broadcasting strategies.

First, let us compute the optimal revenue that we can get by broadcasting strategies *i.e.* $u_i(t_k) = u_j(t_k) \triangleq \alpha_k, \forall i, j \in \mathcal{V}$. We suppose that the graph representing the social network contains a directed spanning tree (*i.e.* a directed graph in which, except the root which is not influenced, each node is influenced by a single other node called parent). Let v be the left eigenvector of L associated with the eigenvalue 0 and satisfying $v^\top \mathbf{1}_N = 1$. Therefore, in the absence of any control action, one has that $\lim_{t \rightarrow \infty} x(t) = v^\top x(0) \mathbf{1}_N \triangleq x_0^\infty$. Let us also introduce the following notation:

$$x_k^\infty = \lim_{t \rightarrow \infty} e^{-L(t-t_k)} x(t_k) = v^\top x(t_k) \mathbf{1}_N, \quad \forall k \in \mathbb{N}. \quad (11)$$

Following (3) and using $\delta_k = t_{k+1} - t_k$, $D_k = \text{diag}(u(t_k))$ one deduces that:

$$x_{k+1}^\infty = v^\top x(t_{k+1}) \mathbf{1}_N = v^\top \left[u(t_{k+1})d + (I_N - D_{k+1})x(t_{k+1}^-) \right] \mathbf{1}_N = v^\top \left[u(t_{k+1})d + (I_N - D_{k+1})e^{-L\delta_k} x(t_k) \right] \mathbf{1}_N. \quad (12)$$

Since $v^\top L = \mathbf{0}_N$ one has that $v^\top e^{-L\delta_k} = v^\top$ and consequently one obtains that

$$x_{k+1}^\infty - x_k^\infty = v^\top \left(u(t_{k+1})d - D_{k+1}e^{-L\delta_k} x(t_k) \right) \mathbf{1}_N. \quad (13)$$

In the case of broadcasting one has $u(t_k) = \alpha_k \mathbf{1}_N$ and $D_k = \alpha_k I_N$, where $\alpha_k \in [0, \bar{u}]$ for all $k \in \{0, \dots, M\}$. Therefore, using $v^\top \mathbf{1}_N = 1$, (13) becomes

$$x_{k+1}^\infty - x_k^\infty = \alpha_{k+1} (d \mathbf{1}_N - x_k^\infty), \quad (14)$$

which can be equivalently rewritten as

$$(d \mathbf{1}_N - x_{k+1}^\infty) = (1 - \alpha_{k+1})(d \mathbf{1}_N - x_k^\infty). \quad (15)$$

Using (15) recursively one obtains that

$$J^\infty(\alpha) = |\mathbf{1}_N^\top (d \mathbf{1}_N - x_M^\infty)| = \prod_{\ell=0}^M (1 - \alpha_\ell) |\mathbf{1}_N^\top (d \mathbf{1}_N - x_0^\infty)| = (Nd - \mathbf{1}_N^\top x_0^\infty) \prod_{\ell=0}^M (1 - \alpha_\ell). \quad (16)$$

where $J^\infty(\alpha)$ denotes the cost associated with a broadcasting strategy using α_k at campaign t_k .

Proposition 4.1. *The cost $J^\infty(\alpha)$ obtained when implementing broadcasting is minimized by using the maximum possible investments as soon as possible, i.e., for all $k \in \{0, \dots, M\}$,*

$$\alpha_k = \begin{cases} \bar{u} & \text{if } k \leq \lfloor \frac{B}{N\bar{u}} \rfloor \\ \frac{B}{N} - \bar{u} \lfloor \frac{B}{N\bar{u}} \rfloor & \text{if } k = \lfloor \frac{B}{N\bar{u}} \rfloor + 1 \\ 0 & \text{otherwise} \end{cases} \quad (17)$$

Proof: Minimizing $J^\infty(\alpha)$ under broadcasting strategy assumption is equivalent with the minimization of $\prod_{\ell=0}^{k+1} (1 - \alpha_\ell)$. This is equivalent to minimizing

$$C(\alpha) = \log \left(\prod_{\ell=0}^{k+1} (1 - \alpha_\ell) \right) \quad (18)$$

and we have

$$\frac{\partial C}{\partial \alpha_\ell} = -\frac{1}{1 - \alpha_\ell} \quad (19)$$

This results in an OP which satisfies the conditions to use Lemma 1 Case 2. ■

It is noteworthy that for $u_i \in [0, 1)$ one has that

$$\prod_{\ell=0}^{k+1} (1 - \alpha_\ell) \geq 1 - \sum_{\ell=0}^{k+1} \alpha_\ell \geq 1 - \frac{B}{N}. \quad (20)$$

The last inequality in (20) comes from the broadcasting hypothesis $u_i(t_\ell) = \alpha_\ell$, $\forall i \in \mathcal{V}$ and consequently the budget spent in the ℓ -th campaign is $N \cdot \alpha_\ell$. Therefore, the total budget for $k+2$ campaigns is $N \sum_{\ell=0}^{k+1} \alpha_\ell$ and has to be smaller than B .

Thus

$$J = |\mathbf{1}_N^\top (d\mathbf{1}_N - x_{k+1}^\infty)| \geq (1 - \frac{B}{N}) |\mathbf{1}_N^\top (d\mathbf{1}_N - x_0^\infty)|. \quad (21)$$

The interpretation of (20) is that for broadcasting strategy the minimal cost J is obtained when the whole budget is spent in one marketing campaign (provided this is possible i.e. $B \leq N\bar{u}$), otherwise the first inequality in (20) becomes strict meaning that

$$J > (1 - \frac{B}{N}) |\mathbf{1}_N^\top (d\mathbf{1}_N - x_0^\infty)|. \quad (22)$$

Let us now suppose that the graph under consideration is a directed spanning tree having the first node as root. Then, using a targeted marketing in which the external entity influences only the root, we will show that, under the same budget constraints, the cost J will be smaller. Indeed, for this graph topology one has $v = (1, 0, \dots, 0)^\top$ yielding $x_k^\infty = x_1(t_k)\mathbf{1}_N$. Moreover, the dynamics of $x_1(\cdot)$ writes as:

$$\begin{cases} \dot{x}_1(t) = 0, & t \in [t_k, t_{k+1}) \\ x_1(t_k) = u_1(t_k)d + (1 - u_1(t_k))x_1(t_k^-), \end{cases} \quad \forall k \in \mathbb{N}. \quad (23)$$

Therefore,

$$x_1(t_k) = u_1(t_k)d + (1 - u_1(t_k))x_1(t_{k-1}) \quad (24)$$

yielding

$$d - x_1(t_k) = [1 - u_1(t_k)][d - x_1(t_{k-1})], \quad (25)$$

which is equivalent to (15). As we have seen before, in the broadcasting strategies one has $\sum_{\ell=0}^{k+1} \alpha_\ell \leq \frac{B}{N}$ while targeting only the root, the constraint becomes $\sum_{\ell=0}^{k+1} u_1(t_\ell) \leq B$. Therefore, for any given broadcasting strategy (u_1, u_2, \dots, u_k) there exists a targeted on the root strategy that consists in repeating N times (u_1, u_2, \dots, u_k) . Doing so, one obtains

$$(d\mathbf{1}_N - x_{k+1}^\infty) = \left[\prod_{\ell=0}^{k+1} (1 - \alpha_\ell) \right]^N (d\mathbf{1}_N - x_0^\infty). \quad (26)$$

which leads to a cost which is seen to be less than the one obtained when using broadcasting-based marketing.

5. General optimal control strategy

First, we rewrite the optimal control problem as an optimization problem by treating the control $u(t_k)$ as an NM -dimensional vector to optimize. We denote $u_{i,k} = u_i(t_k)$ to represent the control for agent $i \in \mathcal{V}$ at time t_k . Then our problem can be rewritten as

$$\begin{aligned} & \underset{u \in \mathbb{R}^{NM}}{\text{minimize}} && J^\infty(u) \\ & \text{subject to} && u_{i,k} - \bar{u} \leq 0, \quad \forall i \in \mathcal{V}, k \in \{0, \dots, M\} \\ & && -u_{i,k} \leq 0, \quad \forall i \in \mathcal{V}, k \in \{0, \dots, M\} \\ & && \sum_{i=1}^N \sum_{k=1}^M u_{i,k} - B \leq 0 \end{aligned} \quad (27)$$

Here, $J^\infty(u)$ is seen as a multilinear function. Before solving problem (27) we want to get further insights on structure of the optimal solution, which will lead to important simplifications. Therefore, instead of solving the general optimization problem (27), we first consider splitting our problem into time-allocation and space-allocation.

Assumption 1. *The graph $\mathcal{G} = (\mathcal{V}, L)$ is weakly connected (sometimes referred to as quasi-strongly connected) i.e. it contains at least one directed spanning tree.*

Assumption 1 is standard in the analysis of multi-agent systems and guarantees that information flows over the entire network. In our analysis this assumption is not essential but we start analyzing networks that satisfy it and, in a second step, we solve the budget allocation problem over disconnected networks. When Assumption 1 holds, if we know that for campaign k a maximum budget of $\beta_k \leq B$ has been allocated (i.e., for a given time-allocation), we find the optimal control strategy for the k -th campaign. Moreover, for long campaign duration (i.e. $t_{k+1} - t_k$ large) and given time budget allocation $(\beta_0, \dots, \beta_M)$, we provide a computationally oriented optimal space allocation of the budget. Based on these results, we propose a discrete-action space-time control strategy. Next, we extend the results to the case when Assumption 1 does not hold and the network consists of a union of weakly connected clusters.

5.1. Minimizing the per-campaign cost

In this section we consider that Assumption 1 holds and the budget β_k for each campaign is a priori given. The objective is to find the spatial allocation of the budget that optimizes the cost $|\mathbf{1}_N^\top (d\mathbf{1}_N - x_k^\infty)|$ associated with the time allocation $(\beta_0, \dots, \beta_M)$. Consequently, the following budget constraint has to be considered at campaign k :

$$\sum_{i=1}^N u_i(t_k) \leq \beta_k. \quad (28)$$

The associated cost for the campaign k is written as

$$J_k^\infty(u(t_k)) = |\mathbf{1}_N^\top (d\mathbf{1}_N - x_k^\infty)| = N \cdot \left| d - \sum_{i=1}^N v_i x_i(t_k) \right| = N \cdot \left| d - \sum_{i=1}^N v_i (u_i(t_k)d + [1 - u_i(t_k)]x_i(t_k^-)) \right|. \quad (29)$$

This rewriting of the cost allows us to define the right quantity to measure the influence power of an agent, which translates the gain the marketer can make by investing on this agent. The corresponding quantity is defined and used in the next proposition.

Proposition 5.1. *Define the influence power of Agent i as $p_i^k = v_i |d - x_i(t_k^-)|$. Denote by $\pi_k : \mathcal{V} \rightarrow \mathcal{V}$, a bijection which sorts the agents based on decreasing p_i^k , i.e. $p_{\pi_k(1)}^k \geq p_{\pi_k(2)}^k \geq \dots \geq p_{\pi_k(N)}^k$. Under Assumption 1 the cost $J_k^\infty(u(t_k))$ is minimized by the following investment profile*

$$u_{\pi(i)}^*(k) = \begin{cases} \bar{u} & \text{if } i \leq \left\lfloor \frac{\beta_k}{\bar{u}} \right\rfloor \\ \beta_k - \bar{u} \left\lfloor \frac{\beta_k}{\bar{u}} \right\rfloor & \text{if } i = \left\lfloor \frac{\beta_k}{\bar{u}} \right\rfloor + 1 \\ 0 & \text{otherwise} \end{cases} \quad (30)$$

Proof: Note that minimizing $J_k^\infty(u(t_j))$ is equivalent with the minimization of

$$C(u(t_k)) = \left(d - \sum_{i=1}^N v_i(u_i(t_k)d + (1 - u_i(t_k))x_i(t_k^-)) \right)^2$$

with the constraints $0 \leq u_i(t_k) \leq \bar{u}$ for all $i \in \mathcal{V}$ and (28). We notice that

$$\frac{\partial C}{\partial u_i(t_k)} = -2v_i(d - x_i(t_k^-))(d - x_k^\infty) \quad (31)$$

where we used the notation $x_k^\infty = v^\top x(t_k)$.

If $d = 1$, then $x_i(t_k^-) \leq 1$, $\forall i \in \mathcal{V}$ and

$$(d - x_i(t_k^-))(d - x_k^\infty) \geq 0.$$

On the other hand, if $d = 0$, we have $x_i(t_k^-) \geq 0$, $\forall i \in \mathcal{V}$ and

$$(d - x_i(t_k^-))(d - x_k^\infty) \geq 0.$$

Therefore, we can rewrite the above equation as

$$\frac{\partial C}{\partial u_i(t_k)} = -\gamma_i g(u(t_k)), \quad g(u(t_k)) = |d - x_k^\infty| \quad (32)$$

which satisfies the conditions to use Case 1 of Lemma 1. ■

5.2. Space allocation for long campaign duration

In the following we consider that Assumption 1 holds and a finite number of marketing campaigns with a priori fixed budget are scheduled such that $t_{k+1} - t_k$ is very large for each $k \in \{0, 1, \dots, M-1\}$. Due to Assumption 1 and the long duration of the campaigns, we can assume that $x_i(t_{k+1}^-) = x_k^\infty$ for all $i \in \mathcal{V}$ and $k \in \{0, 1, \dots, M-1\}$. Under this assumption, we write

$$x_i(t_1^-) = x_0^\infty(u(t_0)) = \sum_{i=1}^N v_i(du_i(t_0) + x_i(t_0^-)(1 - u_i(t_k))) \quad (33)$$

for any $i \in \mathcal{V}$. Subsequently, we have

$$x_k^\infty(u(t_0), u(t_1), \dots, u(t_k)) = \sum_{i=1}^N v_i \left[du_i(t_k) + x_{k-1}^\infty(u(t_0), \dots, u(t_{k-1}))(1 - u_i(t_k)) \right] \quad (34)$$

for all $k \in \{1, 2, \dots, M\}$. Our objective is to minimize

$$J^\infty(u) = N \cdot |x_M^\infty(u(t_0), \dots, u(t_M)) - d|$$

and this can be done using the Proposition 5.2 below.

Proposition 5.2. Define $\rho_k : \mathcal{V} \rightarrow \mathcal{V}$ a bijection such that $\rho_0 = \pi_0$ (defined in Proposition 5.1) and for all $k \in \{1, 2, \dots, M\}$, ρ_k gives the agent index after sorting over v_i , i.e. $v_{\rho_k(1)} \geq v_{\rho_k(2)} \geq \dots \geq v_{\rho_k(N)}$. Let Assumption 1 hold and the time budget allocation be given by $\beta = (\beta_0, \dots, \beta_M)$ such that $\sum_{k=1}^M \beta_k \leq B$ and $\beta_k \leq N\bar{u}$. Then, the optimal allocation per agent minimizing the cost $J(u)$ is given by

$$u_{\rho_k(i)}^*(k) = \begin{cases} \bar{u} & \text{if } i \leq \left\lfloor \frac{\beta_k}{\bar{u}} \right\rfloor \\ \beta_k - \bar{u} \left\lfloor \frac{\beta_k}{\bar{u}} \right\rfloor & \text{if } i = \left\lfloor \frac{\beta_k}{\bar{u}} \right\rfloor + 1 \\ 0 & \text{otherwise} \end{cases} \quad (35)$$

Proof: Minimizing $J^\infty(u)$ is equivalent with the minimization of

$$C(u) = (d - x_M^\infty(u(t_0), \dots, u(t_M)))^2$$

with the constraints given in (27).

We have

$$\frac{\partial C(u)}{\partial u_{i,k}} = 2(d - x_M^\infty) \frac{\partial x_M^\infty}{\partial u_{i,k}}$$

for any $i \in \mathcal{V}$, $k \in \{0, 1, \dots, M\}$. Observe that

$$\frac{\partial x_k^\infty}{\partial u_{i,k}} = v_i(d - x_{k-1}^\infty)$$

for $k \in \{1, \dots, M\}$ and

$$\frac{\partial x_0^\infty}{\partial u_{i,0}} = \pm \gamma_i$$

with the sign being negative if $x_i(t_0^-) > d$ and positive otherwise. We also have

$$\frac{\partial x_k^\infty}{\partial u_{i,k-1}} = \frac{\partial x_{k-1}^\infty}{\partial u_{i,k-1}} \sum_{i \in \mathcal{V}} v_i(1 - u_{i,k})$$

Using the equations above iteratively, we have

$$\frac{\partial x_M^\infty}{\partial u_{i,k}} = v_i(d - x_{k-1}^\infty) \prod_{j=k+1}^M \sum_{i \in \mathcal{V}} v_i(1 - u_{i,j})$$

for $k \geq 1$ and

$$\frac{\partial x_M^\infty}{\partial u_{i,0}} = \pm \gamma_i \prod_{j=1}^M \sum_{i \in \mathcal{V}} v_i(1 - u_{i,j})$$

Therefore,

$$\frac{\partial C(u)}{\partial u_{i,k}} = 2(d - x_M^\infty) v_i(d - x_{k-1}^\infty(u)) \prod_{j=k+1}^M \sum_{i \in \mathcal{V}} v_i(1 - u_{i,j})$$

for $k \geq 1$ and

$$\frac{\partial C(u)}{\partial u_{i,0}} = 2|d - x_M^\infty(u)| \gamma_i \prod_{j=1}^M \sum_{i \in \mathcal{V}} v_i(1 - u_{i,j})$$

Assume that the optimal time-allocation of budget is known and is given by $\beta = (\beta_0, \beta_1, \dots, \beta_M)$ such that

$$\sum_{i=1}^N u_{i,k} = \beta_k, \quad \forall k \in \{0, 1, \dots, M\}$$

Then, the optimal spatial allocation problem within any campaign k is an OP satisfying Case 1 of Lemma 1. ■

Due to the long stage duration, the opinions are in consensus for all t_k , except for the case of $k = 0$, which is the first campaign. This means that in the first stage, the agents are sorted based on both their initial opinions and their centralities, but from the next stage onwards, only their centralities are considered.

5.3. Budget allocation for clusterized networks

In practice, when the social network becomes large, several issues may appear. Computational complexity limitations may prevent from operating with NM -dimensional spaces. Also, uncertainties or inaccuracies on L or the centrality vector may appear. Motivated by these two observations, we formulate here the solution in the case where the network is structured into clusters which may be naturally present or arise from a dynamical process. Clusters can be determined from cluster detection algorithms such as the one proposed in [25, 12]) and emphasize a time-scale separation in the dynamics of the system. The subsequent analysis could use the time-scale modeling ([26, 27, 28]) but we would have to carefully deal with the jumps introduced by the advertising campaigns. Instead of doing that, in this paper we simplify the analysis and the presentation by neglecting the weakest interconnections that may appear between agents belonging to different clusters. In other words we analyze the behavior of the reduced-order dynamics. Therefore, the network is assumed to be the union of a certain number of weakly connected clusters; in that case, we show that the cluster size has to be taken into account to optimally allocate the available budget.

Let us consider that $\mathcal{V} = \bigcup_{i=1}^m C_i$, where C_1, \dots, C_m are disjoint subsets of agents which are weakly connected. For any $i \in \{1, \dots, m\}$ we also denote by N_i the cardinality of cluster C_i and by $x_{C_i} \in \mathbb{R}^{N_i}$ the column vector collecting all the states of the agents in cluster C_i . Since Assumption 1 does not hold, only local agreements corresponding to each cluster are obtained. Basically, $L = \text{diag}(L_1, \dots, L_m)$ where $L_i \in \mathbb{R}^{N_i \times N_i}$ is the Laplacian matrix associated with the interactions in cluster C_i . In the sequel we denote by $x_{C_i, k}^\infty$ the agreement value of cluster C_i starting from the initial condition $x_{C_i}(t_k)$. In other words, if the last advertising campaigns takes place at time t_k than the system converges to the following state:

$$x_k^\infty = \lim_{t \rightarrow \infty} e^{-L(t-t_k)} x(t_k) = \begin{pmatrix} x_{C_1, k}^\infty \mathbf{1}_{N_1} \\ \vdots \\ x_{C_m, k}^\infty \mathbf{1}_{N_m} \end{pmatrix}, \quad \forall k \in \mathbb{N}. \quad (36)$$

A fixed number of advertising campaigns is assumed and the last campaign takes place at time t_M . The overall cost to minimize can be expressed as:

$$J^\infty(u) = \sum_{i=1}^m N_i \cdot |x_{C_i, M}^\infty(u(t_0), \dots, u(t_M)) - d|. \quad (37)$$

In order to characterize the optimal control strategy in the case of disconnected networks we introduce some additional notation: v^i is the left eigenvector of L_i associated with the simple eigenvalue 0 and satisfying $(v^i)^\top \mathbf{1}_{N_i} = 1$. It can be noticed that

$$x_{C_i, k}^\infty(u(t_0), u(t_1), \dots, u(t_k)) = \sum_{j \in C_i} v_j^i \left[du_j(t_k) + x_{C_i, k-1}^\infty(u(t_0), \dots, u(t_{k-1})) (1 - u_j(t_k)) \right] \quad (38)$$

for all $k \in \{1, 2, \dots, M\}$. This observation is exploited in the next proposition to define an appropriate measure of the influence power of an agent. Here again, the corresponding quantity is used to express the optimal budget allocation policy.

Proposition 5.3. *Let $\mathcal{V} = \bigcup_{i=1}^m C_i$, where C_1, \dots, C_m are disjoint subsets of agents which are weakly connected. Let also the time budget allocation be given by $\beta = (\beta_0, \dots, \beta_M)$ such that $\sum_{k=1}^M \beta_k \leq B$ and $\beta_k \leq N\bar{u}$. For each agent $j \in \mathcal{V} \cap C_i$ define the influence power as $s_j^k = N_i \cdot v_j^i \cdot |d - x_j(t_k^-)|$. At last, define $\sigma_k : \mathcal{V} \mapsto \mathcal{V}$ a bijection which sorts the agents based on decreasing s_j^k i.e. $s_{\mathcal{H}(1)}^k \geq s_{\mathcal{H}(2)}^k \geq \dots \geq s_{\mathcal{H}(N)}^k$. Then, the optimal allocation per agent minimizing the cost $J(u)$ is given by*

$$u_{\sigma_k(i)}^*(k) = \begin{cases} \bar{u} & \text{if } i \leq \left\lfloor \frac{\beta_k}{\bar{u}} \right\rfloor \\ \beta_k - \bar{u} \left\lfloor \frac{\beta_k}{\bar{u}} \right\rfloor & \text{if } i = \left\lfloor \frac{\beta_k}{\bar{u}} \right\rfloor + 1 \\ 0 & \text{otherwise} \end{cases} \quad (39)$$

Proof: The result follows again by applying Case 1 of Lemma 1. We avoid unnecessary details and we just point out that

$$\frac{\partial J(u)}{\partial u_{j,k}} = -s_j^k \quad (40)$$

leading to the desired result when g is identically 1 in Lemma 1. ■

6. Discrete-action space-time control strategy

Motivated by the results in Propositions 5.1 and 5.2, in this section we consider that $u_i(t_k) \in \{0, \bar{u}\}, \forall i \in \mathcal{V}, k \in \mathbb{N}$ and $B = Q\bar{u}$ with $Q \in \mathbb{N}$ given a priori. The objective is to numerically find the best space-time control strategy for a given initial state x_0 of the network.

6.1. Brute force algorithm

We will consider in turn the cases of short and long campaigns.

In the short-campaign case, given a time allocation consisting of the budgets $\beta_k = b_k \bar{u}$ at each campaign, either Proposition 5.1 (if the directed graph is weakly connected) or Proposition 5.3 (if the network is clustered) tells us how to allocate each campaign budget optimally across the agents. Denote all possible budgets at one campaign by $\mathcal{B} = \{0, \dots, \min\{N, Q\}\}$. A simple algorithm is then to search in a brute-force manner all possible time allocations $\mathbf{b} = (b_0, \dots, b_M) \in \mathcal{B}^{M+1}$, subject to the constraint $\sum_k b_k \leq Q$. For each such vector \mathbf{b} , we simulate the system from x_0 with dynamics (1) where the budget b_k is allocated with Proposition 5.1 or 5.3. After the last campaign, we compute with the appropriate formula the final, infinite-time state of the network $x_F(\mathbf{b})$. We retain a solution with the best average cost:

$$\min_{\mathbf{b}} \frac{1}{N} \sum_{i=1}^N |x_{i,F}(\mathbf{b}) - d|$$

where subscript i is the agent index (recall that the agents all have the same opinion at infinite time if the network is weakly connected, and if it is clustered each cluster has its own opinion). Note that we report the average cost over the agents, J^∞/N , instead of the sum J^∞ because this version is easier to interpret as a mean deviation of each agent from the target state. Furthermore, the simulation can be done in closed form, using the fact that $x^-(t_{k+1}) = e^{-L\bar{\delta}_k} x(t_k)$. The complexity of this search is $O(N^3(M+1)(\min\{N, Q\}+1)^{M+1})$, dominated by the exponential term. Therefore, this approach will only be feasible for small values of N or Q , and especially of M .

Considering now the long-campaign case for weakly connected networks, we could still implement a similar brute-force search, but using dynamics (34) for inter-campaign propagation and Proposition 5.2 for allocation over agents. However, now we can do better by taking advantage of the fact that for all $k > 1$, the opinions of all the agents reach identical values. Using this, we will derive a more efficient, dynamic programming solution to the optimal control problem:

$$\min_{\mathbf{b}} |x_{1,F}(\mathbf{b}) - d|$$

where the long-campaign dynamics apply but by a slight abuse of notation we keep it the same as above.

6.2. Dynamic programming algorithm

When the graph is weakly connected and we are in the long-campaign case, we are able to provide a dynamic programming (DP) algorithm that is much more efficient than the brute force search above. Owing to the long campaigns, the agents have already reached consensus by t_1 and so we can use a scalar variable $y_k = |d - x_{k-1}^\infty|$ to represent the cost (or equivalently, the state of the network) before the campaign for all $k \in \{1, 2, \dots, M+1\}$, i.e., from the second campaign onwards. After the initial-campaign decision at $k = 0$, y_1 is computed in a special way since the network is not yet at consensus. To see how, consider a fixed, given initial opinion $x(t_0)$. Then, $y_1 = f_0(b_0)$ where f_0 describes the evolution of the network after allocating the first-campaign budget b_0 using Proposition 5.1:

$$f_0(b_0) = \sum_{i=1}^{b_0} v_{\rho_0(i)} (1 - \bar{u}) |x_{\rho_0(i)}(t_0) - d| + \sum_{i=b_0+1}^N v_{\rho_0(i)} |x_{\rho_0(i)}(t_0) - d| \quad (41)$$

Now, to compute y_k after the decisions at subsequent campaigns $k \in \{1, 2, \dots, M\}$, i.e. for t_1, t_2, \dots, t_M , define function $f : \mathbb{Z}_{\geq 0} \rightarrow (0, 1]$:

$$f(b) = \left(1 - \bar{u} \sum_{i=1}^b v_{\rho_1(i)} \right). \quad (42)$$

If b_k denotes the budget allocated to campaign k , then it can be shown that after the optimal spatial allocation described in Proposition 5.2,

$$y_{k+1} = y_k f(b_k), \quad (43)$$

for all $k \in \{1, 2, \dots, M\}$. This lets us write the final cost as

$$y_{M+1} = y_1 \prod_{k=1}^M f(b_k). \quad (44)$$

In addition to the network state, we will need an additional integer state $r_k \in \{0, \dots, Q\}$ that keeps track of the remaining budget. This variable is initialized to the total budget $r_0 = Q$ and evolves according to $r_{k+1} = r_k - b_k$.

For any given y_1 obtained after using a budget b_0 during campaign 0, minimization of the final cost involves minimizing $\prod_{k=1}^M f(b_k)$, with the constraint $\sum_{k=1}^M b_k \leq r_1$. Since $f(b) > 0$ for any $b \geq 0$, we can minimize the final cost by minimizing the logarithm of the product mentioned above, i.e., the minimization of the final cost is equivalent to the following optimization problem:

$$\begin{aligned} & \min_{b_0, b_1, \dots, b_M} \log(f_0(b_0)) + \sum_{k=1}^M \log(f(b_k)), \\ & \text{such that } \sum_{k=0}^M b_k \leq Q, \\ & \text{where } b_k \in \{0, 1, \dots, \min\{Q, N\}\}, \forall k \in \{0, \dots, M\} \end{aligned} \quad (45)$$

where the budget allocated to each campaign is upper-bounded by the budget available and the number of agents in the network as assumed in Proposition 5.3.

In order to implement the DP algorithm, we keep a value function V_k that represents at each k the sum of logarithms from step k onwards. This value function only depends on the remaining budget:

$$\begin{aligned} V_M(r_M) &= \log(f(r_M)), \\ V_k(r_k) &= \min_{b_k \in \{0, 1, \dots, \min\{N, r_k\}\}} [\log(f(b_k)) + V_{k+1}(r_k - b_k)] \quad \text{for } k \in \{M-1, M-2, \dots, 1\}, \\ V_0 &= \min_{b_0 \in \{0, 1, \dots, \min\{N, Q\}\}} [\log(f_0(b_0)) + V_1(Q - b_0)] \end{aligned} \quad (46)$$

To understand this algorithm, note first that because f is a strictly decreasing function, the optimal budget to use at campaign M is all the remaining budget, which leads to the final-campaign DP initialization $V_M(r_M) = \log(f(r_M))$. Then, at intermediate campaigns, we simply minimize the summation in (45) with a DP rule. Finally, for the initial campaign $k = 0$, instead of using f , the cost for any initial budget b_0 is calculated explicitly by taking $\log(f_0(b_0))$ defined in (41). Note that V_0 is a scalar constant, that represents the total optimal cost of the solution.

Once V_k is available, an optimal solution is found by a forward pass, as follows:

$$\begin{aligned} b_0^* &= \arg \min_{b_0 \in \{0, 1, \dots, \min\{Q, N\}\}} [\log(f_0(b_0)) + V_1(Q - b_0)] \\ b_k^* &= \arg \min_{b_k \in \{0, 1, \dots, \min\{r_k, N\}\}} [\log(f(b_k)) + V_{k+1}(r_k - b_k)], \quad \text{for } k = 1, \dots, M-1 \\ b_M^* &= r_M \end{aligned} \quad (47)$$

The complexity of the backward pass for value function computation is $O(MNQ \min\{N, Q\})$ (the complexity of the forward pass is much smaller). To develop an intuition, take the case $N < Q$; then the algorithm is quadratic in N and linear in M and Q . This allows us to apply the algorithm to much larger problems than the brute-force search above.

6.3. Numerical results

In this section, we begin by exemplifying on a small-scale problem: the short-campaign brute-force algorithm; long-campaign DP; and the clustered case (again with the brute-force method). After that, to illustrate the scalability

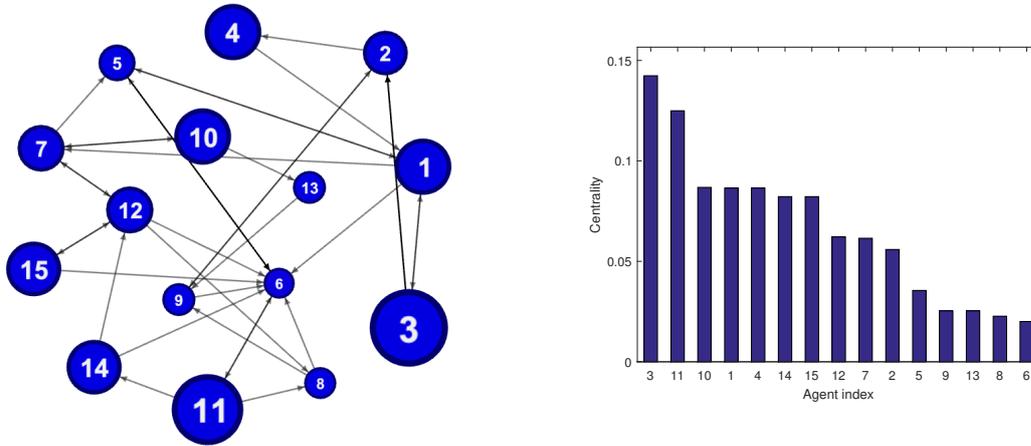


Figure 1. Left: small-scale weakly connected graph. Right: agent centralities, sorted in descending order.

of the proposed DP method with respect to the size of the network and to the number of campaigns, we devise a larger experiment with more agents and campaigns, and apply DP to it.

The small-scale problem has $N = 15$ agents, and we start with the weakly connected graph from Figure 1, left. The target opinion d is 1. The initial opinions of the agents are distributed on an equidistant grid, starting from 0 for agent 1, up to 1 for agent 15 (so agents with smaller indices have opinions closer to zero). The centrality of each agent is shown in Figure 1, right. There are 4 campaigns, corresponding to $M = 3$, and the budget $Q = N = 15$ and $\bar{u} = 0.2$. For a short campaign length $\delta_k = 0.5 \forall k$, the brute-force approach gets the results from Figure 2. The final cost (each individual agent's difference from the desired opinion) is 0.3259. Examining now the list of agents influenced, we see that these agents are generally among those with large centralities. Nevertheless, relatively lower-centrality agents are preferred when their opinion is far from the target (as is the case for agents 1 and 4, whose initial opinion is small).

To better see the advantages of a well-designed advertising strategy, we compare the results above with the uncontrolled case (no advertising), and to the broadcast strategy (which consists of spending the entire budget at the initial time, with \bar{u} allocated to each agent). The cost without using any control in this situation is 0.5135, and the cost with the broadcast strategy is 0.4108, i.e., we observe a 20% gain over the uncontrolled system using the broadcast and around 20% gain over the broadcast using the optimal strategy.

For the second experiment, in the same network of agents, we consider long campaigns, i.e., $t_{k+1} - t_k \rightarrow \infty$. We apply DP, with the results shown in Figure 3. The solution is different from the short-campaign case, which is especially visible at campaigns $k \geq 1$, where only the most central agents are influenced. The final cost is slightly larger, 0.3457. To better understand the meaning of the long campaigns, note that the network can be (informally speaking) associated with a time constant T equal to the inverse of the smallest real part among all the eigenvalues of the Laplacian L excluding the zero eigenvalue, and as soon as $t_{k+1} - t_k$ is significantly larger than $4T$, the network effectively reaches consensus in-between campaigns so we may consider we are in the long-campaign case. For the particular graph here, $T \approx 3.28$. Note that we can directly compare this long-campaign result with no-control and broadcasting above (since those strategies are independent of campaign length), and we still see significant improvement over both.

Next, to illustrate the results of Section 5.3, we take the graph in Figure 1 and remove all the links between agents 1 to 4 and the rest of the agents, obtaining the graph in Figure 4. This new graph has two clusters, the first consisting of agents 1 to 4, and the second of the rest of the agents. Four campaigns of length 0.5 are considered, like before. The brute-force algorithm is applied with $Q = N = 15$, starting from an initial state of the network where all agents have opinion 0.5 (this is done so that they all have the same initial deviation from the desired state, which better exposes the influence of their centrality and group size). The results are shown in Figure 5. It is interesting to observe that

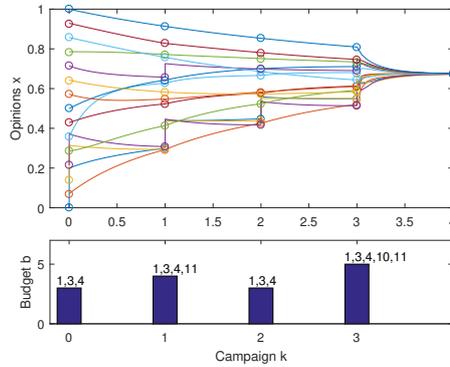


Figure 2. Results for short campaigns. The bottom bar plot shows the budget allocated by the algorithm at each campaign, with the agents influenced in each campaign shown above each bar. The top plot shows the opinions of the agents, with an additional, long campaign converging to the average opinion (so the last campaign duration is not to scale). The circles indicate the opinions *right before* applying the control at each campaign; note the discontinuous transitions of the opinions after control.

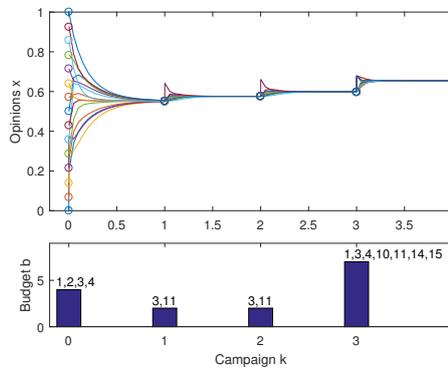


Figure 3. Results for long campaigns. The continuous opinion dynamics is plotted for a sufficiently long time to observe the long campaign behavior, *i.e.* the convergence of opinions of the agents (which means the horizontal axis is not to scale).

despite their lower centrality, many agents in cluster 2 (e.g. 7, 10, and 14) are given preference over agents 1, 2, and 4 in cluster 1. This is because the number of agents is larger in the second cluster, and the selection criterion (39) takes this into account. To compare, we have the final opinion without control to be 0.5 (since all agents start with the same opinion), with the broadcast strategy to be 0.4, and with the optimal strategy to be around 0.34, which corresponds to 85% of the cost with the broadcast strategy.

Finally, we move on to the problem where we test the scalability of DP for large graphs and many campaigns. Specifically, we take 100 agents and 20 campaigns. Link weights generated from a uniform distribution over $[0, 1]$, after which any link with a weight smaller than 0.3 is removed. Initial opinions are equidistantly spaced in $[0, 1]$ as before, and the total budget is $Q = N = 100$. The final cost here is 0.38 instead of 0.5 without any control. The obtained cost is close to the one related to the broadcast strategy (which is 0.4) because all nodes have very similar

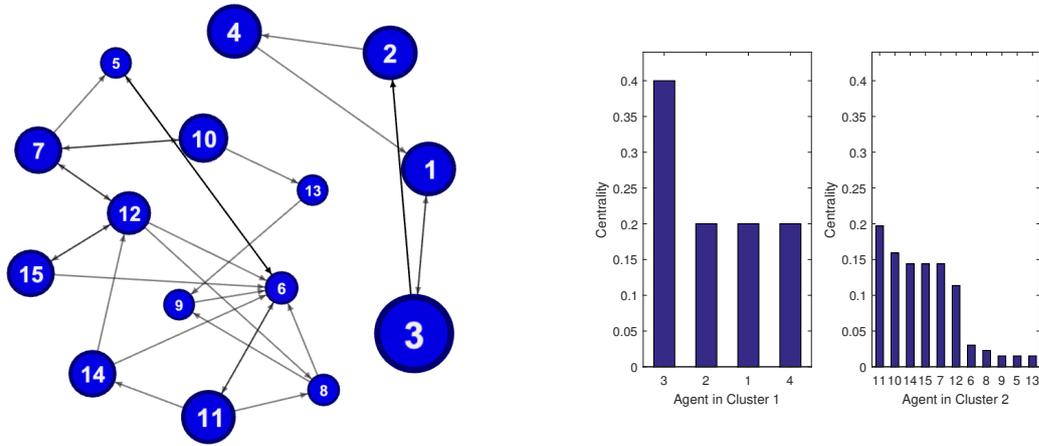


Figure 4. Left: clustered graph. Right: agent centralities in the two clusters.

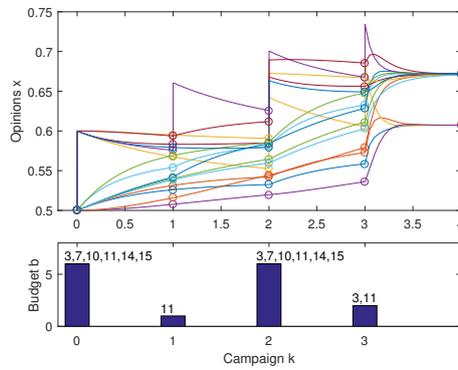


Figure 5. Results for the clustered problem.

centrality. Consequently, the DP cost is 76% of the cost without any control and 95% of the cost with broadcast strategy. As expected large part of the budget (47%) is used on the first campaign towards agents having initial condition closer to 0. Note that the brute-force approach would be entirely unfeasible in this problem, while the execution time of DP in Matlab is around 1.7 s on an Intel i7-3540M CPU.

7. Conclusions

In this paper, we have proposed a mathematical formulation of the problem of target marketing over social networks. We show how to exploit some properties of the social network graph in the design of the marketing budget

allocation over the agents and over marketing campaigns. A marketer should mainly consider the initial opinion of an agent, its centrality, and (when relevant) the number of agents in the cluster it belongs to. Based on this, we have defined appropriate quantities which measure the influence power of an agent and which allows the marketer to define an order in which it has to allocate its influence budget. The derived budget allocation policies are shown to have a water-filling-type structure. The conducted numerical analysis allows one to extract many precious insights on how to invest a budget over consumers and time. For instance, key consumers to be influenced immediately appear, the number of campaigns to be performed is easily obtained, and the impact of having an advanced marketing campaign (versus allocating the available budget uniformly over consumers and campaigns) can be quantified.

References

- [1] M. H. DeGroot, Reaching a consensus, *Journal of the American Statistical Association* 69 (345) (1974) 118–121.
- [2] N. E. Friedkin, E. C. Johnsen., Social influence and opinions., *Journal of Mathematical Sociology*. 15 (1990) 193–206.
- [3] G. Deffuant, D. Neau, F. Amblard, G. Weisbuch, Mixing beliefs among interacting agents, *Advances in Complex Systems* 3 (2000) 87–98.
- [4] R. Hegselmann, U. Krause, Opinion dynamics and bounded confidence models, analysis, and simulation, *Journal of Artificial Societies and Social Simulation* 5 (3).
- [5] C. Altafini, Consensus problems on networks with antagonistic interactions, *IEEE Transactions on Automatic Control* 58 (4) (2013) 935–946.
- [6] N. R. Chowdhury, I.-C. Morărescu, S. Martin, S. Srikant, Continuous opinions and discrete actions in social networks: a multi-agent system approach, in: *Proceedings 55th IEEE Conference on Decision and Control*, 2016, pp. 1739–1744.
- [7] J. H. Davis, *Understanding group behavior : Consensual action by small groups*, Vol. 1, Lawrence Erlbaum Associates, 1996, Ch. Group decision making and quantitative judgments : A consensus model, pp. 35–60.
- [8] Y. Ohtsubo, A. Masuchi, D. Nakanishi, Majority influence process in group judgment : Test of the social judgment scheme model in a group polarization context., *Group processes & intergroup relations* 5 (3) (2002) 249–261.
- [9] C. V. Kerckhove, S. Martin, P. Gend, P. J. Rentfrow, J. M. Hendrickx, V. D. Blondel, Modelling influence and opinion evolution in online collective behaviour, *PLoS ONE* 11 (6) (2016) 1–25.
- [10] R. Axelrod, The dissemination of culture: A model with local convergence and global polarization, *The Journal of Conflict Resolution*, 41 (2) (1997) 203–226.
- [11] S. Galam, S. Moscovici, Towards a theory of collective phenomena: Consensus and attitude changes in groups, *European Journal of Social Psychology* 21 (1) (1991) 49–74.
- [12] I.-C. Morărescu, A. Girard, Opinion dynamics with decaying confidence: Application to community detection in graphs, *IEEE Transactions on Automatic Control* 56 (8) (2011) 1862 – 1873.
- [13] M. Caponigro, B. Piccoli, F. Rossi, E. Trélat, Sparse feedback stabilization of multi-agent dynamics., in: *Proceedings of the 55th IEEE Conference on Decision and Control*, 2016, pp. 4278–4283.
- [14] F. Dietrich, S. Martin, M. Jungers, Control via leadership of opinion dynamics with state and time-dependent interactions, *IEEE Transactions on Automatic Control* 10.1109/TAC.2017.2742139.
- [15] R. Hegselmann, S. Kurz, C. Niemann, J. Rambau, Optimal opinion control : The campaign problem, *Journal of Artificial Societies and Social Simulation* 18 (3).
- [16] B. Fogg, *Persuasive technology: Using computers to change what we think and do.*, Morgan Kaufmann, 2002.
- [17] M. L. Fransen, E. G. Smit, P. W. J. Verlegh., Strategies and motives for resistance to persuasion: an integrative framework, *Frontiers in Psychology* 6.
- [18] R. Abebe, J. Kleinberg, D. Parkes, C. Tsourakakis, Opinion dynamics with varying susceptibility to persuasion, <https://arxiv.org/pdf/1801.07863.pdf>.
- [19] J. Leskovec, L. A. Adamic, B. A. Huberman, The dynamics of viral marketing, *ACM Transactions on the Web (TWEB)* 1 (1) (2007) 5–es.
- [20] D. Arthur, R. Motwani, A. Sharma, Y. Xu, Pricing strategies for viral marketing on social networks, in: *International workshop on internet and network economics*, Springer, 2009, pp. 101–112.
- [21] A. M. Masucci, A. Silva, Strategic resource allocation for competitive influence in social networks, in: *Communication, Control, and Computing (Allerton)*, 2014 52nd Annual Allerton Conference on, IEEE, 2014, pp. 951–958.
- [22] V. S. Varma, I.-C. Morărescu, S. Lasaulce, S. Martin, Opinion dynamics aware marketing strategies in duopolies, in: *IEEE Conference on Decision and Control*, 2017, pp. 3859–3864.
- [23] S. D. McClurg, The electoral relevance of political talk: Examining disagreement and expertise effects in social networks on political participation, *American Journal of Political Science* 50 (3) (2006) 737–754.
- [24] P. Bonacich, P. Lloyd, Eigenvector-like measures of centrality for asymmetric relations, *Social Networks* 23 (2001) 191–201.
- [25] V. Blondel, J.-L. Guillaume, R. Lambiotte, E. Lefebvre, Fast unfolding of communities in large networks, *Journal of Statistical Mechanics: Theory and Experiment* P10008 (10).
- [26] J. Chow, P. Kokotović, Time scale modeling of sparse dynamic networks, *IEEE Transactions on Automatic Control* 30 (8) (1985) 714–722.
- [27] E. Bıyık, M. Arcak, Area aggregation and time-scale modeling for sparse nonlinear networks, *Systems & Control Letters* 57 (2) (2007) 142–149.
- [28] S. Martin, I.-C. Morărescu, D. Nešić, Time scale modeling for consensus in sparse directed networks with time-varying topologies., in: *Proceedings 55th IEEE Conference on Decision and Control*, 2016, pp. 7–12.