On the Optimal Dissemination of Information in Social Networks

Makan Fardad, Xi Zhang, Fu Lin, and Mihailo R. Jovanović

Abstract— We consider social networks which contain agents that spread misinformation and refuse to change their opinion. For a fixed number of information disseminating agents, we formulate an optimization problem to find their optimal location within the network such that the spread of misinformation is countered and public awareness is maximally raised. Once the location of the information disseminators is identified, we examine how to maximize their social influence either by creating new social links or by strengthening their existing links. Our formulation leads to a combinatorial optimization problem that is solved using the alternating direction method of multipliers. Illustrative examples are provided to demonstrate our theoretical developments.

Index Terms— Alternating direction method of multipliers (ADMM), leader selection, misinformation, optimization, social networks, stochastic matrices.

I. INTRODUCTION

Influencing public opinion is the holy grail of political parties, advertising agencies, media outlets, community leaders, and economists. One of the main questions facing these entities is how to most effectively sway public opinion in their favor, under the constraints of limited resources and rivaling entities.

This paper is motivated by recent work [1] on opinion dynamics and the spread of (mis)information in social networks. A successful example of spreading misinformation is the Swift Boat Veterans for Truth campaign, which ran ads during the 2004 US presidential race. The paper [1] investigates how the presence of forceful agents, i.e., agents that refuse to update their beliefs, affects the propagation of misinformation in a society.

In contrast to [1], which performs an analysis of social networks in a stochastic setting, we consider problems of *optimal design* in a deterministic framework. Given a social network with preassigned forceful misinforming agents, we consider the problem of optimally choosing agents to serve as forceful information disseminators in order to maximally raise public awareness and counter misinformation.

To this end, we model the opinion of each individual as a real scalar. Such a characterization has been noted to be a good approximation of individuals' opinions [2], [3]; for example, it is argued in [2] that an individual's opinions on a wide range of social and political issues can be characterized by how liberal or conservative he/she is. We model the discrete time evolution of opinions using stochastic matrices, such that at every time step each agent updates his/her beliefs by taking a weighted average of the beliefs of those agents he/she socially interacts with.

The rest of the paper is organized as follows. In Sec. II we address the problem of how to choose a fixed number of information disseminators to maximize public awareness in the presence of misinforming agents. Our formulation leads to a combinatorial optimization problem that we tackle using the alternating direction method of multipliers. In Sec. III we consider the problem of optimally increasing the social influence of our information disseminators. Illustrative examples are provided in Sec. IV.

II. SELECTION OF INFORMATION DISSEMINATORS

Consider a social network composed of n agents whose beliefs evolve according to

$$\bar{x}(k+1) = A\,\bar{x}(k),$$

where $\bar{x}$ is a column n-vector composed of nonnegative values that represent the beliefs of the agents with regards to a particular social issue. A is a (right) stochastic matrix that satisfies $A_{ij} \geq 0$, $A\mathbb{1} = \mathbb{1}$, and $\mathbb{1}$ is the column vector of all ones. The justification for A being a right stochastic matrix is the following: At every time step, each agent updates his/her beliefs by taking a weighted average (equivalently, a convex combination) of the beliefs of those agents he/she socially interacts with. This implies that every row of A is composed of positive elements that sum up to one. Note that A is not necessarily a symmetric matrix, which means that the weight agent i puts on the beliefs of agent j is not necessarily equal to the weight agent j puts on the beliefs of agent i.

Next we consider a scenario in which k_ϕ agents are designated to be misinformers (MIs). The MI agents always hold the belief zero and refuse to update it. Let ϕ be an n-vector whose ith element is equal to one if agent i is an MI and zero otherwise, and define

$$D_\phi = \mathrm{diag}\{\phi\}.$$

The refusal of the MIs to update their beliefs means that the evolution of the entire social network is described by

$$\bar{x}(k+1) = (I - D_\phi)\,A\,\bar{x}(k).$$

Furthermore, since the beliefs of the MIs are always equal to zero, then a similar procedure to that described in [4] leads

to the new evolution equations

$$x(k + 1) = T x(k),$$

where x is a column $(n - k_\phi)$-vector composed of the beliefs of the non-MI agents, and T is the $(n - k_\phi) \times (n - k_\phi)$ principal submatrix of A obtained by eliminating all rows and columns corresponding to the MI agents.

We now describe the evolution model, which is used to formulate the main optimization problem. Given $x(k + 1) = T x(k)$, we seek k_ψ agents to act as information disseminators (IDs), with the purpose of informing the public and countering the spread of misinformation by the MIs. The ID agents always hold the belief one and refuse to update it. Let ψ be an $(n - k_\phi)$-vector whose ith element is equal to one if agent i is an ID and zero otherwise, and define

$$D_\psi = \mathrm{diag}\{\psi\}.$$

The refusal of the IDs to update their beliefs means that the evolution of the entire social network, minus the MIs, is described by

$$x(k + 1) = (I - D_\psi) T x(k) + D_\psi \mathbb{1}.$$

After a reordering and partitioning of the states, the above equation transforms into

$$\begin{bmatrix} \hat{x}(k+1) \\ \tilde{x}(k+1) \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ U_\psi & T_\psi \end{bmatrix} \begin{bmatrix} \hat{x}(k) \\ \tilde{x}(k) \end{bmatrix} + \begin{bmatrix} \mathbb{1} \\ 0 \end{bmatrix},$$

where $\hat{x}$ is a column k_ψ-vector corresponding to the beliefs of the ID agents, and $\tilde{x}$ is a column $(n - k_\phi - k_\psi)$-vector corresponding to the beliefs of those agents who are neither MIs nor IDs (i.e., $\tilde{x}$ encapsulates the beliefs of the general public). Defining $m := n - k_\phi - k_\psi$, the matrix T_ψ is the $m \times m$ principal submatrix of T obtained by eliminating all rows and columns corresponding to the ID agents, and U_ψ is the $m \times k_\psi$ submatrix of T obtained by eliminating all rows corresponding to the ID agents and all columns corresponding to the non-ID agents.

Since the beliefs of the IDs are always equal to one, the public beliefs $\tilde{x}$ evolve according to

$$\tilde{x}(k+1) = T_\psi \tilde{x}(k) + U_\psi \hat{x}(k)$$
$$= T_\psi \tilde{x}(k) + U_\psi \mathbb{1}.$$

The matrix T_ψ is a principal submatrix of the stochastic matrix A and therefore its eigenvalues reside in the open unit disk. Thus the above time evolution reaches steady-state $\tilde{x}_\infty$ for the public beliefs, $\tilde{x}_\infty = T_\psi \tilde{x}_\infty + U_\psi \mathbb{1}$. Therefore

$$\tilde{x}_\infty = (I - T_\psi)^{-1} U_\psi \mathbb{1}.$$

We are now in a position to formulate one of the main optimization problems considered in this work. We assume that limited resources available to the information disseminators only allow k_ψ agents to act as IDs. Furthermore, we assume that the overall level of 'public awareness' is measured by the sum of the beliefs of all individuals in society.[1] Defining the objective function as the steady-state

[1]Note that all elements of $\tilde{x}_\infty$ have nonnegative values.

level of public awareness, we have

$$J := \mathbb{1}^T \tilde{x}_\infty$$
$$= \mathbb{1}^T (I - T_\psi)^{-1} U_\psi \mathbb{1}.$$

Therefore, to maximize the steady-state value of public awareness using k_ψ information disseminating agents, we solve the problem

$$\begin{aligned} \text{minimize} \quad & -\mathbb{1}^T (I - T_\psi)^{-1} U_\psi \mathbb{1} \\ \text{subject to} \quad & \psi_i \in \{0, 1\}, \quad i = 1, \ldots, n - k_\phi \\ & \mathbb{1}^T \psi = k_\psi, \end{aligned} \quad (1)$$

where the optimization variable is the binary vector ψ.

In the following proposition we reformulate the optimization problem (1) in a way that lends itself to the application of the alternating direction method of multipliers (ADMM) algorithm [5], [6].

Proposition: The optimization problem (1) is equivalent to

$$\begin{aligned} \text{minimize} \quad & -\mathbb{1}^T (I - (I - D_\psi) T)^{-1} D_\psi \mathbb{1} \\ \text{subject to} \quad & \psi_i \in \{0, 1\}, \quad i = 1, \ldots, n - k_\phi \\ & \mathbb{1}^T \psi = k_\psi, \end{aligned}$$

in which ψ is the optimization variable, and $D_\psi = \mathrm{diag}\{\psi\}$.

Proof: Without loss of generality, we assume that a reordering of the elements of the state vector x is performed, so that the new state vector can be partitioned as $\begin{bmatrix} \hat{x}^T & \tilde{x}^T \end{bmatrix}^T$. Correspondingly, the matrices T and D_ψ are partitioned as

$$T = \begin{bmatrix} R_\psi & Q_\psi \\ U_\psi & T_\psi \end{bmatrix}, \qquad D_\psi = \begin{bmatrix} I & 0 \\ 0 & 0 \end{bmatrix}.$$

We have

$$(I - (I - D_\psi) T)^{-1}$$
$$= \left(\begin{bmatrix} I & 0 \\ 0 & I \end{bmatrix} - \begin{bmatrix} 0 & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} R_\psi & Q_\psi \\ U_\psi & T_\psi \end{bmatrix} \right)^{-1}$$
$$= \begin{bmatrix} I & 0 \\ (I - T_\psi)^{-1} U_\psi & (I - T_\psi)^{-1} \end{bmatrix},$$

and thus

$$\mathbb{1}^T (I - (I - D_\psi) T)^{-1} D_\psi \mathbb{1}$$
$$= \mathbb{1}^T \begin{bmatrix} I & 0 \\ (I - T_\psi)^{-1} U_\psi & (I - T_\psi)^{-1} \end{bmatrix} \begin{bmatrix} I & 0 \\ 0 & 0 \end{bmatrix} \mathbb{1}$$
$$= \mathbb{1}^T (I - T_\psi)^{-1} U_\psi \mathbb{1} + k_\psi.$$

Since k_ψ is a fixed parameter, it is clear that the maximization of $\mathbb{1}^T (I - (I - D_\psi) T)^{-1} D_\psi \mathbb{1}$ is equivalent to the maximization of $\mathbb{1}^T (I - T_\psi)^{-1} U_\psi \mathbb{1}$. Therefore the optimization problem in the proposition statement is equivalent to that in (1). This completes the proof. ∎

Using the proposition, we reformulate (1) as

$$\text{minimize} \quad -\text{trace}(\mathbb{1}\,\mathbb{1}^T\,(I-(I-D_\psi)\,T)^{-1}\,D_\psi)$$
$$\text{subject to} \quad \psi_i \in \{0,1\}, \quad i=1,\dots,n-k_\phi \tag{2}$$
$$\mathbb{1}^T\,\psi = k_\psi,$$

where $D_\psi = \text{diag}\{\psi\}$.

Application of ADMM

We proceed as in [7]. We define the indicator function of the constraint set in (2) as

$$g(\psi) = \begin{cases} 0 & \text{if } \psi_i \in \{0,1\}, \; \mathbb{1}^T\,\psi = k_\psi \\ \infty & \text{otherwise} \end{cases}$$

and the objective function as

$$f(\psi) = -\text{trace}(\mathbb{1}\,\mathbb{1}^T\,(I-(I-D_\psi)\,T)^{-1}\,D_\psi).$$

Hence, (2) can be formulated as the unconstrained optimization problem

$$\text{minimize} \quad f(\psi) + g(\psi).$$

which, in order to make more suitable for application of the ADMM algorithm, we rewrite in the following form

$$\text{minimize} \quad f(\psi) + g(\xi)$$
$$\text{subject to} \quad \psi - \xi = 0,$$

with ξ being an auxiliary variable. We form the augmented Lagrangian [5]

$$\mathcal{L}(\psi,\xi,\lambda) = f(\psi) + g(\xi) + \lambda^T(\psi-\xi) + \frac{\rho}{2}\|\psi-\xi\|_2^2,$$

where the vector λ is a dual variable, and the positive scalar ρ is a penalty weight.

The ADMM algorithm now proceeds as follows [5]. Given ξ^0, λ^0, for $\kappa = 0,1,\dots$ we carry out the updates

$$\psi^{\kappa+1} := \underset{\psi}{\text{argmin}}\; \mathcal{L}(\psi,\xi^\kappa,\lambda^\kappa), \tag{3}$$

$$\xi^{\kappa+1} := \underset{\xi}{\text{argmin}}\; \mathcal{L}(\psi^{\kappa+1},\xi,\lambda^\kappa), \tag{4}$$

$$\lambda^{\kappa+1} := \lambda^\kappa + \rho(\psi^{\kappa+1} - \xi^{\kappa+1}),$$

until both of the conditions

$$\|\psi^{\kappa+1} - \xi^{\kappa+1}\|_2 \le \varepsilon, \qquad \|\xi^{\kappa+1} - \xi^\kappa\|_2 \le \varepsilon,$$

are satisfied.

We now elaborate on each of the minimization problems (3) and (4):

In general, problem (3) is nonconvex in ψ, but local minima of the augmented Lagrangian can be found using descent algorithms. It can be shown that

$$(\nabla\mathcal{L})_i = -(\mathbb{1}\,\mathbb{1}^T\,(I-E_\psi\,T)^{-1})_{ii}$$
$$+ (T\,(I-E_\psi\,T)^{-1}\,D_\psi\,\mathbb{1}\,\mathbb{1}^T\,(I-E_\psi\,T)^{-1})_{ii}$$
$$+ \lambda_i + \rho\,(\psi_i - \xi_i),$$

and

$$\nabla^2\mathcal{L} = 2\,\text{sym}\{(T\,(I-E_\psi\,T)^{-1})\circ(\mathbb{1}\,\mathbb{1}^T\,(I-E_\psi\,T)^{-1})^T\}$$
$$- 2\,\text{sym}\{(T\,(I-E_\psi\,T)^{-1})$$
$$\circ(T\,(I-E_\psi\,T)^{-1}D_\psi\,\mathbb{1}\,\mathbb{1}^T\,(I-E_\psi\,T)^{-1})^T\}$$
$$+ \rho\,I,$$

where

$$E_\psi = I - D_\psi,$$

the operator sym gives the symmetric part of its matrix argument,

$$\text{sym}(W) = (1/2)\,(W + W^T),$$

and $\circ$ denotes elementwise matrix multiplication.

Problem (4) can be solved analytically; we present the solution here and refer the reader to [7] for details of its derivation. The optimal ξ solving (4) is found from

$$\xi_i = \begin{cases} 1 & \text{if } \bar{\xi}_i \ge [\bar{\xi}]_k \\ 0 & \text{if } \bar{\xi}_i < [\bar{\xi}]_k \end{cases}$$

where

$$\bar{\xi} := \psi^{\kappa+1} + \frac{1}{\rho}\lambda^\kappa,$$

and $[\bar{\xi}]_k$ is the kth largest entry of $\bar{\xi}$.

III. Increasing the Influence of Information Disseminators

In this section we assume that after choosing k_ψ nodes to serve as information disseminators, we now want to invest in increasing their influence on public opinion. We model this as creating new links (respectively, strengthening existing links) between the IDs and the public.

Let us recall that in the framework considered in this work, each agent updates his/her belief by taking a weighted average over the beliefs of those he/she interacts with. Therefore, if an ID wants to create a link between itself and an agent it did not interact with previously (or if an ID wants to increase its influence on an agent it is already interacting with), it would mean that the agent has to 'make room' (or 'make more room') in his/her weighted averaging scheme for the belief of the ID.

To model this we first notice that since all IDs are identical (i.e., all have belief value equal to one at all times), from a mathematical point of view it does not make any difference which ID establishes a link with a given agent.[2] Therefore, we proceed by finding merely the weight of the new links between public agents and the IDs (or the added weight between agents and IDs already interacting with them), without specifying exactly which ID each agent will be linked to. Once the optimal weights are found, one can use metrics such as physical proximity or availability to decide how to link (or strengthen existing links between) IDs and agents.

[2]Of course, from a practical point of view, considerations such as physical proximity would suggest that for a given agent the closest ID should be the one to initiate social interactions.

Suppose, for example, that a new social link is to be established between an agent and an ID. And suppose that the agent is to give a 10% weight in his/her averaging scheme to the belief of the ID. Since the agent's existing weights sum to one, he/she makes room for the ID's belief by scaling all his/her existing weights by 0.9. (If a social link already exists between an agent and an ID, the procedure stays the same in the sense that all weights, including that corresponding to the ID, are scaled down to make room for the added weight.)

The description above corresponds to the following evolution of beliefs

$$\begin{bmatrix} \hat{x}(k+1) \\ \tilde{x}(k+1) \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ S_\sigma U_\psi + U_\sigma & S_\sigma T_\psi \end{bmatrix} \begin{bmatrix} \hat{x}(k) \\ \tilde{x}(k) \end{bmatrix} + \begin{bmatrix} \mathbb{1} \\ 0 \end{bmatrix},$$

where

$$\sigma = \begin{bmatrix} \sigma_1 & \cdots & \sigma_m \end{bmatrix}^T, \quad 0 \leq \sigma_i \leq 1 \text{ for } i = 1, \ldots, m,$$

σ_i = weight agent i puts on the belief of any ID,

U_σ = matrix whose only nonzero entry in its ith row is σ_i,

$$S_\sigma = I - \text{diag}\{\sigma\},$$

and $m = n - k_\phi - k_\psi$ is the number of public (non-MI non-ID) agents. Note that the exact location in which σ_i appears in the ith row of U_σ is not important to us, as we only care about the weight the ith agent puts on the belief of *some* ID (i.e., the index of the ID is not important).

Since the beliefs of the IDs are always equal to one, the public beliefs $\tilde{x}$ evolve according to

$$\begin{aligned} \tilde{x}(k+1) &= S_\sigma T_\psi \tilde{x}(k) + (S_\sigma U_\psi + U_\sigma) \mathbb{1} \\ &= S_\sigma T_\psi \tilde{x}(k) + S_\sigma U_\psi \mathbb{1} + \sigma. \end{aligned}$$

The above time evolution reaches steady-state $\tilde{x}_\infty$ for the public beliefs, $\tilde{x}_\infty = S_\sigma T_\psi \tilde{x}_\infty + S_\sigma U_\psi \mathbb{1} + \sigma$. Therefore

$$\tilde{x}_\infty = (I - S_\sigma T_\psi)^{-1} (S_\sigma U_\psi \mathbb{1} + \sigma).$$

Note that if $\sigma_i = 1$ then the ith node effectively becomes an ID itself, since it fully ignores the beliefs of others (including his/her own beliefs) and adopts that of the IDs at all times.

We choose the objective function of our optimization problem to account for both the steady-state level of public awareness, $\mathbb{1}^T \tilde{x}_\infty$, which we wish to maximize, and the total weight of the links made to the IDs, $\mathbb{1}^T \sigma$, which we would like to keep small. Thus to maximize steady-state public awareness while keeping down the resources spent by the IDs, we solve the problem

$$\begin{aligned} \text{minimize} \quad & -\mathbb{1}^T (I - S_\sigma T_\psi)^{-1} (S_\sigma U_\psi \mathbb{1} + \sigma) + \gamma \mathbb{1}^T \sigma \\ \text{subject to} \quad & 0 \leq \sigma_i \leq 1, \quad i = 1, \ldots, m, \end{aligned}$$

(5)

where the optimization variable is the vector of weights σ, with $S_\sigma = I - \text{diag}\{\sigma\}$ and γ a positive scalar that determines the relative importance of the two terms in the objective function.

Customized Interior Point Method

Consider the following approximation of (5) using log-barrier functions [8]

$$\begin{aligned} \text{minimize} \quad & -\mathbb{1}^T (I - S_\sigma T_\psi)^{-1} (S_\sigma U_\psi \mathbb{1} + \sigma) + \gamma \mathbb{1}^T \sigma \\ & + \frac{1}{r} \sum_{i=1}^{n} (-\log(\sigma_i) - \log(1 - \sigma_i)) \end{aligned}$$

(6)

A sequence of optimization problems (6) are solved in which r is gradually increased in every step, and at each step the numerical scheme (e.g., Newton's method) is initialized with the solution of the previous iteration.

Defining $\mathcal{J}$ as the objective function in (6), we have

$$\begin{aligned} (\nabla \mathcal{J})_i = \; & (T_\psi (I - S_\sigma T_\psi)^{-1} (S_\sigma U_\psi \Pi + D_\sigma \mathbb{1} \mathbb{1}^T) \\ & \qquad\qquad\qquad\qquad (I - S_\sigma T_\psi)^{-1})_{ii} \\ & + ((U_\psi \Pi - \mathbb{1} \mathbb{1}^T)(I - S_\sigma T_\psi)^{-1})_{ii} \\ & + \gamma - (1/r)(\sigma_i^{-1} + (\sigma_i - 1)^{-1}), \end{aligned}$$

and

$$\begin{aligned} \nabla^2 \mathcal{J} = \; & -2 \, \text{sym}\{(T_\psi (I - S_\sigma T_\psi)^{-1}) \circ (T_\psi (I - S_\sigma T_\psi)^{-1} \\ & \qquad (S_\sigma U_\psi \Pi + D_\sigma \mathbb{1} \mathbb{1}^T)(I - S_\sigma T_\psi)^{-1})^T\} \\ & - 2 \, \text{sym}\{(T_\psi (I - S_\sigma T_\psi)^{-1}) \circ ((U_\psi \Pi - \mathbb{1} \mathbb{1}^T) \\ & \qquad\qquad\qquad\qquad (I - S_\sigma T_\psi)^{-1})^T\} \\ & + (1/r) \, \text{diag}\{\sigma_i^{-2} + (\sigma_i - 1)^{-2}\}, \end{aligned}$$

where Π is the $k_\psi \times m$ matrix of all ones, and

$$D_\sigma = \text{diag}\{\sigma\}.$$

IV. EXAMPLES

In this section, we provide four examples: the first two illustrate the problem of selecting IDs (cf. Sec. II), and the other two illustrate the problem of increasing the influence of IDs (cf. Sec. III).

Examples I & II: Selection of IDs

We consider the graph of the social network given in [4] and shown here in Fig. 1; we ignore the coloring of the nodes for the time being. We assume that certain nodes have been determined *a priori* to serve as MIs, and we seek the optimal locations for a given number of ID nodes.

Let M denote the (symmetric) adjacency matrix of the graph, i.e., let M_{ij} take the value 1 if node i is connected to node j and the value 0 otherwise. We also assume that each node has a self-loop (not shown in the figure) and thus $M_{ii} = 1$ for $i = 1, \ldots, n$. Let d_i denote the number of neighbors of node i, $d_i = \sum_{j=1}^{n} M_{ij}$.[3] We define ijth entry of the social interaction matrix A as

$$A_{ij} = \frac{M_{ij}}{d_i}.$$

(7)

Clearly A is a stochastic matrix.

Figures 1 and 2 demonstrate two scenarios with one and

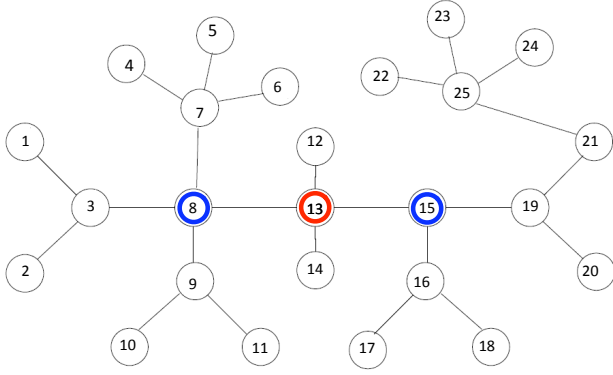[3]Since $M_{ii} = 1$, every node is a neighbor to itself.

Fig. 1: The social network given in [4]. In Example I, when node 13 is determined *a priori* to act as an MI, then the optimal locations for two IDs are nodes 8 and 15.
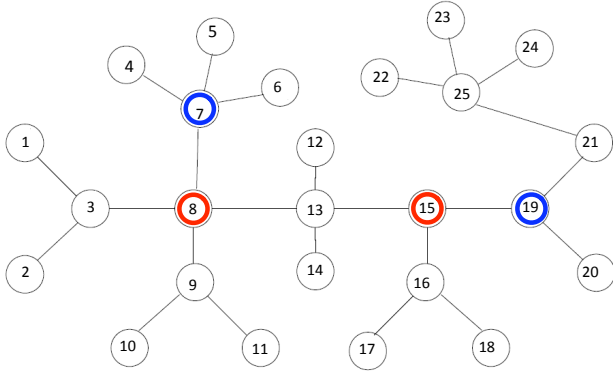

Fig. 2: In Example II, when nodes 8 and 15 are determined *a priori* to act as MIs, then the optimal locations for two IDs are nodes 7 and 19.

two MI agents (colored in red), respectively, and two ID agents (colored in blue). In both examples the solution was found by applying the ADMM algorithm to the optimization problem (2) of Sec. II. We note that the solution found in both cases is in agreement with that obtained using an exhaustive search, which implies that the ADMM algorithm (with $\rho = 10^4$) has indeed found the global minimum. Furthermore, the solutions make sense from an intuitive point of view; the IDs do their best to 'isolate' the MI node(s) and limit the spread of misinformation to as few nodes as possible.

We make a few notes and comparisons with respect to the existing literature on social networks. The papers [2], [3] characterize the social influence of agents in terms of the entries of the left eigenvector v of A corresponding to the eigenvalue one, $v^T A = v^T$; the magnitude of the positive

number v_i shows how much influence agent i has on public opinion. Indeed, for the special class of graphs described by (7), where agents use the average value of their neighbors' beliefs to update their own, we have [2]

$$v_i = \frac{d_i}{\sum_{i=1}^n d_i}.$$

This means that the influence of each agent is proportional to his/her number of neighbors. Based on this result, one would expect nodes $7, 8, 13, 25$ in Fig. 1, 2 to make good MIs and/or IDs, as they are the nodes with the largest number of neighbors. However, note that once some nodes are assigned to be MIs then the social interaction matrix A no longer satisfies (7) and the given formula for v_i is invalid. Furthermore, computing v directly using $v^T A = v^T$ for a network with an MI agent (or any forceful agent that refuses to update his/her belief) at node l renders v equal to a vector whose only nonzero element appears in location l. This means that only the MI agent is socially influential. Therefore, the eigenvector approach does not give a good indication of which nodes would serve as powerful IDs in the presence of MIs. On the other hand, the formulation proposed in this work for choosing IDs is general enough to be applied to social networks that contain MI agents.

Examples III & IV: Increasing Influence of IDs

We now consider the scenario in which both the MI and ID agents have already been selected and it is desired to find the optimal σ to create new links (or strengthen existing links) between the IDs and the public, such that the influence of the ID agents is maximally increased. The solutions are found by applying the customized interior point method in (6) to the optimization problem (5) of Sec. III.

Consider the MI and ID configuration of Fig.1, where node 13 is an MI and nodes 8 and 15 are IDs. For $\gamma = 1$ the optimal value of the vector σ obtained from solving (5) is given by

$$\sigma_i = \begin{cases} 1 & \text{if } i = 12, 14, \\ 0 & \text{otherwise.} \end{cases}$$

This result makes sense, since the only nodes that are not already isolated from the MI by the IDs are nodes 12 and 14. For the obtained value of σ, the IDs aggressively influence the opinion of nodes 12 and 14 by choosing $\sigma_{12} = \sigma_{14} = 1$.

As another example, consider the MI and ID configuration of Fig. 3, where node 13 is an MI and nodes 7 and 19 are IDs. Fig. 4 demonstrates the entries of the optimal σ for $\gamma = 0.1$ obtained from solving (5). Once again, the result makes sense, since the IDs attempt to aggressively influence the opinion of those nodes through which the MI communicates its misinformation to the rest of the network.

Not surprisingly, the optimal σ changes for different values of γ in (5). Fig. 5, 6 demonstrate the entries of the optimal σ for $\gamma = 1$ and $\gamma = 10$. Finally, it is interesting to note that, while the IDs become less aggressive as γ increases, the sparsity of σ does not change significantly as γ varies.
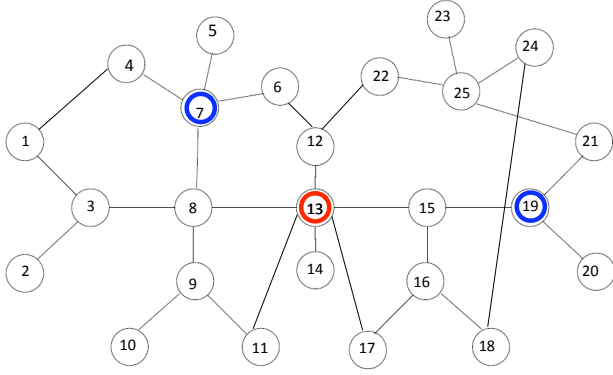
Fig. 3: In Example IV, node 13 is an MI and nodes 7 and 19 are IDs. For different values of γ, the entries of the optimal σ vector are given in Fig. 5–6.
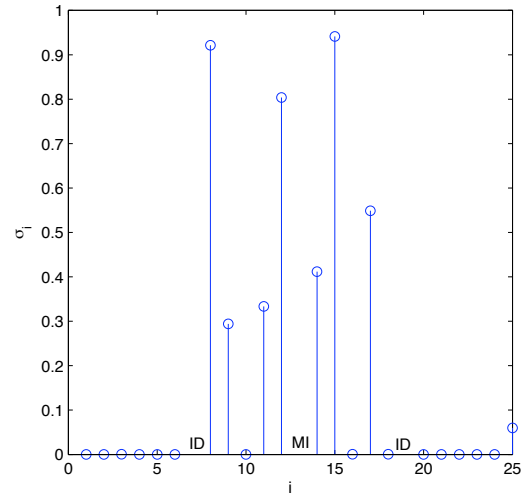

Fig. 5: The entries of the vector σ for the social network in Fig. 3, obtained by solving (5) for $\gamma = 1$.
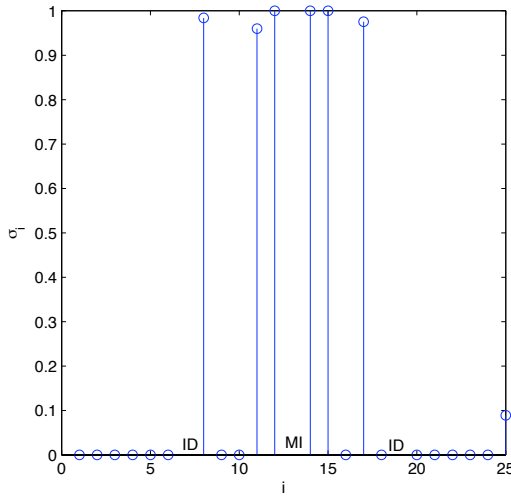

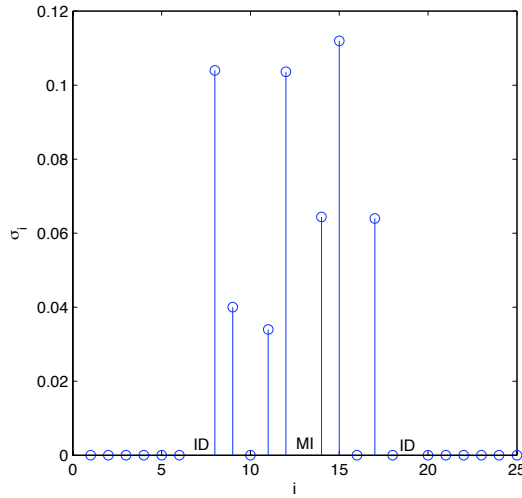Fig. 4: The entries of the vector σ for the social network in Fig. 3, obtained by solving (5) for $\gamma = 0.1$.

Fig. 6: The entries of the vector σ for the social network in Fig. 3, obtained by solving (5) for $\gamma = 10$.

V. CONCLUSIONS AND FUTURE WORK

We consider two problems: First, given a social network with misinforming agents, we find nodes that would best serve as information disseminators to the public. Second, once the best information disseminating nodes are found, we consider the problem of optimally increasing their influence on public opinion. The first problem is combinatorial in nature, and we demonstrate that the alternating direction method of multipliers is well-suited for addressing it.

Future work in this area would include deriving new objective functions that account for not only steady-state behavior but also the transient characteristics and convergence rates of public opinion.

REFERENCES

[1] D. Acemoglu, A. Ozdaglar, and A. ParandehGheibi, "Spread of (mis)information in social networks," *Games and Economic Behavior*, vol. 70, pp. 194–227, 2010.

[2] P. M. DeMarzo, D. Vayanos, and J. Zwiebel, "Persuasion bias, social influence, and unidimensional opinions," *The Quarterly Journal of Economics*, vol. 118, no. 3, pp. 909–968, 2003.

[3] B. Golub and M. Jackson, "Naive learning social networks and the wisdom of crowds," *American Economic Journal: Microeconomics*, vol. 2, no. 1, pp. 112–149, 2010.

[4] S. Patterson and B. Bamieh, "Leader selection for optimal network coherence," in *Proceedings of the 49th IEEE Conference on Decision and Control*, 2010, pp. 2692–2697.

[5] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.

[6] F. Lin, M. Fardad, and M. R. Jovanović, "Design of optimal sparse feedback gains via the alternating direction method of multipliers," *IEEE Trans. Automat. Control*, 2011, submitted; also arXiv:1111.6188v2.

[7] F. Lin, M. Fardad, and M. R. Jovanović, "Algorithms for leader selection in large dynamical networks: Noise-corrupted leaders," in *Proceedings of the 50th IEEE Conference on Decision and Control*, 2011, pp. 2932–2937.

[8] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.