

Space-time budget allocation for marketing over social networks [★]

Irinel-Constantin Morărescu ^{*}, Vineeth Satheeskumar Varma ^{*},
Lucian Buşoniu ^{**}, Samson Lasaulce ^{***}

^{*} *Université de Lorraine, CNRS, CRAN, F-54000 Nancy, France. email: (constantin.morarescu,vineeth.satheeskumar-varma)@univ-lorraine.fr*

^{**} *Automation Department, Technical University of Cluj–Napoca, Romania. email:lucian@busoniu.net*

^{***} *Laboratoire des Signaux et Systemes (L2S, CNRS - CentraleSupélec - Univ. Paris Sud), Gif-sur-Yvette, France. email: samson.lasaulce@lss.supelec.fr*

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 to its own side by using a specific influence budget during discrete-time advertising campaigns; consequently, the overall closed-loop 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 a desired opinion; for instance, the marketer may prioritize certain agents over others based on their influence in the social graph. 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.

© 2018, IFAC (International Federation of Automatic Control) Hosting by Elsevier Ltd. All rights reserved.

Keywords: Social networks, hybrid systems, optimal control.

1. INTRODUCTION

The last decades have witnessed an increasing interest in the study of opinion dynamics in social networks. This is mainly motivated by the fact that people's opinions are increasingly influenced through digital social networks. Therefore, governmental institution but also private companies consider that marketing over social networks becomes a key tool for promoting new products or ideas. However, 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 dynamics. Various mathematical models (DeGroot, 1974; Friedkin and Johnsen., 1990; Deffuant et al., 2000; Hegselmann and Krause, 2002; Altafini, 2013; Chowdhury et al., 2016) have been proposed to capture more features of these complex dynamics. Empirical models based on in vitro and in vivo experiments have also been developed (Davis, 1996; Ohtsubo et al., 2002; Kerckhove et al., 2016).

The emergence of consensus received a particular attention in opinion dynamics (Axelrod, 1997; Galam and Moscovici, 1991). While some mathematical models naturally lead to consensus (DeGroot, 1974; Friedkin and Johnsen., 1990), others lead to network clustering (Hegselmann and Krause, 2002; Altafini, 2013; Morărescu and Girard, 2011). In order to enforce consensus, some recent studies propose the control of one or a few agents, see Caponigro et al. (2016); Dietrich et al. (2017).

[★] This work was supported by projects PEPS INS2I IODINE and PEPS S2IH INS2I YPSOC funded by the CNRS.

Beside 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 (Hegselmann et al., 2015). In (Masucci and Silva, 2014), 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, (Varma et al., 2017) 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 a 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. 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 of the network. (see Bonacich and Lloyd (2001) for a formal definition of centrality).

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 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. Both cases point out that the optimal control contains only 0 or \bar{u} actions. These results motivate us to study in Section 6 the spatio-temporal distribution of the budget under the assumption that all the components of $u(t_k)$ are either 0 or \bar{u} . Numerical examples and concluding remarks end the paper.

2. PROBLEM STATEMENT

We consider an entity (for example, a governmental institution or a company) that is interested in attracting consumers to a certain opinion. 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 $n \in \mathcal{V}$ we assign a normalized scalar opinion $x_n(t) \in [0, 1]$ that represents its opinion level and can be interpreted e.g., as the probability for an agent to act as desired. 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 $u(t_k) = (u_1(t_k), \dots, u_N(t_k))^T \in \mathcal{U}_i$ in marketing campaigns at discrete time instants $t_k \in \mathcal{T}$, $k \in \mathbb{N}$. A given action therefore corresponds to a given marketing campaign aiming at influencing the consumer's opinion. The instants corresponding to the campaigns are known and are collected in the set $\mathcal{T} = \{t_0, t_1, \dots, t_M\}$. Between two consecutive campaigns, the consumer's opinion is only influenced by the other consumers of the network. We assume that $t_k - t_{k-1} = \delta_k \in [\delta_m, \delta_M]$ where $\delta_m < \delta_M$ are two fixed real numbers. Throughout the paper we refer to $d \in \{0, 1\}$ as the desired opinion that the external entity wants to be adopted. We consider $\forall i \in \mathcal{V}$ the following dynamics:

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

where $u_i(t_k) \in [0, \bar{u}]$ with $\bar{u} < 1$, $\forall i \in \mathcal{V}$ and

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

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

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

$$\begin{cases} \dot{X}(t) = -\mathcal{L}X(t) \\ X(t_k) = \mathcal{P}X(t_k^-) \end{cases}, \quad (2)$$

where

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

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.

Remark 1. It is noteworthy 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 are connected 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 tree having the external entity as a parent of all the other nodes. Consequently, without budget constraints, the network reaches, at least asymptotically, the value d .

Several space-time control strategies can be implemented under the budget constraints. For instance, we can spend the same budget for each agent i.e., $u_i(t_k) = u_j(t_k)$, $\forall i, j \in \mathcal{V}$, 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

$$J^T = \sum_{i=1}^N |x_i(T) - d| \quad (3)$$

for some $T > t_M$, and we have the cost associated with the asymptotic opinion given by

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

This can be interpreted as follows. If the entity (a governmental institution 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 to 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 . It is worth mentioning that between campaigns and after the last campaign the opinions evolve according to consensus dynamics.

3. PRELIMINARIES

We first state a very useful Lemma which will help us to find the optimal solutions for many sub-cases of our problem.

Lemma 1. Given an optimization problem (OP) of the form

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

where $C(y)$ is a decreasing convex function in y_i such that one of the following two conditions hold

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

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

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

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

$$y_{\mathcal{R}(i)}^* = \begin{cases} \bar{y} & \text{if } i \leq \left\lfloor \frac{B}{L\bar{y}} \right\rfloor \\ B - \bar{y}L \left\lfloor \frac{B}{L\bar{y}} \right\rfloor & \text{if } i = \left\lceil \frac{B}{L\bar{y}} \right\rceil \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

where $\mathcal{R} : \{1, \dots, L\} \mapsto \{1, \dots, L\}$ is any bijection for Case 2 and, a bijection satisfying

$$c_{\mathcal{R}(1)} \geq c_{\mathcal{R}(2)} \geq \dots \geq c_{\mathcal{R}(L)}.$$

for Case 1.

4. THE BROADCASTING CASE STUDY

To emphasize the relevance of the problem under consideration, we will show that for some particular network topologies we can obtain a significant improvement of the revenue by using targeted marketing instead of broadcasting-based marketing in which the marketer allocates the same amount of resource to all the agents.

First, we derive the optimal revenue that can be obtained by implementing a broadcasting strategy *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 spanning tree. Let v be the right 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 $\lim_{t \rightarrow \infty} x(t) = v^\top x(0) \mathbf{1}_N \triangleq \underline{x}_0^\infty$. Let us also introduce the following notation:

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

Following (2) and using $\delta_k = t_{k+1} - t_k$, $D_k = \text{diag}(u(t_k))$ one deduces that:

$$\begin{aligned} x_{k+1}^\infty &= v^\top x(t_{k+1}) \mathbf{1}_N \\ &= v^\top (u(t_{k+1})d + (I_N - D_{k+1})x(t_{k+1}^-)) \mathbf{1}_N \\ &= v^\top (u(t_{k+1})d + (I_N - D_{k+1})e^{-L\delta_k} x(t_k)) \mathbf{1}_N. \end{aligned}$$

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

$$x_{k+1}^\infty - x_k^\infty = v^\top (u(t_{k+1})d - D_{k+1}e^{-L\delta_k} x(t_k)) \mathbf{1}_N. \quad (7)$$

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 $v^\top \mathbf{1}_N = 1$, (7) becomes

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

which can be equivalently rewritten as

$$(d\mathbf{1}_N - x_{k+1}^\infty) = (1 - \alpha_{k+1})(d\mathbf{1}_N - x_k^\infty). \quad (8)$$

Using (8) recursively one obtains that

$$J^\infty(\alpha) = (d\mathbf{1}_N - x_M^\infty) = \prod_{\ell=0}^M (1 - \alpha_\ell)(d\mathbf{1}_N - \underline{x}_0^\infty). \quad (9)$$

where $J^\mathbf{B}(\alpha)$ denotes the cost associated with a broadcasting strategy using α_k at stage t_k .

Proposition 1. The broadcasting cost $J^\infty(\alpha)$ is minimized by using the maximum possible investments as soon as possible, *i.e.*

$$\alpha_k = \begin{cases} \bar{u} & \text{if } k \leq \left\lfloor \frac{B}{N\bar{u}} \right\rfloor \\ B - \bar{u}N \left\lfloor \frac{B}{N\bar{u}} \right\rfloor & \text{if } k = \left\lceil \frac{B}{N\bar{u}} \right\rceil \\ 0 & \text{otherwise} \end{cases} \quad (10)$$

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

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

and we have

$$\frac{\partial C}{\partial \alpha_\ell} = -\frac{1}{1 - \alpha_\ell} \quad (11)$$

This results in an OP which satisfies the conditions to use Lemma 1 case 2. \blacksquare

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

$$\prod_{\ell=0}^{k+1} (1 - \alpha_\ell) \geq 1 - \sum_{\ell=0}^{k+1} \alpha_\ell \geq 1 - \frac{B}{N}. \quad (12)$$

The last inequality in (12) comes from the broadcasting hypothesis $u_i(t_\ell) = \alpha_\ell$, $\forall i \in \mathcal{V}$ which mean that 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

$$J = \mathbf{1}_N^\top |(d\mathbf{1}_N - x_{k+1}^\infty)| \geq (1 - \frac{B}{N}) \mathbf{1}_N^\top |(d\mathbf{1}_N - \underline{x}_0^\infty)|.$$

The interpretation of (12) is that for the 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 (12) becomes strict meaning that

$$J > (1 - \frac{B}{N}) \mathbf{1}_N^\top |(d\mathbf{1}_N - \underline{x}_0^\infty)|.$$

Let us now suppose that the graph under consideration is a 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:

$$\begin{cases} \dot{x}_1(t) = 0, & t \in [t_k, t_{k+1}) \\ x_1(t_k) = u_1(t_k)d + (1 - u_1(t_k))x_1(t_k^-), \end{cases} \quad \forall k \in \mathbb{N}. \quad (13)$$

Therefore,

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

yielding

$$d - x_1(t_k) = (1 - u_1(t_k))(d - x_1(t_{k-1})),$$

which is equivalent to (8). As we have seen before, in the broadcasting strategy one has $\sum_{\ell=0}^{k+1} \alpha_\ell \leq \frac{B}{N}$ whereas 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 strategy targeted on the root that consists of repeating N times (u_1, u_2, \dots, u_k) . Doing so, one obtains

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

which leads to a much smaller cost J i.e., the strategy is more efficient.

5. GENERAL OPTIMAL SPACE-TIME 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

$$\text{Minimize}_{u \in \mathbb{R}^{NM}} J^T(u)$$

$$\text{Subject to } 0 \leq u_{i,k} \leq \bar{u} \quad \forall i \in \mathcal{V}, k \in \{0, \dots, M\}, \quad (14)$$

$$\text{and } \sum_{i=1}^N \sum_{k=1}^M u_{i,k} \leq B$$

Here, $J^T(u)$ is a multilinear function in u . Before solving problem (14) we want to get further insights on the solution's structure, which will lead to important simplifications. Therefore, instead of solving the general optimization problem (14), we consider splitting our problem into time-allocation and space-allocation. For a given time-allocation, i.e. if we know that for stage k a maximum budget of $\beta_k \leq B$ has been allocated, we find the optimal control strategy for the k -th stage. Moreover, for long stage durations (i.e., $t_{k+1} - t_k$ large) and given temporal budget allocation $(\beta_0, \dots, \beta_M)$, we characterize the optimal space allocation of the budget. Based on these results, we propose a discrete-action spatio-temporal control strategy.

5.1 Minimizing the per-stage cost

In this section we consider that the budget β_k for each campaign is a priori given, and optimize the corresponding $|d\mathbf{1}_N - x_k^\infty|$. Denote the budget for stage k by β_k such that

$$\sum_{i=1}^N u_i(t_k) \leq \beta_k \quad (15)$$

The corresponding cost for the stage k is written as

$$J_k^\infty(u(t_k)) = (d\mathbf{1}_N - x_k^\infty) = \left| d - \sum_{i=1}^N v_i x_i(t_k) \right| \quad (16)$$

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

We use $\gamma_i = v_i |d - x_i(t_k^-)|$ to denote the gain by investing in agent $i \in \mathcal{V}$. Define by $\mathcal{R} : \mathcal{V} \rightarrow \mathcal{V}$, a bijection which sorts the agents based on decreasing γ_i , i.e.

$$\gamma_{\mathcal{R}(1)} \geq \gamma_{\mathcal{R}(2)} \geq \dots \geq \gamma_{\mathcal{R}(N)}$$

Proposition 2. The cost $J_k^\infty(u(t_k))$ is minimized by the following investment profile

$$u_{\mathcal{R}(i)}^*(k) = \begin{cases} \bar{u} & \text{if } i \leq \left\lfloor \frac{\beta_k}{\bar{u}} \right\rfloor \\ \beta_k - \bar{u} \left\lfloor \frac{\beta_k}{\bar{u}} \right\rfloor & \text{if } i = \left\lfloor \frac{\beta_k}{\bar{u}} \right\rfloor + 1 \\ 0 & \text{otherwise} \end{cases} \quad (17)$$

Proof. Due to space limitation the proof is omitted.

5.2 Space allocation for long stage duration

In the following we consider that 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 $k \in \{0, 1, \dots, M-1\}$. In this case, 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

$$x_i(t_1^-) = x_0^\infty(u(t_0)) = \sum_{i=1}^N v_i (du_i(t_0) + x_i(t_0^-)(1 - u_i(t_0))) \quad (18)$$

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

$$x_k^\infty(u(t_0), u(t_1), \dots, u(t_k)) = \sum_{i=1}^N v_i [du_i(t_k) + x_{k-1}^\infty(u(t_0), \dots, u(t_{k-1})) (1 - u_i(t_k))] \quad (19)$$

for all $k \in \{1, 2, \dots, M\}$.

Our objective is to minimize

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

and this can be done using the proposition below. First, let us define $\mathcal{S}_k : \mathcal{V} \rightarrow \mathcal{V}$ a bijection such that $\mathcal{S}_0 = \mathcal{R}$ and for all $k \in \{1, 2, \dots, M\}$, \mathcal{S}_k gives the agent index after sorting over v_i , i.e.,

$$v_{\mathcal{S}_k(1)} \geq v_{\mathcal{S}_k(2)} \geq \dots \geq v_{\mathcal{S}_k(N)}$$

Proposition 3. Let the temporal budget allocation be given by

$$\beta = (\beta_0, \dots, \beta_M) \text{ such that } \sum_{k=1}^M \beta_k \leq B \text{ and } \beta_k \leq N\bar{u}. \text{ Then,}$$

the optimal allocation per agent minimizing the cost $J(u)$ is given by

$$u_{\mathcal{S}_k(i)}^*(k) = \begin{cases} \bar{u} & \text{if } i \leq \left\lfloor \frac{\beta_k}{\bar{u}} \right\rfloor \\ \beta_k - \bar{u} \left\lfloor \frac{\beta_k}{\bar{u}} \right\rfloor & \text{if } i = \left\lfloor \frac{\beta_k}{\bar{u}} \right\rfloor + 1 \\ 0 & \text{otherwise} \end{cases} \quad (20)$$

Proof. Due to space limitation the proof is omitted.

6. DISCRETE-ACTION SPACE-TIME CONTROL STRATEGY

Motivated by the results in Propositions 2 and 3, in this section we consider that $u_i(t_k) \in \{0, \bar{u}\}$, $\forall i \in \mathcal{V}, k \in \mathbb{N}$ and $B = K\bar{u}$ with $K \in \mathbb{N}$ 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 Algorithms

Let us consider in turn the cases of short and long stages. In the short-stage case, given a time allocation consisting of the budgets $\beta_k = b_k \bar{u}$ at each stage, Proposition 2 tells us how to allocate each stage budget optimally across the agents. Denote all possible budgets at one stage by $\mathcal{B} = \{0, \dots, \min\{N, K\}\}$. A very 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 $\sum_k b_k \leq K$. 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 2, and we obtain a final state $x_F(\mathbf{b}) = v^\top x(t_M) \mathbf{1}_N$ (the infinite-time

state of the network after the last campaign). We retain a solution with the best cost:

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

where subscript 1 denotes the first agent (recall that the agents all have the same opinion at infinite time). Note that this cost is J^∞/N ; we do not sum it over all the agents because this version is easier to interpret as a deviation of each agent from the target state. Furthermore, the simulation can be done in closed form, using the fact that $x^-(t_{k+1}) = e^{-L\delta_k}x(t_k)$. The complexity of this search is $O(N^3(M+1)(\min\{N, K\}+1)^{M+1})$, dominated by the exponential term. Therefore, this approach will only be feasible for small values of N or K , and especially of M .

Considering now the long-stage case, we could still implement a similar brute-force search, but using dynamics (19) for inter-stage propagation and Proposition 3 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-stage dynamics apply but by a slight abuse we keep it the same as in the previous section.

To obtain the DP algorithm, define for $k = 1, \dots, M, M+1 =: F$ a new state signal $z_k = [y_k, r_k]^\top \in \mathcal{Z} := [0, 1] \times \{0, \dots, K\}$. In this signal, $y_k = x_{k-1}^\infty$, the opinion resulting from long-term propagation after the $k-1$ th campaign, and r_k is the remaining budget to apply (we will start from $r_0 = K$). Here, F is associated with the infinite-time state of the network. We will compute a value function $V_k(z_k)$, representing the best cost attainable from stage k , if the agent state is y_k and budget r_k remains:

$$V_F(z_F) = |x_F - d|, \forall r_F$$

$$V_k(z_k) = \min_{b_k=0}^{\min\{r_k, N\}} V_{k+1}(g(z_k, b_k)), k = M, \dots, 1$$

Here, the dynamics $g: \mathcal{Z} \times \mathcal{B} \rightarrow \mathcal{Z}$ are given by:

$$y_{k+1} = v^\top x(t_k), \quad r_{k+1} = r_k - b_k$$

where $x_i(t_k) = u_i(b_k)d + (1 - u_i(b_k))y_k$ is the network state after campaign k , in which the agent allocations $u_i(b_k)$ are computed by distributing budget b_k with Proposition 3.

At stage 0, special dynamics $g_0: \mathcal{X} \times \mathcal{B} \rightarrow \mathcal{Z}$ apply, because the initial state of the network cannot be represented by a single number:

$$y_1 = v^\top x(t_0), \quad r_1 = K - b_0$$

where $x_i(t_0) = u_i(b_0)d + (1 - u_i(b_0))x_{i,0}$ and – differently from the other steps – $u_i(b_0)$ is found with Proposition 2.

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

$$b_0^* = \arg \min_{b_0=0}^{\min\{K, N\}} V_1(g_0(x_0, b_0)), \quad z_1^* = g_0(x_0, b_0^*)$$

$$b_k^* = \arg \min_{b_k=0}^{\min\{r_k, N\}} V_{k+1}(g(z_k^*, b_k)), \quad z_{k+1}^* = g(z_k^*, b_k^*)$$

$$\text{for } k = 1, \dots, M$$

and the optimal cost of the overall solution is simply the minimum value at the first step.

To implement this algorithm in practice, we will discretize the continuous state y into Y points, and interpolate the value function on the grid formed by these points, see Buşoniu et al. (2010) for details. The complexity of the backward pass for value function computation is $O(MY(\min\{N, K\}+1)N)$ (we

disregard the complexity of the forward pass since it is much smaller). To develop an intuition, take the case $N < K$; then the algorithm is quadratic in N and linear in M and Y . This allows us to apply the algorithm to much larger problems than the brute-force search above. Finally, note that in principle we could develop such a DP algorithm for the short-stage problem, but there we cannot condense the network state into a single number. Each agent state would have to be discretized instead, leading to a memory and time complexity proportional to Y^N , which makes the algorithm unfeasible for more than a few agents.

6.2 Numerical results

We begin by evaluating the brute-force algorithm on a small-scale problem with short stages. Then, we move to the long-stage case, where for the same network we compare DP and the brute-force method, confirming that they get the same result.

Consider $N = 15$ agents connected on the graph from Figure 1, where the size of the node corresponds to its centrality. There are 4 stages, corresponding to $M = 3$, and the budget $K = 15 = N$. The initial states of the agents are random. For a short stage 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 0.2485. Table 1, left shows the agents influenced at each stage. Reexamining Figure 1, we see that most of these agents have a large centrality, which is also the reason for which the algorithm selects them. An exception is agent 6 which has a low centrality, but is still influenced at many stages as $x_6(0) \approx 0.1$.




Figure 1. Small-scale graph.

Campaign	Agents	Campaign	Agents
0	3,5,6,7,8,14	0	3,6,7,8
1	3,6,7	1	2,3,7,9
2	3,6,7,9	2	2,3,7,9
3	3,9	3	2,3,9

Table 1. Agents influenced in each campaign. Left: short stages. Right: long stages.

We take now the same problem and make a single change: the stages become long (i.e., $t_{k+1} - t_k \rightarrow \infty$). We apply DP, with a discretization of y into 10 points. The results are in Figure 3, with the specific agents being controlled shown on the right side of Table 1. Note the solution is different from the short-stage case, and the final cost is 0.2553, slightly worse, which indicates that giving up the fine-grained control of the




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

agents over time leads to some losses, but they are small. To evaluate the impact of function approximation (interpolation), we also run the brute-force search, since in this problem it is still feasible. It gives the same strategy and cost as DP, so we do not show the result. Note that unlike before, agent 6 is only influenced at $k = 0$ as the stage durations are long and its opinion value plays a role only at the first stage.

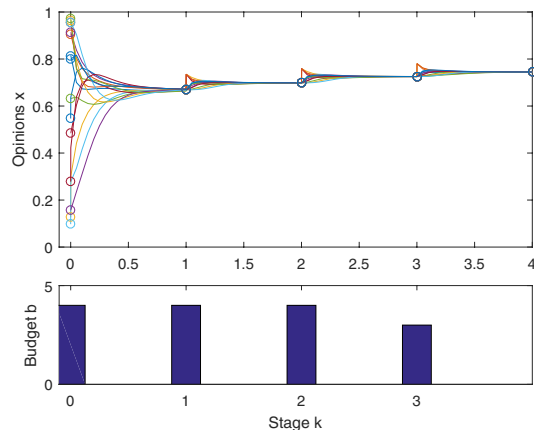


Figure 3. Results for long stages. The continuous opinion dynamics is plotted for $t \in [t_k, t_k + 25)$ per stage k , which is sufficient to observe the long stage behavior, *i.e.*, the convergence of opinions of the agents.

7. CONCLUSIONS

In its full generality, the problem of space-time budget allocation problem over a social network is seen to be non-trivial. However, it can be solved in several special cases of practical interest. If for every marketing campaign, the budget is allocated uniformly over the agents, the problem becomes a pure time budget control and can be solved. On the other hand, for a given time budget control, the problem becomes a pure space problem and the optimal way of allocating the budget is proved to be a water-filling allocation policy. Thirdly, if one goes for a binary budget allocation *i.e.*, the marketer either allocates a given amount of budget to an agent or nothing, the space-time budget allocation problem can be solved by using dynamic programming-based numerical techniques. Numerical results illustrate how the available budget should be used by the marketer to reach its objective in terms of desired opinion for the network.

REFERENCES

- Altafini, C. (2013). Consensus problems on networks with antagonistic interactions. *IEEE Transactions on Automatic Control*, 58(4), 935–946.
- Axelrod, R. (1997). The dissemination of culture: A model with local convergence and global polarization. *The Journal of Conflict Resolution*, 41(2), 203–226.
- Bonacich, P. and Lloyd, P. (2001). Eigenvector-like measures of centrality for asymmetric relations. *Social Networks*, 23, 191–201.
- Bușoniu, L., Ernst, D., De Schutter, B., and Babuška, R. (2010). Approximate dynamic programming with a fuzzy parameterization. *Automatica*, 46(5), 804–814.
- Caponigro, M., Piccoli, B., Rossi, F., and Trélat, E. (2016). Sparse feedback stabilization of multi-agent dynamics. In *Proceedings of the 55th IEEE CDC*, 4278–4283.
- Chowdhury, N.R., Morărescu, I.C., Martin, S., and Srikant, S. (2016). Continuous opinions and discrete actions in social networks: a multi-agent system approach. In *Proc. of 55th IEEE Conf. on Decision and Control*.
- Davis, J.H. (1996). *Understanding group behavior: Consensual action by small groups*, volume 1, chapter Group decision making and quantitative judgments: A consensus model, 35–60. Lawrence Erlbaum Associates.
- Deffuant, G., Neau, D., Amblard, F., and Weisbuch, G. (2000). Mixing beliefs among interacting agents. *Advances in Complex Systems*, 3, 87–98.
- DeGroot, M.H. (1974). Reaching a consensus. *Journal of the American Statistical Association*, 69(345), 118–121.
- Dietrich, F., Martin, S., and Jungers, M. (2017). Control via leadership of opinion dynamics with state and time-dependent interactions. *IEEE Trans. on Automatic Control*.
- Friedkin, N.E. and Johnsen, E.C. (1990). Social influence and opinions. *Journal of Mathematical Sociology*, 15, 193–206.
- Galam, S. and Moscovici, S. (1991). Towards a theory of collective phenomena: Consensus and attitude changes in groups. *European J. of Social Psychology*, 21(1), 49–74.
- Hegselmann, R. and Krause, U. (2002). Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of Artificial Societies and Social Simulation*, 5(3).
- Hegselmann, R., Kurz, S., Niemann, C., and Rambau, J. (2015). Optimal opinion control: The campaign problem. *Journal of Artificial Societies and Social Simulation*, 18(3).
- Kerckhove, C.V., Martin, S., Gend, P., Rentfrow, P.J., Hendrickx, J.M., and Blondel, V.D. (2016). Modelling influence and opinion evolution in online collective behaviour. *PLoS ONE*, 11(6), 1–25.
- Masucci, A.M. and Silva, A. (2014). Strategic resource allocation for competitive influence in social networks. In *52nd IEEE Annual Allerton Conference on Communication, Control, and Computing, 2014*, 951–958.
- Morărescu, I.C. and Girard, A. (2011). Opinion dynamics with decaying confidence: Application to community detection in graphs. *IEEE Trans. on Automatic Ctrl*, 56(8), 1862 – 1873.
- Ohtsubo, Y., Masuchi, A., and Nakanishi, D. (2002). Majority influence process in group judgment: Test of the social judgment scheme model in a group polarization context. *Group processes & intergroup relations*, 5(3), 249–261.
- Varma, V.S., Morărescu, I.C., Lasaulce, S., and Martin, S. (2017). Opinion dynamics aware marketing strategies in duopolies. In *IEEE Conf. on Decision and Control*.