Hierarchical Influence Maximization for Advertising in Multi-agent Markets

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

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

Abstract—Maximizing product adoption within a customer social network under a constrained advertising budget is an important special case of the general influence maximization problem. Specialized optimization techniques that account for product correlations and community effects can outperform network-based techniques that do not model interactions that arise from marketing multiple products to the same consumer base. However, it can be infeasible to use exact optimization methods that utilize expensive matrix operations on larger networks without parallel computation techniques. In this paper, we present a hierarchical influence maximization approach for product marketing that constructs an abstraction hierarchy for scaling optimization techniques to larger networks. An exact solution is computed on smaller partitions of the network, and a candidate set of influential nodes is propagated upward to an abstract representation of the original network that maintains distance information. This process of abstraction, solution, and propagation is repeated until the resulting abstract network is small enough to be solved exactly. Our proposed method scales to much larger networks and outperforms other influence maximization techniques on marketing products.

I. INTRODUCTION

Advertising in today's market is no longer viewed as a matter of simply convincing a potential customer to buy the product but of convincing their social network to adopt a lifestyle choice. It is well known that 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. By considering factors such as homophily and possible unobserved confounding variables, it is possible to examine these behavior correlations in a social network statistically [1]. The aim of viral marketing strategies is to leverage these behavior correlations to create information cascades in which a large number of customers imitate a much smaller set of informed people, who are initially convinced by targeting marketing schemes.

Marketing with a limited budget can be modeled as a specialized version of the influence maximization problem in which the aim is to advertise to the optimal set of seed nodes to modify opinion in the network, based on a known influence propagation model. Commonly used propagation models such as LTM (Linear Threshold Model) and ICM (Independent Cascade Model) assume that a node's adoption probability is conditioned on the opinions of the local network

neighborhood [2]. Much of the previous influence maximization work [3], [4], [5] uses these two interaction models. Since the original LT model and IC model, other generalized models have been proposed for different domains and specialized applications. For instance, the decreasing cascade model simulates processes used in the sociology and economics communities where a behavior spreads in a cascading function according to a probabilistic rule, beginning with a set of nodes that adopt the behavior [2]. In contrast with the original IC model, in the decreasing cascade model the probability of influence propagation from an active node is not constant. Similarly, generalized versions of the linear threshold model have been introduced (e.g., [6], [7]). The simplicity of these propagation models facilitates theoretical analysis but does not realistically model specific marketing considerations such as the interactions between advertisements of multiple products and the effects of community membership on product adoption.

To address these problems, in previous work [8], we developed a model of product adoption in social networks that accounts for these factors, along with a convex optimization formulation for calculating the best marketing strategy assuming a limited budget. These social factors can emerge from different independent variables such as ties between friends and neighbors, social status, and the economic circumstance of the agents. We believe that in marketing, all these factors affect the customers' susceptibility to influence and their ability to influence others. As an example, [9] analyzes the effect of social status on the influence factor of people on Facebook. Having a more realistic model is particularly useful for overcoming negative advertisement effects in which the customers refrain from purchasing any products after being bombarded with mildly derogatory advertisement from multiple advertisers trying to push their own products. It is critical to model the propagation of negative influence as well since it propagates and can be stronger and more contagious than positive influence in affecting people's decisions [10].

The main limitation of this and similar types of optimization approaches is that they involve computationally expensive matrix operations which prevent these algorithms from scaling to larger networks. In this paper, we propose a hierarchical influence maximization approach that advocates "divide and conquer"—the network is partitioned into multiple smaller networks that can be solved exactly with optimization techniques, assuming a generalized IC model, to identify a candidate set of seed nodes. The candidate nodes are used to create

a distance-preserving abstract version of the network that maintains an aggregate influence model between partitions. Here we demonstrate how this abstraction technique can be used to create a scalable algorithm Hierarchical Influence Maximization (HIM) for maximizing steady-state product adoption by customers connected by a social network. Moreover, we present a theorem which shows that the realistic social system model has a fixed-point, validating the strategy of optimizing product adoption at the steady state.

The paper is organized as follows. Section II provides an overview of the related work in influence maximization. Section III introduces our proposed method, Hierarchical Influence Maximization (HIM), as well as summarizing the operation of the realistic product adoption model introduced by [8]. We evaluate our method vs. another optimization technique and a set of centrality based network analysis techniques in Section IV. We end the paper with a discussion of future work.

II. RELATED WORK

Influence maximization can be described as the problem of identifying a small set of nodes capable of triggering large behavior cascades that spread through the network. This set of nodes can be discovered using probabilistic approaches (e.g., [11], [12]) or optimization-based techniques. [13], [8] treat influence maximization as a convex optimization problem; this is feasible for influencing small communities but does not scale to larger scale problems. Due to the matrix computation requirements, these approaches fail when the number of agents in the system increases. Our HIM algorithm overcomes this deficiency by using a hierarchical approach to factor the system into smaller matrices.

The HIM model is designed to work on a complex social system where multiple factors affect the propagation of influence. The simpler case, where the network topology alone dictates activation spread, has been examined by multiple research groups, seeking to improve on Kempe's early work on greedy approaches for influence maximization [14]. Examples of possible speedups include innovations such as the use of a shortest-path based influence cascade model [15] or a lazy-forward optimization algorithm [16] to reduce the number of evaluations on the influence spread of nodes. Clever heuristics have been used very successfully to speed computation in both the LT model (e.g., the PMIA algorithm [4]) and also the IC model [5]. In this paper, instead of using the original cascade models by Kempe et al. we introduce a cascade model that accounts for product interactions and community differences in influence propagation.

Proposed models for investigating how ideas and influence propagate through the network have been applied to many domains, including technology diffusion, strategy adoption in game-theoretic settings, and the admission of new products in the market [14]. For viral marketing, influential nodes can be identified either by following interaction data or probabilistic strategies. For example, Hartline et al. [17] solve a revenue maximization problem to investigate effective marketing strategies. [18] presented a targeted marketing method based on the interaction of subgroups in social network. Similar to this work, Bagherjeiran and Parekh leverage purchasing homophily in social networks [19]. But instead of finding influential

nodes, they base their advertising strategy on the profile information of users. Achieving deep market penetration can be an important aspect of marketing; Shakarian and Damon present a viral marketing strategy for selecting the seed nodes that guarantees the spread of the word to the entire network [20]. Our work differs from related work in that our model not only considers social factors but also incorporates the negative effect of competing product advertisements and the correlation between demand for different products. Our optimization approach is largely unaffected by the additional complexity since these factors only impact the long-term expected value and not the actual solution method.

III. METHOD

Our proposed hierarchical approach operates as follows:

- 1) Create a local network for each node consisting of its neighbors and neighbors of neighbors;
- 2) Model the effect of the outside network by assigning a virtual node for each boundary node to abstract activity outside the local partition;
- 3) Update the interaction parameters to the virtual node based on the model and the network connections;
- 4) Create a candidate set of influential nodes for each local network using convex optimization to maximize steady state product adoption;
- 5) Propagate the candidate set upward to a higher-level of abstraction and link the abstract nodes based on their shortest paths in the previous network;
- 6) Repeat the abstraction process until the resulting network is small enough to be optimized as a single partition; the resulting set of candidate nodes is then targeted for advertisement.

Figure 1 demonstrates the process of the algorithm with three hierarchies. The selected nodes at each local neighborhood, colored in red, are moved to the upper hierarchy and re-connected based on shortest path distances from the lower-level. The same process is repeated at the next hierarchy to select more influential nodes. The procedure terminates at the last hierarchy when the number of influential nodes finally is smaller than the advertising budget.

A. Market Model

To explore the efficiency of the proposed hierarchical influence maximization (HIM) method in business marketing, we have used the multi-agent system model, presented by [8], to simulate a social system of potential customers. We have slightly changed the definition of some parameters in this model to make a more sensible model with generalized capabilities. To avoid any confusion, we retained the same notation from previous work.

In this model, the population of N agents, represented by the set $A = \{a_1, \dots, a_N\}$, consists of two types of agents ($A = A_R \cup A_P$), named *Regular* and *Product* agents respectively. The *Regular* agents are the potential customers in the market who will occasionally change their attitudes on purchasing products based on the influence they receive either from other neighbors or from the *Product* agents who represent salespeople offering one specific product.

Fig. 1. At each hierarchical level (H_i) local neighborhoods are created and influential nodes (red) are selected using an optimization technique. Nodes that have been selected at least once as an influential node are transferred to the next level of the hierarchy. At the higher levels, the connection between selected nodes is defined using the shortest path distance in the original network. The process is repeated until the final set of influential nodes is smaller than the total advertising budget.

Regular agents belong to a connected social network where the directed weighted links in this network possess a history of past interactions among the agents. This social network 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 and the in-node and out-node degree of agent a_i is the sum of all in-node and out-node weights, respectively.

In this model a vector of $\vec{X}_i$ is assigned to each agent, both *Regular* and *Product* agents, representing the attitude or desire of the agent toward all of the products in the market. Each element of this vector, x_{ip} , is a random variable in the $[-1, 1]$ interval that indicates the desire of agent a_i to buy an item or consume a specific product, p .

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,

$$p_{ij} = \begin{cases} \frac{e_{ij}}{d_{out}^i} & i, j \in A_R \\ \frac{u_{ji}}{Threshold} & i \in A_R, j \in A_P \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where 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. The u_{ji} parameter is an indicator marking whether the *Product* agent is connected to the *Regular* agent.

At each interaction there is a chance for agents to influence each other and change their desire vector for purchasing or

consuming a product. During these interactions the *Product* agents never change their attitude and maintain a fixed desire vector of 1 toward themselves and -1 toward the other advertising companies. The probability that agent i is susceptible to agent j is denoted as α_{ij} and calculated as:

$$\alpha_{ij} = \begin{cases} \frac{e_{ji}}{d_{in}^i} & i, j \in A_R \\ cte & i \in A_R, j \in A_P \end{cases} \quad (2)$$

The other important parameter in the agent influence process is ε_{ij} , which determines how much agent j will influence agent i . This parameter indicates the role of social factors in decision making of agents. In contrast to previous work, we did not restrict this parameter to a specific distribution to provide more flexibility to the model.

Moreover, in real life there is a correlation between the user demand for different products in the market. The desire of customers for a specific product is related to his/her desire toward other similar products. Matrix $\mathbf{M}$ models this correlation, and we consider its effect in our formulation.

The ultimate goal of our marketing problem is to recognize the influential agents in the graph and define 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.

Note that the links between *Product* agents and *Regular* agents are directed links from products to agents and not in the opposite direction.

B. Generalized ICM

We use a generalized version of ICM similar to [8], [21]. The dynamics of the model at each iteration k proceed as follows:

- 1) Agent i initiates the interaction according to 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:

$$\begin{cases} \vec{X}_i(k+1) = \varepsilon_{ij} \mathbf{M} \vec{X}_i(k) + (1 - \varepsilon_{ij}) \mathbf{M} \vec{X}_j(k) \\ \vec{X}_j(k+1) = \vec{X}_j(k) \end{cases} \quad (3)$$

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:

$$\begin{cases} \vec{X}_i(k+1) = \vec{X}_i(k) \\ \vec{X}_j(k+1) = \vec{X}_j(k) \end{cases} \quad (4)$$

It is worthwhile to note that the above interaction model can be degraded to the IC model, if we set $\varepsilon_{ij} = 0$, $\mathbf{M} = \mathbf{I}$, and restrict p_{ij} s to be equal to 1 right after activation of any node and equal to 0 the rest of the time. Also since the values of the desire vector range from $[-1, 1]$, the x_{ip} s $\in [0, 1]$ and x_{ip} s $\in [-1, 0]$ can be quantized to 1 and 0 respectively to match the IC model representation of activation and deactivation.

TABLE I. HIM ALGORITHM

```

HIM (Agent, E, P, A,  $A_R$ ,  $H_{max}$ ,  $r$ )
 $H = 0$ 
 $\mathbf{E}^H = \mathbf{E}$ 
 $N^H = |A_R|$ 
While stopCriteria do
   $H = H + 1$ 
   $\text{inflList} = \text{NULL}$ 
  for  $i = 1$  to  $N^H$  do
     $\text{neighborList} = \text{FindNeighborList}(i, r, \mathbf{E}^H)$ 
     $\mathbf{E}_i^H = \text{Subgraph}(\text{neighborList}, \mathbf{E}^H)$ 
     $\mathbf{E}_i^H = \text{AddOutsideWorld}(\mathbf{E}_i^H, \mathbf{E}_i^H)$ 
     $(\mathbf{P}_i, \mathbf{A}_i) = \text{UpdateMat}(\mathbf{E}_i^H, \mathbf{P}, \mathbf{A}, \text{neighborList})$ 
     $L = \text{Optimize}(\text{Agent}, \mathbf{E}_i^H, \mathbf{P}_i, \mathbf{A}_i)$ 
     $\text{inflList} = \text{inflList} \cup L$ 
     $\text{Agent} = \text{UpdateAgent}(\text{inflList})$ 
  end for
   $N^H = |\text{inflList}|$ 
   $\mathbf{U} = \text{MakeU}(\text{Agent})$ 
  stopCriteria =  $\text{UpdateCriteria}(\text{inflList}, H)$ 
   $\mathbf{E}^H = \text{UpdateHierarchy}(\text{inflList})$ 
end while
return  $\mathbf{U}$ 

```

C. HIM Algorithm

Using these assumptions about customer product adoption dynamics, we devised a new scalable optimization technique, Hierarchical Influence Maximization (HIM). The pseudocode of our proposed HIM algorithm is presented in Table I. Here, matrix $\mathbf{E}$ represents the connection matrix among *Regular* agents, and matrices $\mathbf{P}$ and $\mathbf{A}$ contain all the p_{ij} 's and α_{ij} 's of the market model, respectively. In other words, all the interactions and influence probabilities between two pairs of *Regular* agents, (A_R), are embedded in the elements of these matrices. *Agent* contains all the information about *Regular* and *Product* agent characteristics including desire vectors, ($\vec{X}_i$'s), and influence tag vectors, $\vec{I}_i$'s with size P , where I_{ip} indicates the number of times that agent i has been selected as an influential node for product p . The algorithm receives as input all the available data on the agents and the model, and the output of the algorithm is the $\mathbf{U}$ matrix that contains the assignments of u_{ji} 's and shows the final connection matrix between all the products and influential seed nodes.

The level of the hierarchy is indicated by parameter H which increments until the stopping criteria are satisfied. At each hierarchy (H), we iterate over all the nodes (i s) in the network of that hierarchy, (E^H), and list the neighboring agents around each node. The radius of the neighborhood, denoted with parameter r , indicates the granularity of analysis. Based on radius r , we partition the network into subsections, (E_i^H), and update the probability matrices, $\mathbf{P}_i$ and $\mathbf{A}_i$ for that subsection. HIM selects the influential agents in that local network, E_i^H , using an optimization technique and tags them for future use. The process of node selection is described in detail in III-C2. Then we add these influential nodes to the set of influential nodes that have been identified in other neighborhoods in the same hierarchy.

1) *Outside World Effect*: When a local neighborhood is detached from the complete network, there exist boundary nodes that are connected to nodes outside the neighborhood. These connections that fall outside of the neighborhood can potentially affect the desire vector of agents within the neighborhood. One possible approach is to ignore these effects and only consider the nodes inside the partition. In this paper

we account for these effects by allocating a virtual node to each boundary node. This virtual node is the representative of all nodes outside the neighborhood that are connected to the boundary node. Figure 2 illustrates the abstraction of outside world effect and shows how the model's parameters are calculated between each boundary and virtual node.

2) *Node Selection*: The process of selecting influential nodes is repeated at each hierarchy and at each local neighborhood surrounding node i . Following previous works [13], [21], [8], we model the desire dynamic of all agents as a Markov chain where the state of the local neighborhood is a matrix of all existing 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 local network around agent i at the k^{th} iteration is a vector of random variables, denoted as $\mathbf{X}_i(k) \in \mathbb{R}^{N_{H_i} \times P \times 1}$ (created through a concatenation of N_i^H vectors of size P) and expressed as:

$$\mathbf{X}_i(k) = \begin{pmatrix} [\vec{X}_1(k)] \\ \vdots \\ [\vec{X}_{N_i^H}(k)] \end{pmatrix}$$

We calculate the expected long-term desire of the agents in each local network around agent i and this calculation results in the following formulation:

$$E[\mathbf{X}_i(k+1)] = E[\mathbf{X}_i(k)] + \mathbf{Q}_i E[\mathbf{X}_i(k)] \quad (5)$$

where $\mathbf{Q}_i$ is a block matrix representing the interactions among *Regular* agents in the neighborhood and interactions between the *Regular* agents and all the *Products*.

3) *Convergence*: In previous work [8], we solved Equation 5 at the steady state and in a global fashion, without giving any guarantee that the state of the system actually reaches the steady state. Here, by using Brouwer fixed-point theorem [22], we prove that each local neighborhood has a fixed-point, hence solving Equation 5 at steady state is a valid choice.

The Brouwer fixed-point theorem states that:

Theorem 1. *Every continuous function from a closed ball of a Euclidean space to itself has a fixed point.*

According to the calculation of Equation 5, $E[\mathbf{X}_i(k+1)]$ is a continuous function as it is the sum of two continuous ones. Also since $\vec{X}_i(k+1)$ in Equation 3 is a bounded function in $[-1, 1]$, its expectation ($E[\mathbf{X}_i(k+1)]$) will be bounded as well. As a result we have a bounded, continuous function which is guaranteed a fixed point by the Brouwer fixed-point theorem. Consequently, we can follow all the calculations of [8] and solve our problem with the proposed optimization algorithm to find the assignment of u_{ji} s in a way to maximize the long-term expected desire vector of agents toward all the products in the market.

4) *Update Hierarchy*: When we proceed from one hierarchy to the next one, the selected nodes which are propagated to the upper hierarchy are not necessarily adjacent. Therefore, we need to define the interaction model between them based on their position in the real network. The *UpdateHierarchy*

Fig. 2. The network on the left is an example of a neighborhood around node e ; the network on the right is the equivalent network with virtual nodes representing the outside world effect. Here w can be any interaction parameter such as link's weight, α , or ϵ . The direction of the interaction with the virtual node is based on the type of links the boundary node has with the nodes outside the neighborhood. The value of the parameter is the average over all similar types of interactions with outside world.

function is responsible for building the proper network connection and interaction model for the next hierarchy based on the selected influential nodes in current hierarchy. These nodes were propagated to the higher hierarchy by being selected as influential nodes in at least one local neighborhood. It is possible for a node to be present in multiple partitions and be selected more than once.

Note that the selected nodes are unlikely to be adjacent nodes in the actual network E . Therefore we need to find a way to form their connections to construct E_H . To do so, we look at the shortest path between these nodes in network E and use that to calculate the weight of the edges in E^H . In the E^H network the weight of the link between two selected nodes is the product of the weights of the shortest path between these two nodes in the previous hierarchy. Also the probabilities of interaction and influence between two influential nodes is set to be the product of the probabilities along the shortest path between them.

5) *Termination Criteria*: To terminate the loop, we establish two different criteria in the *UpdateCriteria* function. This function checks the stopping criteria based on the level of the hierarchy and the list of influential nodes. One criterion is based on the maximum number of levels in the hierarchy and the other is based on the ratio of the selected influential nodes and the advertising budget. According to the *stopCriteria* output, the algorithm decides whether to proceed to a higher hierarchy or to stop the search, returning the current U matrix to be used as the advertising assignment.

IV. EVALUATION

A. Experimental Setup

We conducted a set of simulation experiments to evaluate the effectiveness of our proposed node selection method on marketing items in a simulated social system with a static network. The parameters of the interaction model for all the runs are summarized in Table II(a). All the results are computed over an average of 100 runs which represent ten different simulations on each of ten network structures.

In the *Regular* and *Product* agent interactions, parameters α and ϵ are fixed for a given interaction and are presented in

TABLE II. PARAMETER SETTINGS

(a) Market Model Parameters		
Parameter	Value	Descriptions
<i>Threshold</i>	2	Number of links between P and R agents
ϵ	0.4	Influence factor between P and R agents
α	0.8	Probability of influence between P and R agents
<i>R</i>	Variable	Number of <i>Regular</i> agents
<i>P</i>	10	Number of <i>Product</i> agents
<i>NIterations</i>	60,000	Number of iterations
<i>NRun</i>	10	Number of runs
<i>NNet</i>	10	Number of different networks

(b) HIM Parameters		
Parameter	Value	Description
r	3	Neighborhood radius
H_{max}	5	Max level of hierarchy

Table II(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 1 and are parametric. Table II(b) provides the parameters for our HIM algorithm (neighborhood radius and the maximum hierarchy level). The remaining part of the social system setup is given by 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 an item, is equal to one.

B. Results

We compare our hierarchical algorithm with the original optimization method (OIM) described in [8] and a set of centrality-based measures commonly used in social network analysis for identifying influential nodes based on network structure [14]. The comparison methods are:

- **OIM**: The Optimized Influence Maximization method finds the influential nodes globally by using a convex optimization method over the entire network.
- **Degree**: Assuming that high-degree nodes are influential nodes in the network, we calculated the probability of

advertising to a *Regular* agent based on the out-degree of the agents and linked the *Product* agents according to a preferential attachment model. Therefore, nodes with higher degree had an increased chance of being selected as an advertising 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 chance of being selected as an influential node.
- **PageRank:** On the assumption that the nodes with the greatest PageRank score have a higher chance of influencing the other nodes, we based the probability of node selection on its PageRank value.
- **Random:** In this baseline, we simply select the nodes uniformly at random.

To evaluate these methods, we started the simulation with an initial desire vector set to 0 for all agents, and simulated 60000 iterations of agent interactions. The entire process of interaction and influence is governed by Equations 3 and 4 (Section III-B). At each iteration, we calculated the average of the expected desire value of the agents toward all products. This average is calculated over 100 runs (10 simulations on 10 different network structures). Note that the desire vector of *Product* agents remain fixed for all products; in our simulation it was set to 1 for the product itself and -0.1 for all other products (e.g., $\mu_1 = [1 \ -0.1 \ -0.1 \dots \ -0.1]$). We used the same network generation technique described in [8] for generating customer networks.

1) *Performance:* To compare the performance of these methods, the average expected desire value of the agents in a network with 150 agents has been shown over time in Figure 3. Here we selected 150 agents as an optimal number of agents to compare all the algorithms together. With fewer agents, having ten simultaneously marketed products saturates the network while with a larger number of agents OIM suffers from scalability issues and the convex optimization method was not feasible due to the near singular interaction matrix. In Figure 3, by using the marketing-specific optimization methods for allocating the advertising budget, the desire value of the agents toward all products increases the most, resulting in the largest number of sales. Although HIM sacrificed some performance in favor of scalability, it clearly outperforms the centrality measurement methods. The locally-optimal selection approach of HIM results in a slightly lower performance compared to globally optimal OIM.

Figure 4 shows the final average value of the expected desire of agents in the last iteration for different number of *Regular* agents. Although OIM with global optimization method outperforms HIM and other centrality measurement methods, it is incapable of scaling up to 300 and more agents in the network due to near singular interaction matrix. HIM with the ability to scale up linearly to number of nodes provides a sub-optimal and yet practical solution in selecting the influential nodes in large networks.

2) *Run-time:* Table III shows a runtime comparison between the two optimization methods, HIM (proposed) and OIM (original). In small networks the runtime of the global optimization method is less than the hierarchical but as the size of network grows, its run time increases exponentially while

Fig. 3. The average of agents' expected desire vs. number of iterations, calculated across all products and over 100 runs (10 different runs on 10 different networks). The optimization methods have the highest average in comparison to the centrality measurement heuristics. As HIM is a sub-optimal method, it is unsurprising that its performance is worse than the global optimization method, OIM.

Fig. 4. The average of the final expected desire vectors for different numbers of *Regular* agents and 10 *Product* agents. The optimization based methods (OIM and HIM) outperforms the other methods in selecting the seed nodes. While OIM is more successful than HIM in selecting the influential nodes, it is unable to scale-up to networks with 300 agents and higher.

the run time of the HIM increases at a slower rate. The long runtime of OIM for the networks larger than 200 nodes, makes the algorithm impractical for finding influential nodes in very large networks.

TABLE III. RUNTIME COMPARISON BETWEEN OIM AND HIM

Number of agents	OIM	HIM
50	10.67s	74.09s
100	94.76s	160.80s
150	290.67s	208.97s
200	897.51s	354.35s

3) *Jaccard Similarity:* To analyze the differences between the algorithms' selection of influential nodes, we use the Jaccard similarity measurement. This measurement is calculated by dividing the intersection of two selected sets by the union of these sets. Figure 5 shows this measurement for all pairs of algorithms. The OIM and HIM algorithms have the highest similarity compared to the other methods with a similarity value of 0.47. The other pairs of methods have very low similarities, resulting in dark squares in the figure. Not surprisingly, Random has the least similar node selection to other methods. This shows that HIM finds many of the

same nodes as the original OIM algorithm, with a much lower runtime cost.

Fig. 5. The average Jaccard similarity measurements between different methods, calculated over 100 runs (10 runs on 10 different networks). Lighter squares denote greater similarity between a pair of algorithms. Note that HIM's selection of nodes is fairly close to OIM's optimal selection.

V. CONCLUSION AND FUTURE WORK

In this paper, we present a general hierarchical approach for applying optimization techniques to influence maximization and demonstrate its use for product marketing. The advantage our method has over network-only seed selection techniques is that it can account for item correlations and community effects on the product adoption rate. Our method comes close to the optimal node selection, at substantially lower runtime costs. One possible extension of this work is to generalize the market simulation to explicitly model the adversarial effects between competing advertisers as a Stackelberg competition. Also in this paper we assumed that the probability of interaction and influence between two agents is small, compared to the size of the network, which results in the agents sticking to a decision for a reasonable period of time. However if the network is smaller or the probability of interaction increases, there can be large fluctuations in the agents' desire vector. Applying a parameter to the model which forces the agents to retain their decisions for a minimum period, regardless of external interactions, would ameliorate this issue. [23]. Furthermore, working with dynamic networks where the agents can enter and leave the network would be another practical extension of this work.

VI. ACKNOWLEDGMENTS

This research was supported in part by DARPA award D13AP00002 and NSF IIS-08451.

REFERENCES

- [1] A. Anagnostopoulos, R. Kumar, and M. Mahdian, "Influence and correlation in social networks," in *Proceeding of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2008, pp. 7–15.
- [2] D. Kempe, J. Kleinberg, and É. Tardos, "Influential nodes in a diffusion model for social networks," *Automata, Languages and Programming*, pp. 1127–1138, 2005.
- [3] W. Chen, Y. Yuan, and L. Zhang, "Scalable influence maximization in social networks under the linear threshold model," in *Proceedings of the IEEE International Conference on Data Mining (ICDM)*, 2010, pp. 88–97.
- [4] W. Chen, C. Wang, and Y. Wang, "Scalable influence maximization for prevalent viral marketing in large-scale social networks," in *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2010, pp. 1029–1038.
- [5] C. Wang, W. Chen, and Y. Wang, "Scalable influence maximization for independent cascade model in large-scale social networks," *Data Mining and Knowledge Discovery*, pp. 1–32, 2012.
- [6] N. Pathak, A. Banerjee, and J. Srivastava, "A generalized linear threshold model for multiple cascades," in *International Conference on Data Mining (ICDM)*, 2010, pp. 965–970.
- [7] S. Bharathi, D. Kempe, and M. Salek, "Competitive influence maximization in social networks," *Internet and Network Economics*, pp. 306–311, 2007.
- [8] M. Maghami and G. Sukthankar, "Identifying influential agents for advertising in multi-agent markets," in *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*, 2012, pp. 687–694.
- [9] S. Aral and D. Walker, "Identifying influential and susceptible members of social networks," *Science*, vol. 337, no. 6092, pp. 337–341, 2012.
- [10] W. Chen, A. Collins, R. Cummings, T. Ke, and et. al, "Influence maximization in social networks when negative opinions may emerge and propagate," in *Proceedings of the SIAM International Conference on Data Mining*, 2011.
- [11] 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 *Proceedings of the International Conference on Computational Science and Engineering*, 2009, pp. 675–682.
- [12] 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*, pp. 1–8, 2009.
- [13] B. Hung, "Optimization-based selection of influential agents in a rural Afghan social network," *Master's Thesis, Massachusetts Institute of Technology*, 2010.
- [14] D. Kempe, J. Kleinberg, and É. Tardos, "Maximizing the spread of influence through a social network," in *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ACM, 2003, pp. 137–146.
- [15] M. Kimura and K. Saito, "Tractable models for information diffusion in social networks," *Knowledge Discovery in Databases (PKDD)*, pp. 259–271, 2006.
- [16] J. Leskovec, A. Krause, C. Guestrin, C. Faloutsos, J. VanBriesen, and N. Glance, "Cost-effective outbreak detection in networks," in *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2007, pp. 420–429.
- [17] J. Hartline, V. Mirrokni, and M. Sundararajan, "Optimal marketing strategies over social networks," in *Proceeding of the International Conference on World Wide Web*, ACM, 2008, pp. 189–198.
- [18] W. Yang, J. Dia, H. Cheng, and H. Lin, "Mining social networks for targeted advertising," in *Proceedings of the Annual Hawaii International Conference on System Sciences*, IEEE Computer Society, 2006.
- [19] A. Bagherjeiran and R. Parekh, "Combining behavioral and social network data for online advertising," in *IEEE International Conference on Data Mining Workshops (ICDMW)*, 2008, pp. 837–846.
- [20] P. Shakarian and D. Paulo, "Large social networks can be targeted for viral marketing with small seed sets," in *Proceedings of the IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 2012, pp. 1–8.
- [21] B. Hung, S. Kolitz, and A. Ozdaglar, "Optimization-based influencing of village social networks in a counterinsurgency," in *Proceedings of the International Conference on Social Computing, Behavioral-cultural Modeling and Prediction*, 2011, pp. 10–17.
- [22] D. Leborgne, *Calcul différentiel et géométrie*. Presses universitaires de France, 1982.
- [23] L. Liow, S. Cheng, and H. Lau, "Niche-seeking in influence maximization with adversary," in *Proceedings of the Annual International Conference on Electronic Commerce*, ACM, 2012, pp. 107–112.