

Identifying Influential Agents for Advertising in Multi-agent Markets

Mahsa Maghami
Department of EECS
University of Central Florida
Orlando, FL, USA
mmaghami@cs.ucf.edu

Gita Sukthankar
Department of EECS
University of Central Florida
Orlando, FL, USA
gitars@eecs.ucf.edu

ABSTRACT

The question of how to influence people in a large social system is a perennial problem in marketing, politics, and publishing. It differs from more personal inter-agent interactions that occur in negotiation and argumentation since network structure and group membership often play a more significant role than the content of what is being said, making the messenger more important than the message. In this paper, we propose a new method for propagating information through a social system and demonstrate how it can be used to develop a product advertisement strategy in a simulated market. We consider the desire of agents toward purchasing an item as a random variable and solve the influence maximization problem in steady state using an optimization method to assign the advertisement of available products to appropriate messenger agents. Our market simulation accounts for the 1) effects of group membership on agent attitudes 2) has a network structure that is similar to realistic human systems 3) models inter-product preference correlations that can be learned from market data. The results show that our method is significantly better than network analysis methods based on centrality measures.

Categories and Subject Descriptors

I.2.11 [Distributed artificial intelligence]: Multi-agent systems

General Terms

Algorithms

Keywords

Marketing, Optimization, Multi-agent social simulations

1. INTRODUCTION

The gift of persuasion is a powerful and highly-sought after skill, as evidenced by the fact that individual self-help books in this area, the most famous being *How to Win Friends and Influence People* published in 1936, remain popular. The rise of social media outlets and click-through advertisement opened the door for relatively small groups to influence large numbers of people. Combined with modern data analysis techniques, it is possible to create a detailed social simulation of the population of interest, but the problem of

whom to influence remains as an open research question. Particularly in advertisement, indiscriminate mass marketing techniques can lead to negative information cascades about product quality, even if cost efficiency is not an issue. This problem can be framed as a network influence propagation problem; previous work in this area has looked at diverse domains such as information propagation in the Flickr social network [7] and identifying important blogs for marketing [3].

In this paper we present a mathematical analysis of how influence propagation occurs over time and propose a new optimization technique for identifying effective messenger agents in the network that outperforms other network analysis methods while accounting for realistic factors such as group membership and product preference correlation. Following the work of Hung et al. [12, 13], optimization is used along with an analysis of the expected long-term system behavior to assign the advertisement of the available products to appropriate agents in the network. In contrast with previous work on identifying influential nodes for marketing purposes (e.g., [11] and [4]), in this work we model the effects of realistic social factors such as group membership on product adoption. In the analysis presented in [12, 13] for counterinsurgency messaging tactics, there exists a single random variable representing the attitude of agents toward counterinsurgency, but in our work, we use a vector of random variables which represents the desire of each agent toward any single product. This consideration combined with product demand correlations in the market make the analysis and optimization more complicated, but ultimately our approach has the promise of being applicable to a wider variety of social systems.

The paper is organized as follows. Section 2 describes the agent-based model and how it can be used for influencing propagation. Section 3 presents our network generation model that is designed to account for social influence factors present in real human societies such as homophily and group membership. The subsequent section presents our analysis of the system dynamics, and our proposed optimization technique is described in Section 5. We evaluate our method vs. a set of centrality based network analysis techniques in Section 6. We end the paper with an overview of related work in this area and a discussion of future work.

2. MARKET MODEL

To explore the efficiency of the proposed marketing method, we have extended a multi-agent system model, inspired by [12] and [13], to simulate a social system of potential customers. In this model, there is a population of N agents, represented by the set $A = \{a_1, \dots, a_N\}$, that consists of two types of agents ($A = A_R \cup A_P$). The first type of agent, defined as: $A_R = \{a_r \mid a_r \text{ is Mutable and } 1 \leq r \leq R\}$, are the *Regular* agents, who are the potential customers. These agents have a changing atti-

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

Copyright © 2012, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

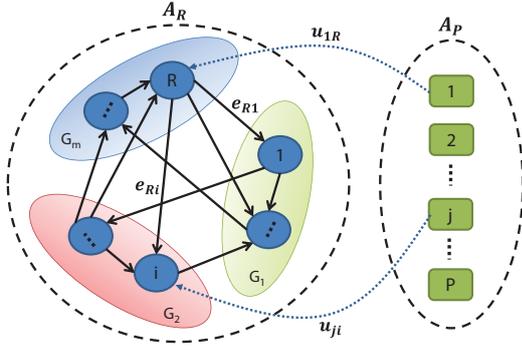


Figure 1: The model of the social system. There exist two types of agents, *Regular agents* (A_R) and *Product agents* (A_P). A static network exists among *Regular agents*, and our problem is to find effective connections between the *Product* (sellers) and *Regular agents* (customers) in order to influence the customers to buy products. *Regular agents* also can belong to different groups in their society (G_m), which modifies the local influence propagation properties.

tude on purchasing products and can be influenced by the *Product* agents who represent salespeople offering one specific product. These agents have an immutable attitude toward a specific product and are defined as: $A_P = \{a_p \mid a_i \text{ is Immutable and } 1 \leq p \leq P\}$. Figure 1 provides an illustration of the market model.

Each *Regular agent* can be considered as a unique node in the social network, connected by directed weighted links based on the underlying interactions with other agents. The connection between the *Regular agents* is modeled by an adjacency matrix, \mathbf{E} , where $e_{ij} = 1$ is the weight of a directed edge from agent a_i to agent a_j . The in-node and out-node degrees of agent a_i are the sum of all in-node and out-node weights, respectively ($d_{in}^i = \sum_{a_j \subseteq A_R} e_{ji}$ and $d_{out}^i = \sum_{a_j \subseteq A_R} e_{ij}$). This network is assumed to follow a power law degree distribution like many human networks, and is generated synthetically as explained in Section 3.

We model the desire of an agent, a_i , to buy an item or consume a specific product, p , as a random variable denoted by $x_{ip} \in [-1, 1]$. As there exist P items in the environment, each agent is assigned a vector of random variables, \vec{X}_i , representing the attitude or desire of the agent toward all of the products in the market.

Within the social network there are different groups of *Regular agents*; these groups could represent demographic groups or other types of subcultures. Agents from the same group are more effective at influencing each other. To model this, the social system contains m different long-lasting groups, G_1, \dots, G_m , and each agent i is designated with a group membership, G_i .

Here, we do not attempt to capture a rich social-cultural behavior model of these interactions, but rather view the model simply as a function $\mathcal{F} : G_i \rightarrow S_i$, mapping the group label of agents, G_i , to a social impression, S_i , that affects link formation and influence propagation, which we designate as the group value judgment. This value represents the agents' judgments on other groups and is based on observable group label of the agent rather than real characteristics of the person. We assume that the impression of different groups has been learnt by agents beforehand therefore each agent has a unique vector of judgment values, noted as $\vec{S}_i = S_1, S_2, \dots, S_m$, to indicate the judgment of each agent on different groups in the simulated society.

Moreover, in real life there is a correlation between the user de-

mand of different products in the market. The desire of customers for a specific product is related to his/her desire toward other similar products. To model this correlation and consider its effect in our formulation, we designate a matrix \mathbf{M} that identifies the relationship between demands among advertised items and can be shown as:

$$\mathbf{M} = \begin{pmatrix} m_{11} & \dots & m_{1P} \\ \vdots & \ddots & \vdots \\ m_{P1} & \dots & m_{PP} \end{pmatrix}$$

where m_{ij} indicates the probability of having desire toward item j assuming the agent already has a desire for item i . We assume that this matrix is known beforehand and has been modeled by the advertisement companies by tracking the users and applying user modeling.

In the market, the companies are trying to select a set of connections between the A_P agents and A_R agents, in such a way to maximize the long term desire of the agents for the products. We define a simple decision variable u_{ji} , where

$$u_{ji} = \begin{cases} 1 & \text{Product } j \text{ connects to Regular agent } i, \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

Note that the links between *Product agents* and *Regular agents* are directed links from products to agents and not in the opposite direction, and that *Product agents* will never connect to other *Product agents*. In the social simulation, each agent interacts with another agent in a pair-wise fashion that is modeled as a Poisson process with rate 1, independent of all other agents. By assuming a Poisson process of interaction, we are claiming that there is at most one interaction at any given time. Here, the probability of interaction between agents a_i and a_j is shown by p_{ij} and is defined as a fraction of the connection weight between these agents over the total connections that agent i makes with the other agents. Therefore,

$$p_{ij} = \begin{cases} \frac{e_{ij}}{d_{out}^i} & i, j \in A_R \\ \frac{u_{ji}}{Threshold} & i \in A_R, j \in A_P \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

where d_{out}^i is the out-node degree of a *Regular agent* i and the *Threshold* parameter is the total number of links that *Product agent* can make with *Regular agents*. The bounds on *Threshold* are a natural consequence of the limited budget of companies in advertising their products.

At each interaction there is a chance for agents to influence each other and change their desire vector for purchasing or consuming a product. In all these interactions *Product agents*, the immutable agents, are the only agents who do not change their attitude and have a fixed desire vector. The probability that agent j influences agent i is denoted as α_{ij} and is calculated based on the out-node degree of agent j as:

$$\alpha_{ij} = \begin{cases} \frac{e_{ji}}{d_{out}^j} & i, j \in A_R \\ cte & i \in A_R, j \in A_P \end{cases} \quad (3)$$

Figure 2 shows a simple example of how to calculate p_{ij} and α_{ij} .

The other important parameter in the agent influence process is ϵ_{ij} , which determines how much agent j will influence agent i . This parameter is derived from a Gaussian distribution assigned to the membership group of agent j based on the experience of agent i with this group. Therefore, this value can easily be extracted from the previously defined vector \vec{S}_i .

As a final note, in this model the agents can access the following information:

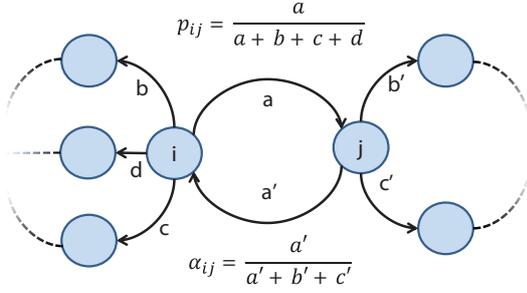


Figure 2: An illustration of how the probability of interaction (p) and the probability of influencing others (α) is calculated between the *Regular* agents.

1. the links connecting agents that possess a history of past interactions. Each agent is aware of its connections with neighbors and their weights;
2. the group membership of neighboring agents and other select members of the community.

The ultimate goal of our marketing problem is to recognize the influential agents in the graph and define u_{jis} in a way to get the maximum benefit of the product advertising.

3. NETWORK GENERATION

To evaluate the performance of our proposed optimization method on identifying influential agents in a variety of networks, we simulate the creation of agent networks formed by the combined forces of homophily and group membership. Since social communities often form a scale-free network, whose degree distribution follows a power law [5], we model our agent networks using the network generation method described in [25]. Note that this network only connects the regular agents ($a_i \in A_R$). The connection between the *Product* and *Regular* agents is identified later in a way to optimize the efficiency of the product marketing.

Following the network data generation method in [23], we control the link density of the network using a parameter, ld , and value homophily between agents using a parameter, dh . The effects of value homophily are simulated as follows:

1. At each step, a link is either added between two existing nodes or a new node is created based on the link density parameter (ld). In general, linking existing nodes results in a higher average degree than adding a new node.
2. To add a new link, first, we randomly select a node as the source node, a_i , and a sink node, a_j ($a_i, a_j \in A_R$), based on the homophily value (dh), which governs the propensity of nodes with similar group memberships to link. Node a_j is selected among all the candidate nodes in the correct group, based on the degree of the node. Nodes with higher degree have a higher chance to be selected.
3. If a prior link exists between agent a_i and a_j , selecting them for link formation will increase the weight of their link by one.

Group membership also governs the process of reciprocal link formation. Once the link generation process starts and the source and sink nodes have been selected, we add a directed link from node a_i to node a_j by default, under the assumption that the first selected agent initiated the interaction. The group value judgment

Table 1: Agent Network Generator

Agent Network Generator ($numNodes, numLabels, ld, dh$)

```

i = 0
E = NULL
while i < numNodes do
  sample r from uniform distribution U(0, 1)
  if r ≤ ld then
    connectNode(E, numLabels, dh)
  else
    addNodes(E, numLabels, dh)
    i = i + 1
  end if
end while
return E

```

of the second node governs whether a reciprocal link is formed or not. We use an evaluation function $F_a(S)$ to map an observed group value S to a binary evaluation of interaction (positive or negative). We assume that all agents use the same evaluation function, which is:

$$F_a(S) = \begin{cases} 1 & S \geq 0.5 \\ -1 & S < 0.5 \end{cases}$$

The result of this process is to create clusters of agents with the same group labels within the network, since group membership affects both the probability of the initial interaction (through the homophily parameter) and also the reciprocal link formation.

To generate a new node, we first select a group label based on a uniform group distribution and assign that group label to the node. Then we add links between the new node and one of the existing nodes as we described above. The algorithm for generating the static network is outlined in Table 1.

4. ANALYSIS OF SYSTEM DYNAMICS

As explained in Section 2, the agent i 's desire toward product p , is modeled as a random variable that assumes a scalar value after each interaction ($x_{ip} \in [-1, 1]$). Therefore, since there exist P different products, each agent has a vector of random variables, \vec{X}_i , which indicates the desire of the agent toward all the available products in market. Following Hung et al. [12, 13], we model the desire dynamic of all agents as a Markov chain where the state of the system is a matrix of all agents' desire vectors at a particular iteration k and the state transitions are calculated probabilistically from the pair-wise interaction between agents connected in a network. The state of the system at the k^{th} iteration is a vector of random variables, denoted as $\mathbf{X}(k) \in \mathbb{R}^{NP \times 1}$ (created through a concatenation of N vectors of size P) and expressed as:

$$\mathbf{X}(k) = \begin{pmatrix} [\vec{X}_1(k)] \\ \vdots \\ [\vec{X}_N(k)] \end{pmatrix}$$

4.1 Interaction and Influence

In this work, we define interactions as any kind of information or belief sharing between two agents about the available products in the market. During these interactions, there is a possibility for one agent to influence the desire of the other one. As explained in Section 2, this possibility is modeled by parameter α_{ij} when agent i initiates the interaction with agent j . Also, in this interaction, we assume that the influenced agent will retain some fraction of

its existing desire. This fraction is different for any single agent i while interacting with agent j , but remains fixed, and is denoted as $\varepsilon_{ij} \in [0, 1]$. The dynamics of the model at each iteration k proceed as described in [13]:

1. Agent i initiates the interaction on a uniform probability distribution over all agents. Then agent i selects another agent among its neighbors with probability p_{ij} . Note that the desire dynamic can occur with probability $\frac{1}{N}(p_{ij} + p_{ji})$ as agent i 's attitude can change whether it initiates the interaction or is selected by agent j .

2. Conditioned on the interaction of i and j :

- With propagability α_{ij} , agent i will change its desire:

$$\begin{cases} \vec{X}_i(k+1) = \varepsilon_{ij} \mathbf{M} \vec{X}_i(k) + (1 - \varepsilon_{ij}) \mathbf{M} \vec{X}_j(k) \\ \vec{X}_j(k+1) = \vec{X}_j(k) \end{cases} \quad (4)$$

Recall that \mathbf{M} is the pre-defined matrix indicating the correlation between the demands of different products.

- With probability of $(1 - \alpha_{ij})$, agent i is not influenced by the other agent:

$$\begin{cases} \vec{X}_i(k+1) = \vec{X}_i(k) \\ \vec{X}_j(k+1) = \vec{X}_j(k) \end{cases} \quad (5)$$

To analyze Equation 4 in detail, we rewrite the matrix calculation for agent i as follows:

$$\vec{X}_i(k+1) = \begin{pmatrix} \sum_{f=1}^P m_{1f} (\varepsilon_{ij} x_{if} + (1 - \varepsilon_{ij}) x_{jf}) \\ \vdots \\ \sum_{f=1}^P m_{Pf} (\varepsilon_{ij} x_{if} + (1 - \varepsilon_{ij}) x_{jf}) \end{pmatrix} \quad (6)$$

A closer look at each row of $(\vec{X}_i(k+1))$ reveals that the desire of agent i toward a product depends on own previous desire, a fraction of the other agent's desire toward that product, and the desire of both agents toward other available products in the market. This is an interesting result showing how our proposed model can express the complexity of real-world markets and capture the dependency of demand for different products [20].

4.2 Expected Long-term Desire

In this work, we determine the long-term desire of the agents for products in the system to find the optimized connection between the *Product* agents and *Regular* agents. In other words, we hypothesize that by examining the expected value of the steady state system $(\mathbf{X}(k))$, we are able to optimize the marketing strategy and identify the most influential nodes in the network. Therefore our goal in this section is to calculate the expectation vector of the system state since it captures all the interactions and the dependencies between the demand of the products.

The conditional expected value of the desire vector of agent i in a single pair-wise interaction between agents i and j , when the current state of the system is observed:

$$\begin{aligned} E[\vec{X}_i(k+1)|\mathbf{X}(k), j] &= (1 - \alpha_{ij}) \vec{X}_i(k) \\ &+ \alpha_{ij} [\varepsilon_{ij} \mathbf{M} \vec{X}_i(k) + (1 - \varepsilon_{ij}) \mathbf{M} \vec{X}_j(k)] \\ &= [\alpha_{ij} \varepsilon_{ij} \mathbf{M} + (1 - \alpha_{ij}) \mathbf{I}] \vec{X}_i(k) \\ &+ \alpha_{ij} (1 - \varepsilon_{ij}) \mathbf{M} \vec{X}_j(k) \end{aligned} \quad (7)$$

By defining matrix $\mathbf{W}(i, j) = \alpha_{ij}(1 - \varepsilon_{ij}) \mathbf{M}$, we rewrite Equation 7 in the form of:

$$\begin{aligned} E[\vec{X}_i(k+1)|\mathbf{X}(k), j] &= \vec{X}_i(k) + \mathbf{W}(i, j) \vec{X}_j(k) \\ &- [\mathbf{W}(i, j) + \alpha_{ij}(\mathbf{I} - \mathbf{M})] \vec{X}_i(k) \end{aligned} \quad (8)$$

Therefore, based on the probability of interaction between two agents $(\frac{1}{N}(p_{ij} + p_{ji}))$, the desire of *Regular* agents dynamically changes as specified in Equation 7. It is worthwhile to mention that matrix \mathbf{W} is a factor of matrix \mathbf{M} , and it has the same dimensions of $P \times P$. Rewriting the dynamics of \vec{X}_i in this way indicates that the desire vector of agent i at iteration $(k+1)$ is equivalent to its own desire plus the weighted desire of agent j at iteration k , minus its own weighted desire at that iteration. This finding shows that, in spite of having the extra matrix \mathbf{M} , extracted from the marketing situation, and a complicated notion of the agents' desire vector, the computation model simply follows [12], although the optimization approach must account for multiple product interactions.

We substitute $\mathbf{W}(i, j) + \alpha_{ij}(\mathbf{I} - \mathbf{M}) = \mathbf{S}(i, j)$, where $\mathbf{S}(i, j)$ again is dimension $P \times P$. Then, Equation 8 simplified as follows:

$$E[\vec{X}_i(k+1)|\mathbf{X}(k), j] = \vec{X}_i(k) - \mathbf{S}(i, j) \vec{X}_i(k) + \mathbf{W}(i, j) \vec{X}_j(k) \quad (9)$$

Next, we write the expected value of agent i 's desire vector at iteration $(k+1)$ over all the possible interactions it initiates or is subject to by other agents' actions, conditioned on the state of the system at k . Recall that the interaction between i and j occurs with probability $\frac{1}{N}(p_{ij} + p_{ji})$.

$$\begin{aligned} E[\vec{X}_i(k+1)|\mathbf{X}(k)] &= \vec{X}_i(k) \\ &- \sum_j \frac{1}{N} (p_{ij} + p_{ji}) \mathbf{S}(i, j) \vec{X}_i(k) \\ &+ \sum_j \frac{1}{N} (p_{ij} + p_{ji}) \mathbf{W}(i, j) \vec{X}_j(k) \end{aligned} \quad (10)$$

Now, we want to express the expected desire of all agents at iteration $(k+1)$ conditioned on all agents' previous desire. This step relies on both the laws of interacting expectations and linearity of expectations. Assembling a vector of all entries for each i results in:

$$E[\mathbf{X}(k+1)|\mathbf{X}(k)] = \mathbf{X}(k) + \mathbf{Q} \mathbf{X}(k) \quad (11)$$

where \mathbf{Q} is a block matrix and each component of $\mathbf{Q} \in \mathbb{R}^{N \times N}$, considering Equation 10, is:

$$Q_{ij} = \begin{cases} \frac{1}{N} (p_{ij} + p_{ji}) \mathbf{W}(i, j) & i \in A_R, j \in A \text{ and } i \neq j \\ -\frac{1}{N} \sum_j (p_{ij} + p_{ji}) \mathbf{S}(i, j) & i \in A_R, j \in A \text{ and } i = j \\ +\frac{1}{N} (p_{ij} + p_{ji}) \mathbf{W}(i, j) & \\ 0 & i \in A_P, j \in A \end{cases} \quad (12)$$

Finally, by calculating the expected value of Equation 11 and using the linearity of expectations, we have:

$$E[E[\mathbf{X}(k+1)|\mathbf{X}(k)]] = E[\mathbf{X}(k+1)] = E[\mathbf{X}(k)] + \mathbf{Q} E[\mathbf{X}(k)] \quad (13)$$

We define $\vec{\mu}_{\mathbf{X}}(k) \in \mathbb{R}^{N \times 1}$ as the expected value vector of $\mathbf{X}(k)$. Therefore, the above equation is simplified as:

$$\vec{\mu}_{\mathbf{X}}(k+1) = \vec{\mu}_{\mathbf{X}}(k) + \mathbf{Q} \vec{\mu}_{\mathbf{X}}(k) \quad (14)$$

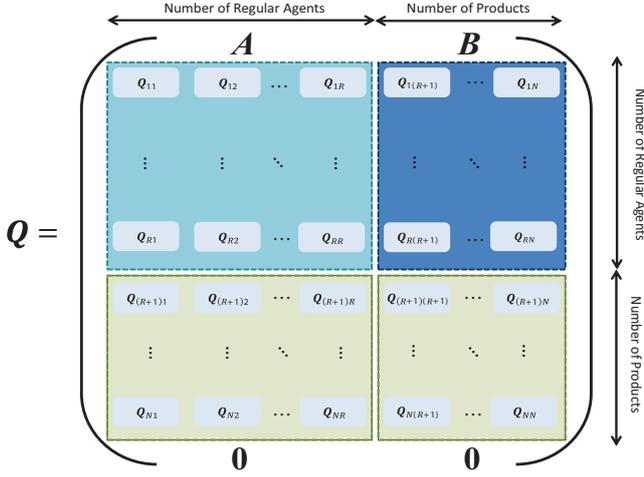


Figure 3: \mathbf{Q} matrix is a block matrix with size $N \times N$ where N is the total number of agents ($R + P$) and each block has the size of $P \times P$. Matrices \mathbf{A} and \mathbf{B} are the non-zero part of this matrix which represent the interactions among *Regular* agents and interactions between *Regular* agents and *Products*, respectively.

Since we are seeking the expected value of $\mathbf{X}(k)$ at steady state, the above equation when $k \rightarrow \infty$ reduces to:

$$\vec{\mu}_{\mathbf{X}}(\infty) = \vec{\mu}_{\mathbf{X}}(\infty) + \mathbf{Q} \vec{\mu}_{\mathbf{X}}(\infty) \Rightarrow \mathbf{Q} \vec{\mu}_{\mathbf{X}}(\infty) = 0 \quad (15)$$

In order to solve this system of equations efficiently, we decompose the matrices:

$$\mathbf{Q} = \begin{pmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{0} & \mathbf{0} \end{pmatrix} \text{ and } \vec{\mu}_{\mathbf{X}}(\infty) = \begin{pmatrix} \vec{\mu}_{\mathbf{R}} \\ \vec{\mu}_{\mathbf{P}} \end{pmatrix} \quad (16)$$

Here $\mathbf{A} \in \mathbb{R}^{RP \times RP}$ is the sub-matrix representing the expected interactions among *Regular* agents while $\mathbf{B} \in \mathbb{R}^{RP \times P^2}$ represents the the expected interactions between *Regular* agents and *Product* agents. Figure 3 shows the breakdown of matrix \mathbf{Q} .

Moreover, $\vec{\mu}_{\mathbf{R}}$ and $\vec{\mu}_{\mathbf{P}}$ are vectors representing the expected long-term desire of *Regular* agents and *Product* agents, respectively, at iteration $k \rightarrow \infty$. Note that vector $\vec{\mu}_{\mathbf{P}}$ is known since the *Product* agents, the advertisers, are the immutable agents, who never change their desire. Solving for $\vec{\mu}_{\mathbf{R}}$ yields the vector of expected long-term desire for all regular agents, for a given set of influence-probabilities on a deterministic social network.

$$\mathbf{A} \vec{\mu}_{\mathbf{R}} + \mathbf{B} \vec{\mu}_{\mathbf{P}} = 0 \Rightarrow \vec{\mu}_{\mathbf{R}} = \mathbf{A}^{-1}(-\mathbf{B} \vec{\mu}_{\mathbf{P}}) \quad (17)$$

Now based on this analytical view of the system, we define an optimization method in following section to maximize the product sales through intelligent selection of the *Product* agent linkages.

5. NODE SELECTION METHOD

Using the analysis from the previous section, we can identify the influential nodes in the network and connect the products to those agents in a way that maximizes the long-term desire of the agents in the social system. Here, we define the objective function as the maximization of the weighted average of the expected long-term desire of all the *Regular* agents in the network toward all the products as:

$$\max_u \sum_{1 \leq k \leq P} \sum_{i \in A_R} (\rho_i \cdot \vec{\mu}_{\mathbf{R},i}) \quad (18)$$

$\vec{\mu}_{\mathbf{R},i}$ is the part of $\vec{\mu}_{\mathbf{R}}$ that belongs to agent i , and ρ_i parameter is simply a weight we can assign to agents based on their importance in the network. In the case of equivalent $\rho_i = 1$ for all the agents, the above function reduces to the arithmetic mean of the expected long-term desire vectors for all agents.

The goal of our proposed method is to assign a fixed number of *Product* agents with limited number of connections to a network of *Regular* agents in a way to optimize the objective function presented above. In Equation 17, matrix \mathbf{A} and vector $\vec{\mu}_{\mathbf{P}}$ are known since the static network among the *Regular* agents and the fixed desire vector of the products are both known. We define the matrix \mathbf{B} based on parameters of u_{ij} s. We substitute the probability of interaction, p_{ij} , occurring between agents i and j in matrix \mathbf{Q} , by Equation 2 of the model.

The partitioning of matrix \mathbf{Q} in Equation 16 and the size of matrices \mathbf{A} and \mathbf{B} (Figure 3), indicates that the elements of matrix \mathbf{B} are all off the diagonal. Therefore substituting the values of p_{ij} and p_{ji} of Equation 2 into Equation 12, $\mathbf{B}_{ij} = \frac{1}{N} u_{ji} \mathbf{W}(i, j) = \hat{u} \otimes \mathbf{M}$. Here, \hat{u} contains all the variables and influence parameters and \otimes indicates the Kronecker product [21].

Therefore, by rewriting Equation 17 as:

$$\vec{\mu}_{\mathbf{R}} = \mathbf{A}^{-1}[\hat{u} \otimes \mathbf{M}] \text{Vec}(\hat{\mu}_{\mathbf{P}}) \quad (19)$$

and using the following identity

$$[\hat{u} \otimes \mathbf{M}] \text{Vec}(\hat{\mu}_{\mathbf{P}}) = \text{Vec}(\mathbf{M} \hat{\mu}_{\mathbf{P}} \hat{u}),$$

Equation 17 becomes $\vec{\mu}_{\mathbf{R}} = \mathbf{A}^{-1} \text{Vec}(\mathbf{M} \hat{\mu}_{\mathbf{P}} \hat{u})$, which is solved using convex optimization methods. Therefore the optimal assignment of *Product* agents to *Regular* agents is obtained through the following optimization problem:

$$\begin{aligned} & \underset{\hat{u}}{\text{maximize}} && \|\mathbf{A}^{-1} \text{Vec}(\mathbf{M} \hat{\mu}_{\mathbf{P}} \hat{u})\|_1 \\ & \text{subject to} && x_{ip} \in [-1, 1], \forall i \in A_R, \\ & && \sum_{j \in A_R} u_{ij} = \text{cte}. \end{aligned} \quad (20)$$

To solve this optimization problem we used the CVX toolbox of Matlab which is useful for convex programming and minimized the dual of our objective function.

6. EVALUATION

6.1 Experimental Setup

We conducted a set of simulation experiments to evaluate the effectiveness of our proposed node selection method on marketing the items in a simulated social system with a static network. The parameters of the model for all the runs are summarized in Table 2(a). All the results are computed over an average of 30 runs with 100 *Regular* agents and 10 *Product* agents.

In this work, we model four long-lasting groups, (G_1, \dots, G_4) , with different feature vector distributions in our social simulation. Moreover, a group value judgment, (S_i) , assigned to each group, is drawn from Gaussian distribution. We assumed that the group model has been learned by agents based on their previous experiences, each agent has its own fixed value judgment toward each group of agents and that value has been selected based on the assigned Gaussian distribution of the model. Consequently, this group value judgment affects the connection of agents during the network generation phase, as we described before. Table 2(b) shows the mean and standard deviations of the Gaussian distributions assigned to each group. Note that the membership in each group is

Table 2: Parameter settings
(a) Experimental parameters

Parameter	Value	Descriptions
R	100	Number of <i>Regular</i> agents
P	10	Number of <i>Product</i> agents
<i>Threshold</i>	2	Number of links between P and R agents
ε	0.4	Influence factor between P and R agents
α	0.6	Probability of influence between P and R agents
$N_{Iterations}$	10000	Number of iterations
N_{Run}	30	Number of runs

(b) Group model

Group	Mean Value	StDev
$G1$	0.9	0.05
$G2$	0.6	0.15
$G3$	0.4	0.15
$G4$	0.3	0.1

permanent for all agents and cannot be changed during the course of one simulation.

In the *Regular* and *Product* agent interaction, parameters α and ε are fixed for any interaction and are presented in Table 2(a). We assume that these parameters can be calculated by advertising companies based on user modeling. The p_{ij} values for this type of interaction are calculated using Equation 2 and are parametric.

Finally, the remaining part of the social system setup is matrix M , which models the correlation between the demand for different products. This matrix is generated uniformly with random numbers between $[0 \ 1]$ and, as it has a probabilistic interpretation, the sum of the values in each row, showing the total demand for one item, is equal to one.

6.2 Results

We compare our optimization-based algorithm with a set of centrality-based measures commonly used in social network analysis for identifying influential nodes based on network structure [15]. The comparison methods are:

Degree Assuming that high-degree nodes are influential nodes in the network is a standard approach for social network analysis. Here, we calculated the probability of joining a *Regular* agent based on the out-degree of the agents and attached the *Product* agents according to preferential attachment. Therefore, nodes with higher degree had an increased chance of being selected as an advertising target.

Closeness This is another commonly used influence measure in sociology, based on the assumption that a node with short paths to other nodes has a higher chance to influence them. Here, we averaged the shortest paths of a node to all the other nodes in the network and sorted the nodes according to this measure. Nodes with shorter average path had a higher chance of being selected as a target.

Betweenness This centrality metric measures the number of times a node appears on the geodesics connecting all the other nodes in the network. Nodes with the highest value of betweenness had the greatest probability of being selected.

Random Finally, we consider selecting the nodes uniformly at random as a baseline.

To evaluate these methods, we started the simulation with an initial desire vector set to 0.02 for all agents, and simulated 10000

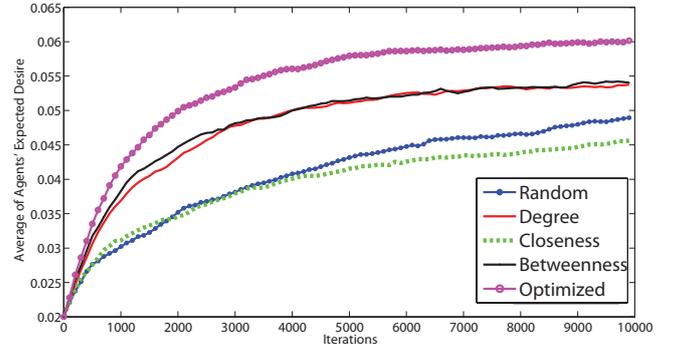


Figure 4: The average of agents' expected desire vs. the iterations. The average is across all the products and over 30 different runs. Our proposed method has the highest average in comparison to other methods which shows its capability as a method for targeted advertisement in a social system.

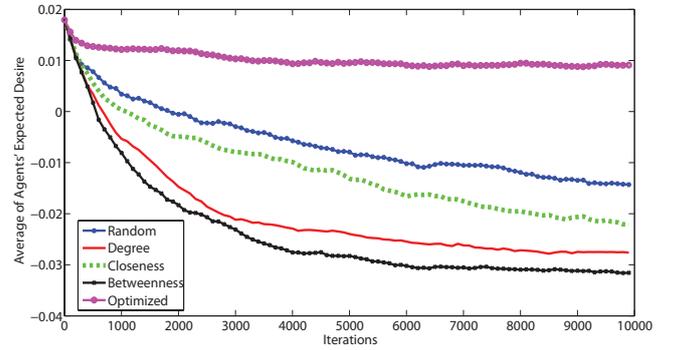


Figure 5: The average of agents' expected desire vs. iterations. In this simulation, the negative effect of advertising products against other products has been increased. This result demonstrates that our proposed method is more robust to the commonly occurring condition where increasing the desire toward one item has a higher negative effect on the desire of agent toward other products.

iterations of agent interactions. The entire process of interaction and influence is governed based on the previous formulas given in Section 4 and extracted parameters from the network. At each iteration, we calculated the average of the expected desire value of agents toward all products. Figure 4 shows this result for 100 agents and 10 advertisements. As explained before, the desire vector of *Product* agents are fixed for all products; in our simulation it was set to 1 for the product itself and -0.05 for all other products (e.g., $\mu_2 = [-0.05 \ 1 \ -0.05 \ \dots \ -0.05]$). The results for this condition show that the proposed method creates a higher total product desire in the social system and is more successful than other methods at selecting influential nodes. To test the robustness of our algorithm we modified the desire vector of *Product* agents and increased the negative effect of advertisements over other products by factor of three (e.g., $\mu_2 = [-0.15 \ 1 \ -0.15 \ \dots \ -0.15]$). The result of this simulation is shown in Figure 5. We can see that in this case the average desire of agents has dropped dramatically for all methods except the proposed algorithm. Even in the cases of having high negative effect toward other products, this algorithm can adapt the node selection in a way to keep the desire of agents high and sell more products.

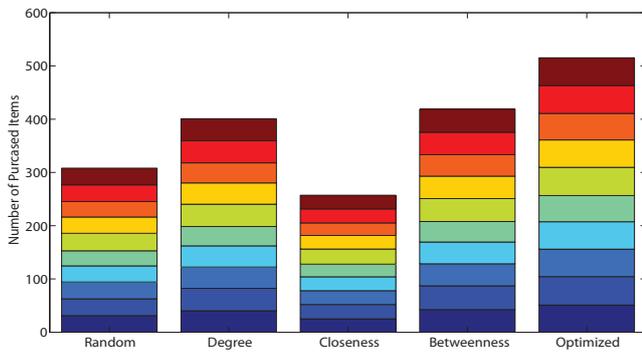


Figure 6: The number of sold items vs. different advertising methods. The assumption is that an agent with expected desire greater than 0.01 will purchase the product. Different colors in each bar indicates the number of sold items of each advertised products. As there exist ten different products, the bar is divided into ten parts.

To estimate the performance of algorithms in selling the products to *Regular* agents, we assumed that agents with expected desire higher than a threshold will purchase the product. Figure 6 shows the average of total purchased items by agents with the purchasing threshold as 0.01. Again, we see that our proposed algorithm is the most successful method in advertising and selling products.

7. RELATED WORK

Social ties between users play an important role in dictating their behavior. One of the ways this can occur is through social influence where a behavior or idea can propagate between friends. In [1], the authors examine the statistical correlation between the actions of friends in a social network by considering factors such as homophily and possible unobserved confounding variables. Hence it follows that it is not only important to advertise to your customer but also to your potential customer’s friends.

One theory about influential nodes is that they can be characterized as a set of initial nodes that trigger behavior cascades. This set of nodes can then be identified either using probabilistic approaches [2, 18] or optimization-based techniques. For example, in [16] the behavior spreads in a cascading fashion according to a probabilistic rule, beginning with a set of initially active nodes. To identify influential agents, they select a set of individuals to target for initial activation, such that the cascade beginning with this active set is as large as possible in expectation. [18] find influential nodes in a complex social network by formulating the likelihood for information diffusion data, the activation time sequence data over all nodes; they propose an iterative method to search for the probabilities that maximize this likelihood.

Apolloni et al. [2] examine the spread of information through personal conversations by proposing a probabilistic model to determine whether two people will converse about a particular topic based on their similarity and familiarity. Similarity is modeled by matching selected demographic characteristics, while familiarity is modeled by the amount of contact required to convey information. On the other hand, [22] propose a learning method for ranking influential nodes and perform behavioral analysis of topic propagation; they compare the results with conventional heuristics that do not consider diffusion phenomena.

In this paper, we present one approach for framing and solving the optimization problem using convex programming. The opti-

mization problem can also be solved using greedy algorithms (e.g., [19, 17]) that find approximate solutions using graph theory. [15] also utilized greedy algorithms to identify the influential nodes. Intelligent heuristics can be used to improve the scalability of influence maximization [8]. [9] made improvements upon existing greedy algorithms to further reduce run-time and proposed new degree discount heuristics that improve influence spread.

The effects of network topology on influence propagation have been studied in several domains, including technology diffusion, strategy adaption in game-theoretic settings, and the admission of new products in the market [15]. It has been demonstrated that the way information spreads is affected by the topology of the interaction network [26] and also that there exists a relationship between a person’s social group and his/her personal behavior [24].

Social network analysis has been used as a tool for implementing effective viral marketing; influential nodes are identified either by following interaction data or probabilistic strategies. For example, Hartline et al. [11] solve a revenue maximization problem to investigate effective marketing strategies. The assumption is that a set of influential nodes can propagate the information about the items to other nodes, and therefore the objective is to find influential buyers in a social network. [27] presented a targeted marketing method based on the interaction of subgroups in social network. Moreover, [4], similar to this work, considered the existing homophily in social networks. But instead of finding influential nodes, they base their advertising strategy on the profile information of users and user-only models.

Our work differs from related work in that, our model not only considers homophily and group membership but also incorporates other important factors such as the positive and negative effect of competing product advertisements and the correlation among demand for different products. Our optimization approach is largely unaffected by the additional complexity since these factors only affect the long-term expected value and not the actual solution method. Outside of social network marketing approaches, there exist many marketing methods based on personalization techniques for delivering advertisements [14] or news [3].

8. CONCLUSION AND FUTURE WORK

Agent-based simulation is an important tool for understanding the behavior of social systems, enabling the analysis of population-level effects over long time horizons that are not easily studied within the confines of the lab. In this paper, we present a model for researching techniques of large-scale persuasion; rather than focusing on the content of the message, we address the problem of identifying agents that have a high probability of indirectly influencing a large number of additional agents. In popular parlance, these influential agents can be considered to be some mixture of Gladwell’s connectors, mavens, and salesman who affect the long-term beliefs of other agents through their actions or communications [10]. Although this problem generalizes to domains such as politics and social media distribution, we demonstrate our method in a market scenario in which product representatives are attempting to maximize their sales through effective advertisement placement. To solve this influence maximization problem, we compute the steady-state expectation of the system and introduce a new optimization approach for selecting influential nodes. Contrary to previous work, our approach is suitable for computing an optimal solution across multiple products and handling the interactions of multiple influencing agents. Our results show that our technique conclusively outperforms a set of centrality-based techniques at selecting influential agents and maximizing total product sales. Our social model accurately reflects many of the complexities of real-

world human systems, such as group membership effects, product preference dependencies, and a network structure driven by multiple conflicting forces. Moreover, in our model, an regular agent need not represent an individual person but can be thought of as an abstraction over communities of people. Then larger systems can simply be solved in a hierarchical, but non-exact, fashion. In our future work, we plan to mathematically analyze the short time behavior of the system. This analysis will allow us to generalized our solution technique to a dynamic network whose structure varies over time [6] which is a characteristic possessed by certain real-world social systems.

9. ACKNOWLEDGMENTS

This research was funded by AFOSR YIP FA9550-09-1-0525.

10. REFERENCES

- [1] A. Anagnostopoulos, R. Kumar, and M. Mahdian. Influence and correlation in social networks. In *Proceeding of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 7–15. ACM, 2008.
- [2] A. Apolloni, K. Channakeshava, L. Durbeck, M. Khan, C. Kuhlman, B. Lewis, and S. Swarup. A study of information diffusion over a realistic social network model. In *2009 International Conference on Computational Science and Engineering*, pages 675–682. IEEE, 2009.
- [3] L. Ardissono, L. Console, and I. Torre. On the application of personalization techniques to news servers on the www. *AI* IA 99: Advances in Artificial Intelligence*, pages 261–272, 2000.
- [4] A. Bagherjeiran and R. Parekh. Combining behavioral and social network data for online advertising. In *IEEE International Conference on Data Mining Workshops (ICDMW'08)*, pages 837–846. IEEE, 2008.
- [5] A. Barabasi and E. Bonabeau. Scale-free networks. *Scientific American*, pages 60–69, May 2003.
- [6] A. Casteigts, P. Flocchini, W. Quattrociocchi, and N. Santoro. Time-varying graphs and dynamic networks. *Arxiv preprint arXiv:1012.0009*, 2010.
- [7] M. Cha, A. Mislove, and K. Gummadi. A measurement-driven analysis of information propagation in the flickr social network. In *Proceedings of the 18th International Conference on World Wide Web*, pages 721–730. ACM, 2009.
- [8] W. Chen, C. Wang, and Y. Wang. Scalable influence maximization for prevalent viral marketing in large-scale social networks. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1029–1038. ACM, 2010.
- [9] W. Chen, Y. Wang, and S. Yang. Efficient influence maximization in social networks. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 199–208. ACM, 2009.
- [10] M. Gladwell. *The tipping point: How little things can make a big difference*. Little, Brown and Company, 2000.
- [11] J. Hartline, V. Mirrokni, and M. Sundararajan. Optimal marketing strategies over social networks. In *Proceeding of the 17th International Conference on World Wide Web*, pages 189–198. ACM, 2008.
- [12] B. Hung. Optimization-based selection of influential agents in a rural afghan social network. *Master's Thesis, Massachusetts Institute of Technology*, 2010.
- [13] B. Hung, S. Kolitz, and A. Ozdaglar. Optimization-based influencing of village social networks in a counterinsurgency. In *Proceedings of the 4th International Conference on Social Computing, Behavioral-cultural Modeling and Prediction*, pages 10–17. Springer-Verlag, 2011.
- [14] P. Kazienko and M. Adamski. Adrosa—adaptive personalization of web advertising. *Information Sciences*, 177(11):2269–2295, 2007.
- [15] D. Kempe, J. Kleinberg, and É. Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 137–146. ACM, 2003.
- [16] D. Kempe, J. Kleinberg, and É. Tardos. Influential nodes in a diffusion model for social networks. *Automata, Languages and Programming*, pages 1127–1138, 2005.
- [17] M. Kimura and K. S. nd R. Nakano. Extracting influential nodes for information diffusion on a social network. In *Proceedings of the National Conference on Artificial Intelligence*, 2007.
- [18] M. Kimura, K. Saito, R. Nakano, and H. Motoda. Finding influential nodes in a social network from information diffusion data. *Social Computing and Behavioral Modeling*, pages 1–8, 2009.
- [19] M. Kimura, K. Saito, R. Nakano, and H. Motoda. Extracting influential nodes on a social network for information diffusion. *Data Mining and Knowledge Discovery*, 20(1):70–97, 2010.
- [20] J. Leskovec. Amazon product co-purchasing network. In <http://snap.stanford.edu/data/amazon0302.html>, 2003.
- [21] J. R. Magnus and H. Neudecker. *Matrix Differential Calculus with Applications in Statistics and Econometrics, 2nd Ed.* Wiley, 1999.
- [22] K. Saito, M. Kimura, K. Ohara, and H. Motoda. Behavioral analyses of information diffusion models by observed data of social network. *Advances in Social Computing*, pages 149–158, 2010.
- [23] P. Sen and L. Getoor. *Link-based classification*. Technical Report, CS-TR-4858, University of Maryland, Reading, Massachusetts, 2007.
- [24] P. Singla and M. Richardson. Yes, there is a correlation:-from social networks to personal behavior on the web. In *Proceeding of the 17th International Conference on World Wide Web*, pages 655–664. ACM, 2008.
- [25] X. Wang, M. Maghami, and G. Sukthankar. Leveraging network properties for trust evaluation in multi-agent systems. In *Proceeding of the 2011 IEEE/WIC/ACM International Conference on Intelligent Agent Technology, Lyon, France*, 2011.
- [26] F. Wu, B. Huberman, L. Adamic, and J. Tyler. Information flow in social groups. *Physica A: Statistical and Theoretical Physics*, 337(1-2):327–335, 2004.
- [27] W. Yang, J. Dia, H. Cheng, and H. Lin. Mining social networks for targeted advertising. In *Proceedings of the 39th Annual Hawaii International Conference on System Sciences-Volume 06*, pages 137–1. IEEE Computer Society, 2006.