Limiting the Spread of Misinformation in Social Networks

Ceren Budak
Department of Computer Science, UCSB
Santa Barbara, USA
cbudak@cs.ucsb.edu

Divyakant Agrawal
Department of Computer Science, UCSB
Santa Barbara, USA
agrawal@cs.ucsb.edu

Amr El Abbadi
Department of Computer Science, UCSB
Santa Barbara, USA
amr@cs.ucsb.edu

ABSTRACT

In this work, we study the notion of competing campaigns in a social network. By modeling the spread of influence in the presence of competing campaigns, we provide necessary tools for applications such as emergency response where the goal is to limit the spread of misinformation. We study the problem of *influence limitation* where a "bad" campaign starts propagating from a certain node in the network and use the notion of limiting campaigns to counteract the effect of misinformation. The problem can be summarized as identifying a subset of individuals that need to be convinced to adopt the competing (or "good") campaign so as to minimize the number of people that adopt the "bad" campaign at the end of both propagation processes. We show that this optimization problem is NP-hard and provide approximation guarantees for a greedy solution for various definitions of this problem by proving that they are submodular. Although the greedy algorithm is a polynomial time algorithm, for today's large scale social networks even this solution is computationally very expensive. Therefore, we study the performance of the degree centrality heuristic as well as other heuristics that have implications on our specific problem. The experiments on a number of close-knit regional networks obtained from the Facebook social network show that in most cases inexpensive heuristics do in fact compare well with the greedy approach.

Categories and Subject Descriptors

F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Non-numerical Algorithms and Problems

Keywords

social networks, information cascades, misinformation propagation, competing campaigns, submodular functions

1. INTRODUCTION

Until very recently, knowledge about the vast majority of public events has been provided by, or filtered through, the mass media which had almost complete autonomy over the decisions as to which piece of information is "newsworthy". This few-to-many information model has been shattered by advances in technology during the last decade, especially with the adoption of the online social networks [11, 28, 30, 33]. Social networks have been shown to have benefits as a medium for fast, widespread information dissemination. They provide fast access to large scale news data, sometimes even before the mass media as in the case of announcement of death of Michael Jackson [26]. They also serve as a medium to collectively achieve a social goal. For instance with the use of group and event pages in Facebook, events such as "Day of Action" protests reached thousands of protestors [12].

While the ease of information propagation in social networks can be very beneficial, it can also have disruptive effects. One such example was observed during the recent shootings at Fort Hood, Texas, when a soldier inside the base sent out messages via Twitter as the event unfolded. Her incorrect reports of multiple shooters and shooting locations quickly spread through the social network and even to the mass media where it was reported on television broadcasts [16]. Another example is the spread of misinformation on swine flu in Twitter [27]. The spread of misinformation in this case reached a very large scale causing panic in the population. As people were being misinformed on the subject, they also contributed in this misinformation trend by repeating it and therefore disseminating it even further. Although social networks like Twitter are the main source of news for many people today [25], they are still not considered reliable due to such problems.

Clearly, in order for social networks to serve as a reliable platform for disseminating critical information, it is necessary to have tools to limit the effect of misinformation. In this study one of our main goals is to address this specific problem. In the presence of a misinformation cascade, we aim to find the most optimal way of disseminating "good information" that will minimize the devastating effects of the misinformation campaign. For instance in the case of [27, 16], we seek ways of making sure that most of the users of the social network hear about the correct information before the bad one. In this way, we can make social networks a more "trustworthy" or "reliable" source of information. In addition to the implication our work has in limiting the effect of misinformation, the methods we introduce can also be applied to any two competing campaigns that are simultaneously spreading throughout the network. Since in a real social network, there are usually two or more correlated information cascades happening simultaneously, we believe capturing this characteristic is crucial to getting a more realistic model of real social networks.

There are several objective functions one might want to optimize in the presence of competing campaigns. One objective might be to try to minimize the population affected by the other campaign referred to as the problem of *eventual influence limitation*. In the case of [27], this function is equivalent to minimizing the number of people that hear about and believe the misinformation about swine flu. Alternatively, an objective could be to make sure that misinformation dissemination on swine flu ends no later then some time period. We refer to this objective function as *time sensitive influence limitation*. In this work we focus specifically on *eventual influence limitation* since it serves more to limiting the impact of misinformation which is our main motivation.

In this work, we study the problem of minimizing number of people that adopt the misinformation and prove that even though the general problem does not exhibit the submodular property, certain restricted versions of it are in fact submodular. We exploit this property to provide efficient solutions with approximation bounds. We also evaluate the performance of our algorithm on a number of close-knit regional networks obtained from the Facebook social network comparing its performance with some well-known heuristics such as degree centrality as well other heuristics and show that in many cases heuristics perform comparable to the computationally more intense greedy method.

We start with a brief overview of information propagation in social networks in Section 2. In Section 3, we first introduce our model of communication and formalize the influence limitation problem. Later, in Section 4 we focus on the *eventual influence limitation* objective function and prove that this problem is NP-hard and submodular. Submodularity guarantees approximation bounds for a greedy algorithm presented in Section 4.3. In Section 5, we provide the experiments that compare the performance of the greedy solution with various heuristics. In Section 6, we present our conclusions. Finally in Section 7, we discuss future work.

2. RELATED WORK

The identification of *influential users* or *opinion leaders* in a social network is a problem that has received a significant amount of attention in recent research. In the *influence maximization problem*, given a probabilistic model of information diffusion such as Independent Cascade Model, a network graph, and a budget k, the objective is to select a set A of size k for initial activation so that the expected value of $f(A)$ (size of cascade created by selecting set A) is maximized [8, 31]. With an efficient, robust solution to this problem, it is possible to widely disseminate important information in a social network. Early works relied on heuristics such as node degree and distance centrality [35] to select the set A. Although the problem of finding an optimal solution in this model is NP-hard, it has been shown that there is a greedy algorithm that yields a spread that is within $1 - 1/e$ of optimal [17]. This solution depends on Monte Carlo simulations which are computationally intense. Work has been done on improving the performance of greedy algorithms for influence maximization [5, 22, 19], but scalability remains a significant challenge. In addition to the scale issues that are inherently there, these definitions of influential users ignore certain aspects of the real social networks such as the existence of competing campaigns. In this work we consider different models of communication that incorporate different aspects of real social networks. The works that are closest to the one introduced in this paper are [17, 22]. Similar to those works, we identify a problem in a social network that involves identifying "influential nodes" and study the feasibility of a solution to this problem. However, our problem formulation is more general in that, we model the existence of competing cascades dissipating in a network.

The existence of competing campaigns has been captured by a number of studies recently. Dubey et al. [9] study the problem as a network game focusing on quasi-linear model and consider various cost, benefit and externalities functions for competing firms. They study the existence of Nash Equilibrium (NE) and show that NE is unique if there is enough competition between firms or if their valuations of clients are anonymous. Bharathi et al. [3] augment the Independent Cascade Model to capture the existence of competing campaigns in a network. Their diffusion model is similar to the one studied in our work and captures the timing issues that are crucial to competing campaigns optimization problems. The algorithmic problem they approach is: Given that there is more than one campaign dissipating in a network and each campaign can select a set of early adopters so as to maximize their benefit, i.e. number of people adopting their product, what is the best strategy for the players? This work studies the problem from both the first and last player's perspectives and shows that the problem of selecting the early adopters for the last player is submodular. They also introduce a FPTAS for the first player when the network structure is a tree. Carnes et al. [4] consider the same problem from the last player's perspective and use one diffusion model where nodes of the network choose the campaign to adopt w.r.t. their distance to the early adopters of the campaigns and another model where the nodes make a uniform random choice among its active neighbors. They present experimental results that show that the greedy approach with the approximation bounds performs better than the heuristics but the difference is not significant which agrees with the results presented in this work. They also experimentally show that the best strategy for the first player is to choose high degree nodes. Kostka et al. [21] study competing campaigns also as a game theoretical problem and show that being the first player, i.e. the first to decide, is not always advantageous. Both [4, 3] use diffusion models where the competing campaigns propagate exactly the same way, i.e. the probability of diffusion on a certain edge is the same for all campaigns and all campaigns start at the same time. In our work, we study the case where the competing campaigns have different acceptance rates and one is in response to the other, and therefore campaign of the last player is started with a delay.

Modeling the behavior of social networks in the presence of competing campaigns is necessary but not sufficient to providing algorithms for limiting the spread of misinformation. The works introduced so far for competing campaigns attack the problem of influence *maximization* as opposed to *limitation*. However, in the presence of a misinformation campaign, the main goal is to minimize the number of people adopting the "bad" campaign rather than maximizing the number of people adopting the "good" campaign. In this work, we aim to provide solutions for *influence limitation* problem. The problem of limiting the effect of misinformation in a social network can be seen in a way similar to the problem of epidemics and inoculation. There are many studies on the spread of infections and immunization [34, 2, 18]. A recent work on identifying influential people in a social network [20] uses SIS(susceptible-infected-susceptible), SIR(susceptible-infected-recovered) models [1, 7, 15] and concludes that the influence of a node is more dependent on its location in the network than the number of connections it has. This work captures the notion of being "immunized" but the immunization ends with the node that is inoculated. Conversely, we consider the case where once a node is inoculated, it can inoculate more people (by spreading the "good" information). Spread of diseases and inoculation has also been studied in game theory literature. Meier et al. [24] studies inoculation games in social networks. The problem is posed again in terms of virus propagation where the owner of each node selfishly decides whether or not to protect itself. This work also assumes that inoculation has direct effect only on to the inoculated node, meaning that the "good" information does not propagate. The decision to "protect" oneself is a distributed process, each node decides for itself and is trying to maximize its own function whereas we consider the problem of finding the optimal solution for the community.

3. DIFFUSION OF MISINFORMATION

A social network can be modeled as a directed graph $G = (N, V)$ consisting of nodes N and edges V. A node v is a *neighbor* of w if and only if there is an edge from w to v in G. In the context of influence spread, the set of nodes, N can be viewed as the users of the social network. If a user w is a "friend" of another user v, there is a direct communication link, an edge $e_{v,w}$ in G. In addition to this, we assign a weight $p_{v,w}$ to each edge $e_{v,w}$ which is used to model the direct influence v has on w or conversely the probability that v will forward certain information it obtains to its neighbor w. Note that in this setting, "friendship" is an asymmetric relationship which enables us to model the case where the influence one user has on a friend is different than the effect this friend has on that user.

3.1 Diffusion Models

Independent cascade is one of the most basic and well-studied diffusion models that has been used in different contexts [10, 23, 13, 14]. In Independent Cascade Model, the process starts with an initial set of active nodes A_0, and the process unfolds in discrete steps according to the following randomized rule. When node v first becomes active in step t, it is given a single chance to activate each currently inactive neighbor w; it succeeds with a probability $p_{v,w}$ independent of the history thus far. If v succeeds, then w will become active in step $t + 1$; but whether or not v succeeds, it cannot make any further attempts to activate w in subsequent rounds. The process runs until no more activations are possible. If w has multiple newly activated neighbors, their attempts are sequenced in an arbitrary order.

We now introduce the *Multi-Campaign Independent Cascade (MCICM)* model. Here, we model the diffusion of two cascades evolving simultaneously in a network. Let C (stands for "campaign") and L (stands for "limiting campaign") denote the two cascades and the initial set of active nodes for cascade L is denoted by A_L. Similarly A_C denotes the initial set of active nodes in C. The process unfolds again in discrete time steps. When a node v first becomes active in campaign L (or C) in step t, it is given a single chance to activate each currently inactive neighbor w in campaign L (or C) and it succeeds with probability $p_{L,v,w}$ (or $p_{C,v,w}$) given that no neighbor of w tries activating w in the competing campaign at the same step. We also refer to $p_{L,v,w}$ (or $p_{C,v,w}$) as the probability of the edge $e_{v,w}$ being *live*. If there are two or more nodes trying to activate w at a given time step, at most one of them can succeed. In independent cascade, when w has several newly activated neighbors, their attempts are sequenced in arbitrary order. However in our studies, we will assume that there is a natural order to the two campaigns, more specifically one is "good" while the other is the "bad" campaign and if the "bad information" and the "good information" reach a node w at the same step, "good information" takes effect. Once a node becomes active in one campaign, it never becomes inactive or changes campaigns and the process continues until there is no newly activated node in either campaign.

We also consider another model of diffusion in which the probabilities of each edge being *live* is independent of the campaign. In this setting we only associate one probability $p_{v,w}$ with each edge $e_{v,w}$ and hence the model becomes almost identical to the *Independent Cascade Model*. No matter which information reaches a node v, v forwards this information to its neighbor w with probability $p_{v,w}$. Although this model is not a perfect fit for inoculation of misinformation (since nodes of the network would be more willing to share "good" information), it is a good fit for modeling competing campaigns where the two information cascades are more likely to be of similar "quality" and the nodes would agree to the campaign that reaches out to them first. Consider for example two articles L and C about the same event spreading through a social network. The probability of a user forwarding article L and C is more dependent on the news itself rather than which agency the news is from. Similar to the *Multi-Campaign Independent Cascade* model, there are three states a node can be in; *inactive, in campaign L, in campaign C* and once a node becomes active in either L or C, it cannot change its state again. As before, we assume that in the case of simultaneous trials of activation at a node, campaign L is ordered before C. We call this model *Campaign-Oblivious Independent Cascade (COICM)*. *COICM* is similar to the diffusion model used in [3]. However here we assume that one of the campaigns is prioritized over the other one in the case of simultaneous activation trials whereas independent and exponentially distributed continuous random variables are assigned to each edge as delay in [3] to make sure there will be no simultaneous activation trials. Note that, the algorithms presented in Section 4 would also work the diffusion model presented in [3].

3.2 Problem Definition

While a substantial amount of research has been done in the context of influence maximization, a problem that has not received much attention is that of limiting the influence of a malicious or incorrect information campaign. One strategy to deal with a misinformation campaign is to limit the number of users who are willing to accept and spread this misinformation. We will assume the *Multi-Campaign Independent Cascade Model* described in Section 3.1 as the model of communication. Without loss of generality we will assume that the spread of influence for campaign C starts from one node n_a and its existence is detected at time step r and at that point the campaign L is started. We refer to r also as the *delay* of campaign L. Our aim is to either limit the effect of campaign C or to maximize the effect of L depending on the specific objective function.

Depending on the context that the influence limitation problem is introduced in, we need to consider different objective functions. The objective can be to try and "save" as many people as possible, to limit the lifespan of the "bad" information campaign or to maximize the effect of the "good" campaign in the presence of the "bad" campaign. In the next section, we will focus on minimizing the number of people that end up adopting campaign C when information cascades from both campaigns are over. We refer to this problem as the *eventual influence limitation problem.* This objective function is our main focus since it has great implications in the area of limitation of effects of misinformation. In addition to the *eventual influence limitation problem (EIL)*, there are other variants of the competing campaigns problem which we briefly introduce in Section 7.

4. EVENTUAL INFLUENCE LIMITATION

Given a network and the *Multi-Campaign Independent Cascade Model* defined in Section 3.1, suppose that a campaign C that is spreading bad information is detected at round r. Given a budget k, select A_L as *seeds* for initial activation with the limiting campaign L such that the expected number of people that adopt campaign C, $\sigma(A_C)$ is minimized. Let $IF(A_C)$ denote the influence set of campaign C in the absence of campaign L, *i.e* the set of nodes that would accept campaign C if there were no limiting campaign. We define the function $\pi(A_L)$ to be the size of the subset of $IF(A_C)$ that campaign L prevents from adopting campaign C. Then, the influence limitation problem is equivalent to selecting A_L such that the expectation of $\pi(A_L)$ is maximized.

We now outline a potential solution to a simplified version of this problem. We assume that there is only a single source of information for campaign C, meaning $|A_C| = 1$. We refer to this single source as the *adversary node* or n_a As it may be much easier to convince a user of the truth than a falsehood, we also assume that the *limiting campaign information* is accepted by users with probability 1 ($p_{L,v,w} = 1$ if there is an edge from v to w and $p_{L,v,w} = 0$ otherwise). We refer to this notion as *high-effectiveness property.* Even with these restrictions, the problem of finding the optimal A_L is NP-hard and therefore finding the optimal solution is inefficient and infeasible. However as we will establish in Section 4.2, the problem can be shown to be submodular which guarantees that we can provide approximation bounds with a simple hill climbing approach.

Later we will investigate a more general form of this problem where we allow arbitrary values for $p_{L,v,w}$ and show that this problem is no longer submodular.

4.1 NP-Hardness of EIL

THEOREM 4.1. *The eventual influence limitation problem is NP-hard even with the high effectiveness property.*

PROOF. Consider an instance of the NP-complete Set Cover problem, defined by a collection of subsets $S_1, S_2, ..., S_m$ for a universe set $U = \{u_1, u_2, ..., u_n\}$; we wish to know whether there exist k of the subsets whose union is equal to U. We show that this can be viewed as a special case of the influence limitation problem. Given an arbitrary instance of the Set Cover problem, we define a corresponding directed bipartite graph with $n+m+1$ nodes: there is a node i corresponding to each set S_i, a node j corresponding to each element u_j, and a directed edge (i, j) whenever $u_j \in S_i$. In addition, there is an adversary node a and a directed edge (a, j) for all u_j with activation probability $p_{a,j} = 1$. The Set Cover problem is equivalent to deciding if there is a set A_L of k nodes in this graph with $\pi(A_L) \geq n + k$ when we become aware of campaign C at time step 0 (when a itself is active in campaign C but has not contacted any of its neighbors yet). Note that for the instance we have defined, activation is a deterministic process, as all probabilities for adversary to infect its neighbors are 0 or 1. Initially activating the k nodes corresponding to sets in a Set Cover solution results in saving all n nodes corresponding to the ground set U, and if any set A_L of k nodes has $\pi(A_L) \geq n + k$, then the Set Cover problem must be solvable. □

4.2 Submodularity of EIL

A function $f(.)$ is said to be submodular or have "diminishing returns" if it satisfies the following property: $f(S \cup v) - f(S) \leq f(T \cup v) - f(T)$, for all elements v and all pairs of sets $S \subset T$, i.e. the marginal gain from adding an element to a set S is at least as high as the marginal gain from adding the same element to a superset of S. As proved by Nemhauser, Wolsey, and Fisher [6, 29], for submodular functions, the greedy hill-climbing algorithm of starting with the empty set, and repeatedly adding an element that gives the maximum marginal gain approximates the optimum solution within a factor of $(1 - 1/e)$. Here we will prove that the influence limitation problem is submodular when the *limiting campaign* has the *high effectiveness property.*

Since influence spread over the graph G is a stochastic process, the influence function for a node or a set of nodes is tricky to define. Following the same approach presented in [17], we view an event of a newly activated node v attempting to activate its neighbor w and succeeding with $p_{C,v,w}$ as flipping a coin with bias $p_{C,v,w}$. Again, similar to [17], it does not matter whether the coin is flipped at the moment when v tries to activate w, or if it was pre-flipped and stored to be examined at the moment when v tries to activate w. So while considering a specific instance of influence spread, we can pre-flip all the coins to determine which edges of the graph G are *live* (meaning if the start node of this edge were to be activated, it would succeed in activating its neighbor) or *blocked* (meaning the attempt would be unsuccessful). In this setting, the spread of "bad campaign" C can be modeled as graph $G' = (N', V')$. where N' represents the set of nodes that are reachable from adversary node n_a via live edges and V' represents the set of live edges amongst the nodes in N'. Similarly, the spread of the limiting campaign L will be captured by the *inoculation graph* G''

We start with the problem of *eventual influence limitation* with the *high effectiveness property*, on a restricted instance where the graph of spread of campaign C, is a tree. Note that we are not necessarily interested in the number of *inoculated* nodes but the *inoculated* nodes that would be *infected* otherwise. We will refer to this set of nodes as *saved*. The *inoculation graph* will be used to determine which set of nodes can be *saved* given an initial set A_L.

Let I denote the set of nodes that are already active in campaign C, or in other words *infected* by time step r when the *limiting campaign* is started. Create a graph $G'' = (N'', V'')$ such that $N'' = \{u | u \in N \wedge u \notin I\}$ and $V'' = \{(u,v) | u \in N'' \wedge v \in S_u\}$ where $S_u = \{v | v \in N' \wedge |SP_G(u,v)| + r \leq |P_{G'}(n_a, v)|\}$ where $SP_G(v,w)$ represent shortest path from v to w in graph G, r denotes the *delay* and $P_{G'}(n_a, v)$ represents the unique path from n_a to v in the tree G'.

CLAIM 1. *A node u is saved if and only if there is a path from some node in A_L to u in G''*

If there is a path from a node v in A_L to u, then u is saved: In the first case, we know that u would be saved since u is in S_v which guarantees that the shortest path from v to u plus the delay r in G is shorter than the path from n_a to u in G'.

If there is no path from a node v in A_L to u in G'', then u is not saved: Proof by contradiction. Assume that there is no path from a node v in A_L to u in G'' and u is saved. There are two ways to save a node u. 1) To make sure u hears about campaign L before campaign C 2) To make sure campaign C does not reach u. For case 1) u must have a shorter path from a node v in A_L than $P_{G'}(n_a, v)$ in which case there would be an edge (v,u) in V''. For case 2) there must exist an ancestor n_j of u in G' s.t. $|SP_G(v, n_j)| + r \leq |P_{G'}(n_a, n_j)|$ since otherwise campaign C would reach out to node u from path $P_{G'}(n_a, u)$. W.l.o.g. let the path from adversary n_a to u in G' consist of nodes $n_a, n_1, n_2, ...n_j, n_{j+1}, ..., u$. Since $V' \subset V$, $|SP_G(v,u)| \leq |SP_G(v, n_j)| + |P_{G'}(n_j, u)|$, Thus $|SP_G(v,u)| + r \leq |P_{G'}(n_a, n_j)| + |P_{G'}(n_j, u)| = |P_{G'}(n_a, u)|$. In both cases, $\exists$ v such that $v \in A_L$ and $|SP_G(v,u)| + r \leq |P_{G'}(n_a, u)|$. Thus, $(v,u) \in V''$ which is a contradiction.

LEMMA 4.2. *Given a network $G = (N, V)$, the eventual influence limitation problem is submodular when the graph of spread of campaign C $G' = (N', V')$ is a tree and the limiting campaign L has the high effectiveness property.*

PROOF. As proved by the claim above, the problem of *saving* a node w is equivalent to the problem of reachability to w from A_L in G''. Therefore the *eventual influence limitation* problem is equivalent to maximizing the number of nodes reachable from the set A_L. Since the problem of reachability is submodular [17], *eventual influence limitation* problem is also submodular. □

The proof of submodularity for the eventual influence limitation problem where the structure of the spread of the campaign C is an arbitrary graph is not as intuitive. Consider the graph of 10 nodes represented in Figure 1(a). Assume that by pre-flipping the coins, we end up with probabilities such that the solid lines are *live* edges and dotted lines are *blocked* edges. In this case a campaign starting from adversary node 0 would reach nodes 0, 1, 2, 3 if there was no limiting campaign. A first look at this graph (or the general eventual limitation problem in general) suggests that in order to save node 3, we need to make sure both 1 and 2 should be saved (or 3 should be saved directly). Superficially, it would seem that submodularity is no longer viable. Since saving only 1 or 2 would not be sufficient to save 3, but their combination would. However, a closer look at this problem reveals that we do not need to secure all the possible paths to a node from an adversary but just the shortest path. If L can reach 3 before C, 3 can never be infected. For instance, for the campaign in Figure 1(a), if campaign L reaches node 1 by $r = 1$, it will be saved (since the adversary will reach node 1 at the same timestamp and by our assumption, when the bad and good campaigns reach a node at the same time stamp, the node accepts the good campaign). In this case the good campaign will reach node 3 at $r = 2$ and even if node 2 is not saved, that still guarantees that node 3 will be saved. Next we provide the formal proof of submodularity for eventual influence limitation. Note that the proof depends on the *high-effectiveness property* of the good campaign. Later on, we will show that when this property does not hold, eventual influence limitation is not, in general, a submodular function.

CLAIM 2. *In MCICM with the high effectiveness property a node w can be saved if and only if $\exists$ v such that $v \in A_L$ and $|SP_G(v,w)| + r \leq |SP_{G'}(n_a, w)|$*

PROOF. *1. If $\exists$ v such that $v \in A_L$ and $|SP_G(v,w)| + r \leq |SP_{G'}(n_a, w)|$, then w is saved:* Assume that such a v exists but w could not be saved. This is only possible if the bad campaign C reaches w strictly before L since otherwise w would be saved at $ts = |SP_G(v,w)| + r$. So there must exist a path $P_{G'}(n_a, w)$ from n_a to w such that $|P_{G'}(n_a, w)| \leq |SP_G(v,w)| + r$. Since $|SP_G(v,w)| + r \leq |SP_{G'}(n_a, w)|$, $|P_{G'}(n_a, w)| \leq |SP_{G'}(n_a, w)|$. This means there is a shorter path from n_a to w in G' than the shortest path which is a contradiction.

2: If $\nexists$ v such that $|SP_G(v,w)| + r \leq |SP_{G'}(n_a, w)|$, then w cannot be saved: Assume contrary, i.e. $\nexists$ v such that $SP_G(v,w) \leq SP_{G'}(n_a, w)$ and node w is saved. If w is saved, at least one of the nodes in one of those shortest paths must have been activated in campaign L since otherwise campaign C would propagate from any of those shortest paths to w and therefore infecting it. W.l.o.g. let one of the shortest paths from n_a to w consist of nodes $n_a, n_1, n_2..., n_i, w$ and $n_j \in SP_{G'}(n_a, w)$ denote such a node that is activated in L. Node n_j can only be activated in campaign L if campaign L reaches n_j at $ts \leq j$ because $SP_{G'}(n_a, j) = n_a, n_1, n_2, ..., n_{j-1}, n_j$. Therefore there must exist $v \in A_L$ s.t. $|SP_G(v, n_j)| + r \leq |SP_{G'}(n_a, n_j)|$. Since $|SP_G(n_j, w)| \leq |SP_{G'}(n_j, w)|$, $|SP_G v, w| + r \leq |SP_G(v, n_j)| + SP_G(n_j, w)| + r \leq |SP_{G'}(n_a, n_j)| + |SP_{G'}(n_j, w)| \leq |SP_{G'}(n_a, w)|$. This contradicts with our initial assumption that $\nexists$ v such that $|SP_G(v,w)| + r \leq |SP_{G'}(n_a, w)|$. □

THEOREM 4.3. *Eventual Influence Limitation is submodular for any form of spread of campaign C when the limiting campaign L has high-effectiveness property*

PROOF. Consider the *inoculation graph* $G'' = (N'', V'')$ such that $N'' = \{u | u \in N \wedge u \notin I\}$ and $V'' = \{(u,v) | (u,v) \in V' \vee v \in S_u\}$ where $S_u = \{v | v \in N' \wedge |SP_G(u,v)| + r \leq |SP_{G'}(n_a, v)\}|$. Based on Claim 2, the *eventual influence limitation* problem is equivalent to maximizing the number of nodes reachable from the set A_L in G'' and as established in [17], this problem is submodular. □

Unfortunately, the general *eventual influence limitation* problem where L does not have the *high-effectiveness property* is not in general submodular. Consider again the graphs in Figure 1. Further assume an instance of this problem where G'' representing the spread of influence for the good campaign L consists of nodes 1,2,5,6 and the edges $e_{5,1}, e_{6,2}$. (Note that if the good campaign had the *high-effectiveness property*, G'' would be identical to G) In this case, $f(5) = 1$, $f(6) = 2$, $f(5, 6) = 1, 2, 3$ since by using 5(6) as a seed, L can save node 1(2). (Since $e_{1,3}(e_{2,3})$ is not a live edge for campaign L, the good campaign will never reach node 3, and 3 will be infected by node 2(1) at the next time step.) On the other hand if $A_L = \{5, 6\}$ both 2,3 will be saved and since these are the only two nodes that could infect 1, node 1 will also be saved. This example shows that eventual influence limitation where *high-effectiveness property* does not hold is not in general a submodular problem.

Alternatively, consider the *Campaign-Oblivious Independent Cascade* introduced in Section 3.1 where the probabilities on the edges are campaign-independent. In this case we associate only one probability $p_{v,w}$ with each edge $e_{v,w}$. As stated before, this model fits competing campaigns where the two campaigns are trying to get users to adopt very similar products or ideas. In this case users are as likely to adopt campaign L as they would adopt campaign C. Note that, this model does not rely on either one of the campaigns being good or bad and therefore can be applied to any two competing campaigns.

Claim 3. *Eventual Influence Limitation is submodular for Campaign-Oblivious Independent Cascade*

Proof. Since a node can only be activated in one campaign one time per instance, an edge $e_{v,w}$ will only be visited at most once. Therefore, using the same idea presented in 4.2, we can pre-flip all the coins to determine which edges are *live* and which ones are *blocked* for an instance of influence dissemination from campaigns C and L. Consider the graph $G' = (N', V')$ where V' is the set of edges in V that are *live* and N' is the set of nodes that are reachable from adversary n_a via *live* edges. In this case, if L reaches a node n_i before C and therefore *saves* it, all the other nodes N_i that have node n_i in their shortest path from the adversary will also be saved since an edge exists for campaign L if it does for campaign C (since the liveness of edges are independent from the campaign). As we have proved in Theorem 4.3, this problem is in fact submodular. □

4.3 Greedy Solution for EIL

Considering the large scale of online social networks today and the NP-hardness of the eventual influence limitation problem, it is crucial to find efficient approximation algorithms that can handle the scale while still guaranteeing error bounds. Since the eventual influence maximization objective function is submodular, a hill climbing approach provides a $(1-1/e)$ approximation [6, 29] while being far less computationally expensive then the NP-hard optimal solution. Figure 2 provides this greedy algorithm that yields an A_L for which $\pi(A_L)$ is within $1 - 1/e$ of optimal.

The algorithm works for a given adversary a, *delay* r and budget k, i.e. number of nodes to initially activate in campaign L. Since independent cascade is a stochastic process, in order to compute π for a given set of nodes, it is necessary to run simulations for random assignments of edge probabilities. As shown in [17], π can be computed with high accuracy for a number of simulations in the order of thousands. As demonstrated in steps (6,7), we run the *InfLimit*

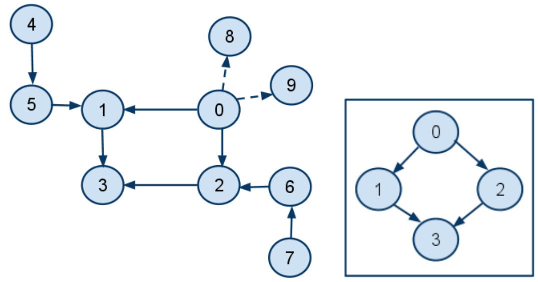

(a) A graph representing spread of campaign C. Solid lines represent the live edges and dotted lines represent dead edges for the spread of information campaign C. Assume that the adversary is node 0. In this case, if there was no opposing campaign, C would reach A_C = {0,1,2,3}

(b) The shortest path structure for spread of influence for the bad campaign

Figure 1: General Influence Spread

1: {Given (n_a,r,k) where n_a denotes the adversary and r denotes the time step campaign C is detected and k denotes the number of nodes to initially activate in L}
2: Initialize A_L to $\emptyset$ R to 10000
3: **for** $i = 1$ to k **do**
4:   **for** each vertex $v \in V - A_L$ **do**
5:     $s_v = 0$
6:     **for** j = 1 to R **do**
7:       $s_v += InfLimit(n_a, r, A_L, v)$
8:     $s_v = s_v/R$
9:   //Choose node i that maximizes $\pi(A_L \cup \{i\}) - \pi(A_L)$
10:   //And set $A_L \leftarrow A_L \cup \{i\}$
11:   $A_L = A_L \cup \{argmax_{v \in V - A_L}\{s_v\}\}$
12: Output A_L

Figure 2: Greedy algorithm to select the set for initial activation in the limiting campaign.

(influence limitation algorithm) for different random assignments for edge probabilities for each node. The procedure *InfLimit (n_a,r,S,v)* entails first assigning random numbers for each edge on the graph and simulating the influence limitation given that the adversary is node n_a, the adversary campaign is caught with delay r, the set of nodes we have already chosen to initially activate in campaign L is S and the node that we are evaluating the influence of is v. This method returns the marginal gain of adding node v to set S, i.e. number of people v could save the set S could not.

5. EVALUATION

Considering the large scale of social networks today and the complexity of the *eventual influence limitation* problem, even the greedy approach that is a polynomial time algorithm is too costly to be used in real social networks. Therefore, we seek other alternatives that can potentially compare well with the greedy approach which, as we have proved, is guaranteed to be a good appoximation. We consider three different heuristics. The first heuristic we consider is the *degree centrality* which has been used in early work to target "influential people".

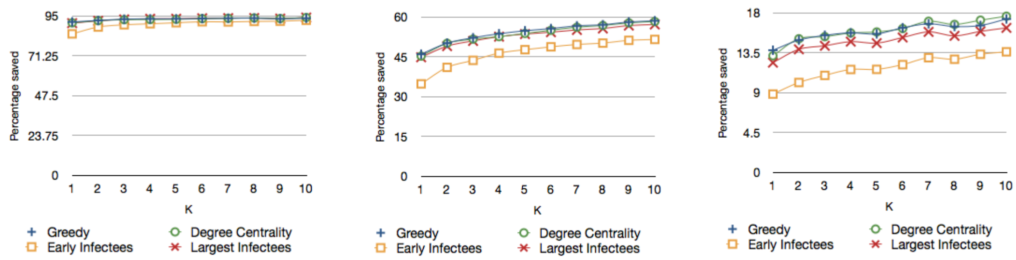

(a) $delay = 20\%$ (b) $delay = 35\%$ (c) $delay = 50\%$

Figure 3: Evaluation of Algorithms on SB08 for Multi-Campaign Independent Cascade Model with high effectiveness property

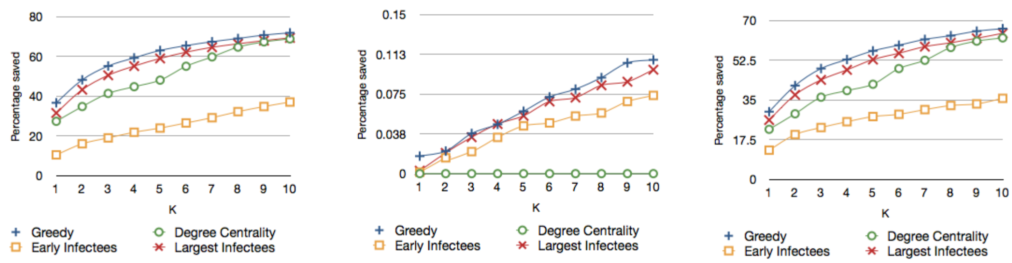

(a) $delay = 20\%$ (b) $delay = 70\%$ (c) $delay = 20\%$, adversary degree ≥ 40

Figure 4: Evaluation of Algorithms on SB08 for Campaign-Oblivious Independent Cascade model

The second heuristic we consider is *early infectees*. This notion refers to choosing *seeds* that are expected to be infected at time step r which we defined as the delay of campaign L. This is equivalent to reaching out to nodes that would be infected early on but after L is started, since those nodes are likely to create a large cascade for campaign C.

The third heuristic is *largest infectees*. This heuristic is very similar to the *early infectees* but rather than simply choosing the nodes that are expected to be infected early on, it aims to choose seeds that are expected to infect the highest number of people if they were to be infected themselves. In this case we restrict ourselves to such nodes that would be infected after timestep r. Note that both *early infectees* and *largest infectees* are more computationally intensive to compute than degree centrality. However they are still far less intense than the greedy method that involves shortest path computations. Computation of these heuristics is very similar to the problem of infection detection studied in [22] and has been shown to scale to very large networks.

Here we evaluate how well the greedy algorithm performs w.r.t. the three heuristics discussed. Note that since influence propagation is a stochastic process, we need to perform Monte Carlo Simulations in the order of thousands as part of our algorithm. This is one of the major scalability issues inherent in this type of problem. However, in our specific problem each simulation involves the expensive computation of shortest paths which is crucial to *eventual influence limitation* and this makes *eventual influence limitation* even more computationally intense then those of [17, 22]. As part of our experiments, we also evaluated how factors like the degree centrality of the adversary, *delay* of campaign L, and the weight distribution for $p_{C,v,w}$ and $p_{L,v,w}$ influence our choice of best fit algorithm. This requires running thousands of experiments on the same network data. Taking these factors into consideration we performed experiments on 4 regional network graphs obtained from Facebook. The data sets are as follows: 2009 snapshot of Santa Barbara regional network with 26455 nodes and 453132 edges; 2008 snapshot of the same network with 12814 nodes and 184482 edges; 2009 snapshot of the Monterey Bay regional network with 14144 nodes and 186582 edges; and 2008 snapshot of the same network with 6117 nodes and 62750 edges.

In Figure 3 we present our evaluation of the 4 methods on MCICM when L has the *high effectiveness property* using Santa Barbara 2008 data set. The y-axis represents the percentage of the population (that would be infected if there was no limiting campaign) that was saved. The x-axis represents the number of nodes that are initially activated in L (A_L). Figure 3(a) demonstrates the case where $delay = 20\%$ i.e. the ratio of the delay of the algorithm L to the duration of the campaign C is 0.2. In this case, all of the methods perform very well *saving* a big portion of the population. Figures 3(b) and 3(c) show the rapid decay of the performance of all the algorithms in the case where delay is 35% and 50% respectively. Here we have omitted the case where delay is 70%. In this case all of the algorithms were doing poorly, especially degree centrality had near-0 value. We conducted the same experiments on the other data sets and the result were similar for all of them. Due to space limitations, we have omitted the graphs for the other data sets. It is evident that for MCICM when L has the *high effectiveness property*, the biggest determining factor is how late the limiting campaign L is started. When L is started early, all the methods perform well whereas when the delay is large, all the algorithms perform poorly. For larger delays, greedy performs better than the other algorithms but the portion of the population saved is so small in all cases that this improvement is insignificant.

In Figure 4 we present our evaluation of the 4 methods on COICM using Santa Barbara 2008 graph. Figures 4(a) and 4(b) present the cases where the delay is 20% and 70% respectively. *Largest infectees* heuristic always performs very similar to the greedy method. *Degree centrality* performs well in general, especially if the delay r is small, but when r is large, meaning that we become aware of campaign C too late, degree centrality is not a good method for determining influential people. The reason for this is, degree centrality is a purely structural heuristic so we do not compute the expectance of being infected for the seeds we choose for A_L. When L is started too late, the highly connected nodes and their neighbors are more likely to have already been activated in campaign C. Comparing Figure 3(a) and 4(a), we observe the importance of *high effectiveness* since for the latter an average of 72% of the population can be saved with 10 seeds whereas the former shows consistent savings of 90-95% even with only one seed. Comparing Figures 4(a) and 4(b) again demonstrates the rapid decay of performance of all of the algorithms with larger delays for L. Figure 4(c) presents the case where the delay of L is 20% and n_a, the adversary that C starts from, has degree ≥ 40 All of the algorithms are less effective when the start node of C is a highly connected node, the reason behind this is that, highly connected adversary is likely to infect more people early on in which case the infection is more likely to grow exponentially and by time step r when L is started a large portion of the population is already infected.

Next, we evaluate MCICM where L does not have the *high effectiveness property*. In this case, the greedy algorithm is too costly to perform since many of the optimizations we performed for the earlier two cases cannot be applied here. Considering the results obtained from the earlier two sets of experiments, we conclude that, at least for close-knit social networks, the heuristics introduced above result in a good performance. Therefore we evaluated how well they perform on a slightly larger social network to see if there was consistency in their behavior. Figure 5 presents the test results for Santa Barbara 2009 data set. Figure 5(a) presents the case where the limiting campaign is started early (delay is 20%) but campaign C is more dominant ($0 \leq p_{C,v,w} \leq 0.5$ for all edges) than L ($0 \leq p_{L,v,w} \leq 0.1$ for all edges) in the sense that nodes are more likely to adopt C than L. Figure 5(b) presents the opposite case where L is more dominant than C. In both cases, the degree centrality and *largest infectees* heuristics have very similar behavior while *early infectees* performs worse than them. The savings are much larger for Figure 5(b) compared to Figure 5(a). We also note the similarity of Figure 5(b) with Figure 3(a), and claim that even if campaign L does not have the *high effectiveness property*, if it is more dominant than C, it is still likely to *save* a large portion of the population.

There are crucial lessons we can extract from the tests we performed. First, in almost all cases, *largest infectees* performs comparable with the greedy algorithm while being far less computationally intense. *Early infectees* heuristic, on the other hand, performs poorly since it *strictly* targets nodes that are expected to be infected at time step r which we observed to be an unstable heuristic. In many cases even the simpler heuristic of degree centrality is a good alternative that provides good results. Second, it is clear that there

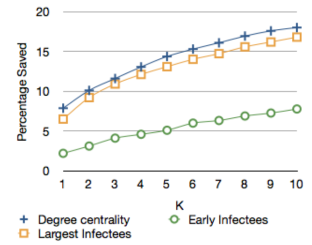

(a) *delay* = 20%, $0 \leq p_{C,v,w} \leq 0.5$, $0 \leq p_{L,v,w} \leq 0.1$

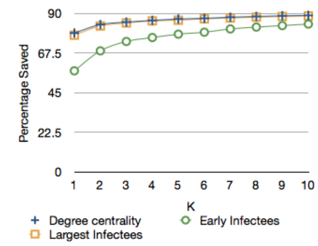

(b) *delay* = 20%, $0 \leq p_{C,v,w} \leq 0.1$, $0 \leq p_{L,v,w} \leq 0.5$

Figure 5: Evaluation of Algorithms on SB09 for Multi-Campaign Independent Cascade Model

are some parameters of the specific competing campaign, such as the delay of L, the connectedness of the adversary n_a, that are crucial to identify correctly to choose the right method for determining influential nodes for limiting a bad campaign C. For instance, in the case where the delay is too large, degree centrality is not a good option whereas it does perform very well for small delays. Having sufficient information about these parameters can help identify the correct method to apply to EIL.

6. CONCLUSION

In this work we have studied the algorithmic problem of limiting the effects of misinformation in a social network. More specifically, we investigated efficient solutions to the following question: Given a social network where a (bad) information campaign is spreading, who are the k "influential" people to start a counter-campaign (the good campaign) if our goal is to minimize the effect of the bad campaign? We call this *eventual influence limitation* problem. In order to study this problem, we first introduced a communication model of social networks that incorporates the notion of correlated campaigns that are disseminating simultaneously in a network. We proved that *eventual influence limitation* problem is NP-hard and therefore an exact solution is infeasible. We also showed that two variations of this problem on two different communication models are submodular and therefore a greedy method is guaranteed to provide a $1/(1-e)$ approximation. Although the greedy algorithm is polynomial time, it is still too costly for today's large scale social networks. Therefore, in addition to the approximation bounds, we also experimentally studied the performance of the greedy algorithm, comparing it with 3 different heuristics one of which is degree centrality. We showed that, in many cases, heuristics perform comparable to the greedy algorithm, even the simple degree centrality heuristic. This may seem counterintuitive at first glance since it does not adhere to many of the studies that claim poor performance for heuristics such as degree centrality. Note however that those performance results have been demonstrated on models of diffusion that do not capture the entire reality of social networks, i.e. the fact that there are multiple campaigns spreading simultaneously in a network. Despite the research claims of poor performance for such heuristics, marketeers have been using those heuristics for a very long time with the claim that "it works for them" [32]. This study provides insights as to why it works in reality. We also identified the cases where *degree centrality* is not a good heuristic and showed that in those cases, the *largest degree heuristic* still performs comparable to the greedy method while being computationally less intense. We studied different aspects of the problem such as the effect of starting the limiting campaign early/late, the effect of the properties of the adversary and how prone the population in general is to accepting either one of the campaigns.

7. FUTURE WORK

Even though it is evident from our experiments that the heuristics perform well for the various close-knit social networks, it is still an open issue if this also applies to networks of larger scale. As future work, we plan to evaluate the performance of our algorithms on networks of larger scale. For the larger networks, we also plan to work on pruning methods that will reduce the complexity. In addition to this, we plan to investigate the applicability of GPGPU computing and Map/Reduce platform to *eventual influence limitation* problem since the computation of $\pi(A_L \cup \{i\}) - \pi(A_L)$ for each node i can be parallelized.

In addition to this, we also plan to investigate other variations of competing campaign problems. One such variation could be *time sensitive influence limitation* where the objective is to find the minimum k such that when we initially activate A_L in campaign L where $|A_L| = k$, we can guarantee the number of steps campaign C spreads throughout the network is less than some constant t_{end}. We note that this function is also NP-hard but not submodular. Therefore providing an efficient algorithm with approximation bounds is a challenge. In the future we plan to study this objective function and evaluate performance of heuristics which we showed to be effective for *eventual influence limitation* problem. In addition to this, we plan to extend our model to study the cases where the nodes of the network can change their minds and switch from one campaign to the other and we plan to study use of nontrivial cost functions for the nodes of the network.

8. REFERENCES

[1] R. M. Anderson and R. M. May. *Infectious Diseases of Humans: Dynamics and Control (Oxford Science Publications)*. Oxford University Press, USA, new ed edition, September 1992.

[2] N. T. J. Bailey. *The mathematical theory of infectious diseases and its applications / [by] Norman T. J. Bailey*. Griffin, London :, 2nd ed. edition, 1975.

[3] S. Bharathi, D. Kempe, and M. Salek. Competitive influence maximization in social networks. In *In WINE*, pages 306–311, 2007.

[4] T. Carnes, C. Nagarajan, S. M. Wild, and A. van Zuylen. Maximizing influence in a competitive social network: a follower's perspective. In *ICEC '07: Proceedings of the ninth international conference on Electronic commerce*, pages 351–360, New York, NY, USA, 2007. ACM.

[5] W. Chen, Y.Wang, and S. Yang. Efficient influence maximization in social networks. In *Proceedings of the 15th ACM International Conference on Knowledge Discovery and Data Mining*, pages 199–208, 2009.

[6] G. Cornuejols, M. L. Fisher, and G. L. Nemhauser. Location of bank accounts to optimize float: An analytic study of exact and approximate algorithms. *Management Science*, 23(8):789–810, 1977.

[7] O. Diekmann and J. A. P. Heesterbeek. *Mathematical Epidemiology of Infectious Diseases: Model Building, Analysis and Interpretation (Wiley Series in Mathematical & Computational Biology)*. Wiley, 1 edition, May 2000.

[8] P. Domingos and M. Richardson. Mining the network value of customers. In *Proceedings of the 7th ACM International Conference on Knowledge Discovery and Data Mining*, pages 57–66, 2001.

[9] P. Dubey, R. Garg, and B. D. Meyer. Competing for customers in a social network. Cowles Foundation Discussion Papers 1591, Cowles Foundation, Yale University, Nov. 2006.

[10] R. Durrett. Lecture notes on particle systems and percolation. Wadsworth Publishing, 1988.

[11] Facebook. `http://www.facebook.com`.

[12] J. Garrison and C. Knoll. Prop. 8 opponents rally across california to protest gay-marriage ban. *Los Angeles Times*, November 2008.

[13] J. Goldenberg, B. Libai, and E. Muller. Talk of the network: A complex systems look at the underlying process of word-of-mouth. *Marketing Letters*, 2001.

[14] M. Granovetter. Threshold models of collective behavior. *American Journal of Sociology*, 1978.

[15] H. W. Hethcote. The mathematics of infectious diseases. *SIAM Rev.*, 42(4):599–653, 2000.

[16] K. M. Heussner. Ft. hood soldier causes stir on twitter. *ABS News (online)*, November 2009.

[17] D. Kempe, J. M. Kleinberg, and É. Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the Ninth ACM International Conference on Knowledge Discovery and Data Mining*, pages 137–146, 2003.

[18] J. O. Kephart and S. R. White. Measuring and modeling computer virus prevalence. In *SP '93: Proceedings of the 1993 IEEE Symposium on Security and Privacy*, page 2, Washington, DC, USA, 1993. IEEE Computer Society.

[19] M. Kimura, K. Saito, R. Nakano, and H. Motoda. *Social Computing and Behavioral Modeling*, chapter Finding Influential Nodes in a Social Network from Information Diffusion Data. Springer US, 2009.

[20] M. Kitsak, L. K. Gallos, S. Havlin, F. Liljeros, L. Muchnik, H. E. Stanley, and H. A. Makse. Identifying influential spreaders in complex networks. Jan 2010.

[21] J. Kostka, Y. A. Oswald, and R. Wattenhofer. Word

of Mouth: Rumor Dissemination in Social Networks. In *15th International Colloquium on Structural Information and Communication Complexity (SIROCCO), Villars-sur-Ollon, Switzerland*, June 2008.

[22] J. Leskovec, A. Krause, C. Guestrin, C. Faloutsos, J. VanBriesen, and N. Glance. Cost-effective outbreak detection in networks. In *Proceedings of the 13th ACM International Conference on Knowledge Discovery and Data Mining*, pages 420–429, 2007.

[23] T. M. Liggett. Interacting particle systems, 1985.

[24] D. Meier, Y. A. Oswald, S. Schmid, and R. Wattenhofer. On the windfall of friendship: inoculation strategies on social networks. In *EC '08: Proceedings of the 9th ACM conference on Electronic commerce*, pages 294–301, New York, NY, USA, 2008. ACM.

[25] Is social networking spreading misinformation? `http://wcbstv.com/topstories/twitter.facebook.misinformation.2.996194.html`.

[26] Michael jackson on tmz, iran on twitter: The social media news. `http://www.blogher.com/spreading-news`.

[27] E. Morozov. Swine flu: Twitter's power to misinform. *Foreign Policy*, April 2009.

[28] MySpace. `http://www.myspace.com`.

[29] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An analysis of approximations for maximizing submodular set functions. *Mathematical Programming*, 14(1):265–294, 1978.

[30] orkut. `http://www.orkut.com`.

[31] M. Richardson and P. Domingos. Mining knowledge-sharing sites for viral marketing. In *Proceedings of the 8th ACM International Conference on Knowledge Discovery and Data Mining*, pages 61–70, 2002.

[32] C. Thompson. Is the tipping point toast? *Fast Company*, January 2008.

[33] Twitter. `http:://www.twitter.com`.

[34] C. Wang, J. C. Knight, and M. C. Elder. On computer viral infection and the effect of immunization. In *ACSAC '00: Proceedings of the 16th Annual Computer Security Applications Conference*, page 246, Washington, DC, USA, 2000. IEEE Computer Society.

[35] S. Wasserman and K. Faust. *Social Network Analysis*. Cambridge University Press, 1994.