

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188		
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>					
1. REPORT DATE (DD-MM-YYYY) 13-09-2016		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 1-Aug-2014 - 31-Jul-2015	
4. TITLE AND SUBTITLE Final Report: Information Dynamics in Networks: Models and Algorithms			5a. CONTRACT NUMBER W911NF-14-1-0366		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER 611102		
6. AUTHORS Kamesh Munagala			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES Duke University C/O Office of Research Support 2200 W. Main St., Ste. 710 Durham, NC 27705 -4677			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSOR/MONITOR'S ACRONYM(S) ARO		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) 64485-NS.4		
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.					
14. ABSTRACT In this project, we investigated how network structure interplays with higher level processes in online social networks. We investigated the appropriateness of existing mathematical models for explaining the structure of retweet cascades on Twitter; we investigated how to detect spam accounts on Facebook and other social networks by graph analytics; and finally we investigated how to design pricing schemes for users on a social network when their valuations are influenced by their neighbors in the network.					
15. SUBJECT TERMS Networks, Cascades, Communities, Opinions					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Kamesh Munagala
a. REPORT UU	b. ABSTRACT UU	c. THIS PAGE UU			19b. TELEPHONE NUMBER 919-660-6598

Report Title

Final Report: Information Dynamics in Networks: Models and Algorithms

ABSTRACT

In this project, we investigated how network structure interplays with higher level processes in online social networks. We investigated the appropriateness of existing mathematical models for explaining the structure of retweet cascades on Twitter; we investigated how to detect spam accounts on Facebook and other social networks by graph analytics; and finally we investigated how to design pricing schemes for users on a social network when their valuations are influenced by their neighbors in the network.

Enter List of papers submitted or published that acknowledge ARO support from the start of the project to the date of this printing. List the papers, including journal references, in the following categories:

(a) Papers published in peer-reviewed journals (N/A for none)

Received

Paper

TOTAL:

Number of Papers published in peer-reviewed journals:

(b) Papers published in non-peer-reviewed journals (N/A for none)

Received

Paper

TOTAL:

Number of Papers published in non peer-reviewed journals:

(c) Presentations

Number of Presentations: 0.00

Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

<u>Received</u>	<u>Paper</u>
09/13/2016	1.00 . A Note on Modeling Retweet Cascades on Twitter, Workshop on Algorithms and Models for the Web Graph. 09-DEC-15, Eindhoven, Netherlands. : ,
09/13/2016	2.00 . Combating Friend Spam Using Social Rejections, 2015 IEEE 35th International Conference on Distributed Computing Systems (ICDCS). 29-JUN-15, Columbus, OH, USA. : ,
09/13/2016	3.00 . Value-Based Network Externalities and Optimal Auction Design, Conference on Web and Internet Economics. 14-DEC-14, Beijing, China. : ,
TOTAL:	3

Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Peer-Reviewed Conference Proceeding publications (other than abstracts):

<u>Received</u>	<u>Paper</u>
-----------------	--------------

TOTAL:

Number of Peer-Reviewed Conference Proceeding publications (other than abstracts):

(d) Manuscripts

<u>Received</u>	<u>Paper</u>
-----------------	--------------

TOTAL:

Number of Manuscripts:

Books

Received Book

TOTAL:

Received Book Chapter

TOTAL:

Patents Submitted

Patents Awarded

Awards

Graduate Students

NAME	PERCENT SUPPORTED	Discipline
Brandon Fain	0.20	
FTE Equivalent:	0.20	
Total Number:	1	

Names of Post Doctorates

NAME	PERCENT SUPPORTED
FTE Equivalent:	
Total Number:	

Names of Faculty Supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>	National Academy Member
Kamesh Munagala	0.25	
FTE Equivalent:	0.25	
Total Number:	1	

Names of Under Graduate students supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Student Metrics

This section only applies to graduating undergraduates supported by this agreement in this reporting period

The number of undergraduates funded by this agreement who graduated during this period: 0.00

The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:..... 0.00

Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):..... 0.00

Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense 0.00

The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields: 0.00

Names of Personnel receiving masters degrees

<u>NAME</u>
Total Number:

Names of personnel receiving PHDs

<u>NAME</u>
Total Number:

Names of other research staff

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Sub Contractors (DD882)

Inventions (DD882)

Scientific Progress

In this project we made 3 main contributions, which we list below. Much of this work is a combination of analytic modeling and data analysis of real social networks such as Twitter. This led to 3 publications in **peer-reviewed** conferences.

The first contribution is the study of information cascades. Such cascades on social networks, such as retweet cascades on Twitter, have been often viewed as an epidemiological process, with the associated notion of virality to capture popular cascades that spread across the network. The notion of structural virality (or average path length) has been posited as a measure of global spread. We argue that this simple epidemiological view, though analytically compelling, is not the entire story. We first show empirically that the classical SIR diffusion process on the Twitter graph, even with the best possible distribution of infectiousness parameter, cannot explain the nature of observed retweet cascades on Twitter. More specifically, rather than spreading further from the source as the SIR model would predict, many cascades that have several retweets from direct followers, die out quickly beyond that.

We show that our empirical observations can be reconciled if we take interests of users and tweets into account. In particular, we consider a model where users have multi-dimensional interests, and connect to other users based on similarity in interests. Tweets are correspondingly labeled with interests, and propagate only in the subgraph of interested users via the SIR process. In this model, interests can be either narrow or broad, with the narrowest interest corresponding to a star graph on the interested users, with the root being the source of the tweet, and the broadest interest spanning the whole graph. We show that if tweets are generated using such a mix of interests, coupled with a varying infectiousness parameter, then we can qualitatively explain our observation that cascades die out much more quickly than is predicted by the SIR model. In the same breath, this model also explains how cascades can have large size, but low “structural virality” or average path length. This work appeared in WAW 2015.

Our second contribution is a graph theoretic measure to weed out spam accounts in social networks. Unwanted friend requests in online social networks (OSNs), also known as friend spam, are among the most evasive malicious activities. Friend spam can result in OSN links that do not correspond to social relationship among users, thus pollute the underlying social graph upon which core OSN functionalities are built, including social search engine, ad targeting, and OSN defense systems. To effectively detect the fake accounts that act as friend spammers, we propose a system called Rejecto. It stems from the observation on social rejections in OSNs, i.e., even well-maintained fake accounts inevitably have their friend requests rejected or they are reported by legitimate users. Our key insight is to partition the social graph into two regions such that the aggregate acceptance rate of friend requests from one region to the other is minimized. This design leads to reliable detection of a region that comprises friend spammers, regardless of the request collusion among the spammers. Meanwhile, it is resilient to other strategic manipulations. To efficiently obtain the graph cut, we extend the Kernighan-Lin heuristic and use it to iteratively detect the fake accounts that send out friend spam. Our evaluation on real social network data shows that Rejecto can discern friend spammers under a broad range of scenarios and that it is computationally practical. This work appeared in ICDCS 2015.

Our final contribution considers designing pricing schemes in a social network. We study the problem of maximizing auctioneer revenue in settings where agents’ valuations exhibit positive network externalities, meaning that agents get positively influenced by social neighbors who also have the item. We give a complete characterization of incentive compatible auctions in this setting. Using this characterization, we show that the optimal auction can be computed in polynomial time when the agents’ signals are independent. We further show a constant factor approximation when the signals of agents are correlated, and an optimal mechanism in this case for a constant number of bidders.

Technology Transfer

ARO Final Report; Project 64485NS

Kamesh Munagala

The goal of this proposal was to explore the structure of *networked interactions* at scale, and how algorithms can be designed to exploit network structure, prevent malicious human behavior, or design pricing schemes for selling products. There are three separate sub-projects we undertook, that we detail below. These are based on our papers [2, 1, 3] that we have attached with the report.

Project 1: Retweet Cascades on Twitter

This work is presented in the paper [2]. Information cascades are among the most widely studied phenomena in social networks. There is a vast literature on modeling the spread of these cascades as diffusion processes, studying the kinds of diffusion trees that arise, as well as trying to predict the global spread (or *virality*) of these cascades. A specific example of such a diffusion process, which is the focus of our research, are retweet cascades on Twitter.

Problem Statement

Extant models of information cascades build on classical epidemiological models for spread of infectious diseases. The simplest of these is the SIR model, where a node in the network can be in one of three states at any time: *Susceptible* (S); *Infected* (I); and *Recovered* (R). Nodes in the network switch their states due to infections transmitted over the network, and the rate of these infections is governed by an infectiousness parameter, p . The SIR model unfolds via the following process: all nodes are initially in state S except the source (or a set of nodes called the “seed set”), which is in state I . Every node which is in state I infects each of its neighbors independently with probability p , before moving itself to state R . If a node in state S gets infected, it moves to state I . This process naturally quiesces with all nodes settling in their final state, and all nodes that were ever in state I are considered to have acquired the infection. There is a natural and trivial mapping of this model to information cascades, where the infectiousness parameter p serves to measure the *interestingness* of the piece of information, in our case, a tweet. In epidemiology, the goal is to differentiate infections that die out quickly from those that spread to the whole network; analogously, information cascades are deemed *viral* if their global reach is large.

The above view of information cascades as the spreading of content through the network is intuitively and analytically appealing. In fact, Goel *et al.* show that when simulated on a scale-free graph, the SIR model statistically mimics important properties of retweet cascades on Twitter. In particular, they use *structural virality*, or average path length in the diffusion tree, as a quantitative measure of “infectiousness” of a cascade, and show that the distribution of cascade sizes (number of users that retweet a tweet plus the author of the tweet) and structural virality are statistically similar to that from the simulations. On the other hand, these empirical studies also show that cascades observed in Twitter are mostly shallow and exceedingly rare: Goel *et al.* show there are no viral cascades in a corpus of a million tweets; and in subsequent work, they show that viral cascades do indeed exist if the corpus size is increased to a billion tweets.

This data contrasts with the observation that social networks like Twitter have a power-law degree distribution, and these networks should have low epidemic threshold, so that even with low infectiousness parameter p , most cascades should be viral. Therefore, explaining the low frequency of viral events on Twitter via an SIR model requires that the infectiousness parameter be quite low almost all the time. Finally, this result also begs the question of whether modeling viral events if even of any interest if these events are so rare.

We therefore asked the following question:

Is there something fundamental about real-world information cascades, particularly those on Twitter, that is not captured by the simple SIR model?

Research Methodology and Contributions

Though this question is about a specific social network, and a specific (simplistic) epidemiological model, even understanding this via suitably designed experiments is challenging, and has not been performed before. In the process of answering the above question, we make the following contributions.

Evaluating Epidemic Models Through Twitter Network. Our main contribution was to show that the SIR model is a *poor* fit for information flow on Twitter. We showed this by empirically testing the hypothesis that retweet cascades on Twitter propagate using the SIR process. Our null hypothesis was that each cascade has an underlying infectiousness p (that could be different for different cascades), and conditioned on receiving the tweet, a user retweets it with probability p . We compared the value of p that we obtain by best-fit for the users directly connected to the source of the tweet (level 1 followers), and those who receive the tweet from a direct follower of the source (level 2 followers). Using a corpus of 8 million cascades, we developed a statistical test to show that these two values of p are different – the second level value is significantly smaller than the first. The technically interesting part of this analysis is the fact that most cascades are shallow. Thus, many tweets generate very few retweets at the first level, and this number dictates the number of tweet impressions and retweets at the second level. The SIR model therefore corresponds to a stochastic process for the retweets that has very low mean but potentially very high variance because of the skewed degree distribution of the graph. We had to therefore devise a statistical test that works around this high variance. Apart from this statistical test, at a coarse level, we found that the median value of first level infection probability is 0.00046, while the median value of second level infection probability is 0 (in other words, half of the tweets do not have second level retweets!). Even among the tweets that have at least 1000 impressions at the first level, more than 80% of them, have that first level p is at least twice the second level p . This suggests that, rather than spreading further from the source, a cascade typically dies out quickly within a few hops. Indeed, the median of first level impressions is 175, while the median of second level impressions is 29. It also suggests an explanation in previous work for truly viral cascades being so rare.

Interest-based SIR Model. Since the SIR model assumption of fixed propagation probability per cascade is statistically violated on Twitter, we proposed an alternative model for retweet cascades. In particular, we presented a tweet propagation model that takes *interests* of users and tweets into account. In order to do this, we revisited a Kronecker graph-based model for social networks first considered in previous work by the PI. In this attribute based model, users have attribute vectors in some d -dimensions, and interests are specified by a subset of these dimensions along with their attribute values. If fewer dimensions are specified, these interests are *broad* and

encompass many users; if many dimensions are specified, these interests are *narrow* with a shallow component around the source. Tweets are also correspondingly labeled with interests, and propagate only in the *subgraph of interested users* via a SIR process with infectiousness drawn from a distribution. We show that if tweets are generated using such a mix of narrow and broad interests, then this coupled with a varying infectiousness parameter can qualitatively explain the level-one infectiousness being larger than subsequent levels. As a simple intuition, observe that cascades corresponding to narrow interests only reside in their shallow subgraphs, while those corresponding to broad interests can be “viral” in the usual sense.

As mentioned above, Goel *et al.* define the notion of *structural virality*, or average path length of a cascade as a measure of its virality. They show that this measure is uncorrelated with the size of the cascade, except when structural virality is large. The proposed explanation in their work is an SIR model on a scale-free graph with extremely low infectiousness parameter. Our model leads to a different explanation: cascades corresponding to narrow interests have low structural virality, but can have large size. This explanation does not depend on any specific setting of the infectiousness parameter, and is therefore of independent interest. Finally, we show that cascades arising for broad interests can have large structural virality, but our model would predict a large expected size as well, which again matches previous empirical findings.

Project 2: Combatting Friend Spam

This work is presented in the paper [1]. Unscrupulous users increasingly find Online Social Networking (OSN) platforms as lucrative targets for malicious activities, such as sending spam and spreading malware. The profitability of such activities and the fact that a large portion of the OSN communication takes place over symmetric social links (e.g., Facebook) motivate attackers to connect to real users. In particular, attackers leverage the open nature of OSNs and send to legitimate users unwanted friend requests, also known as *friend spam*.

Problem Statement and Motivation

Friend spam can result in OSN links that do not correspond to social relationship among users, thus it enables the pollution of the underlying undirected social graph. Because OSN providers build on social graphs their core functionalities that often assume a social graph solely consists of links representing social trust of user pairs, the consequences of falsely accepted requests by unsuspected users are severe. In particular, the false OSN links resulting from friend spam can compromise the accuracy of social ad targeting and search, and the privacy of shared content by users. Moreover, friend spam can be used to undermine the effectiveness of defense systems that are either built upon, or take input signals from social graphs. For example, the additional OSN links that fake accounts (called Sybils) obtain via friend spam can enable part of them to evade the detection of social-graph-based defense systems.

Given how central online social networks have become to information dissemination and opinion formation, preventing malicious behavior in such networks is of paramount importance. We therefore aimed to answer the question:

How can we design a robust system to throttle friend spam in online social networks?

Our **hypothesis** was that we could uncover the fake accounts (Sybils) that indiscriminately send out friend spam in symmetric networks (e.g., Facebook and LinkedIn) by leveraging the rejection of unwanted friend requests. The insight is that although some OSN users accept friend requests from

unknown users, cautious users reject, ignore, or report them to the OSN providers. The attackers behind the spamming accounts usually have limited knowledge about the degree of their targets’ security awareness, due to the massive scale of today’s OSNs. As a result, the spamming accounts inevitably receive a significant number of social rejections from legitimate users. We confirmed this in our study on fake Facebook accounts in the wild. In contrast, friend requests from legitimate users are only sporadically rejected because they are often sent to people the senders know.

Although using social rejections to combat friend spam is intuitive, designing a robust scheme that is strategy-proof poses an algorithmic challenge. First, the spamming fake accounts can attempt to evade the detection by collusion. Specifically, they can accept each other’s requests, decreasing the fraction of the rejected requests of each individual account to that of a legitimate user’s. Second, a part of the fake accounts can mimic legitimate users by rejecting friend requests from other fake accounts. By doing so the attacker sacrifices his accounts that got the rejections. Yet, he whitewashes his rejecting accounts because they now reject requests in the same way legitimate users do. We called this strategy *self-rejection*.

Research Methodology and Contributions

We designed and implemented a system called *Rejecto*. It exploits the readily available social rejections in OSNs and systematically uncovers the fake accounts that act as friend spammers. Rejecto monitors the friend requests sent out by users and augments the social graph with directed social rejections. Once an OSN detects the fake accounts used for friend spam, it can prevent them from sending requests in the future. The goal of our system is to be able to effectively identify fractions of users that participate in.

Graph Partitioning Approach. Rejecto systematically addresses the above mentioned security challenges. To be resilient to collusion, we formulated the friend spammer detection as a graph partitioning problem. Specifically, while the portion of the accepted friend requests among legitimate accounts is high, the aggregate acceptance rate of all the requests sent from fake to legitimate accounts is substantially lower, regardless of the acceptance of the requests among the fake accounts. Therefore, Rejecto extends the Kernighan-Lin graph partitioning approach and uses it to partition the social graph into two regions such that the aggregate acceptance rate of the requests from one region to the other is minimized. It then declares the accounts in the region with the minimum aggregate acceptance rate as *suspicious*. Because this aggregate acceptance rate is independent of the requests and links among fake accounts, an attacker cannot arbitrarily boost this rate by having his accounts befriend each other.

To mitigate the impact of the self-rejection attack strategy and cope with the existence of multiple independent groups of fake accounts, it applies the graph partitioning multiple times and iteratively identify fake-account groups after pruning the detected ones from the social graph.

Implementation. We implemented Rejecto on Spark [36], an efficient in-memory large-data processing platform. We evaluate Rejecto through extensive simulations on real social graphs, and we showed that it can withstand friend spam under a broad range of scenarios and that it is resilient to attack strategies. As a demonstration of Rejecto’s effectiveness on improving OSN-related services, we applied it for an in-depth defense against OSN fake accounts in combination with existing social-graph based schemes. The details can be found in the attached paper [1]. In particular, we showed that Rejecto has better precision/recall performance under a wide range of simulation settings compared to the current standard method called VoteTrust.

Furthermore, OSN providers can apply Rejecto to the detection of other malicious accounts, such as compromised ones. If compromised accounts are manipulated to pollute the social graph via friend spam, their requests follow Rejecto’s friend spam model. Therefore, they are exposed to Rejecto’s detection. In such a deployment scenario, the OSN provider can bucket friend requests and rejections according to the time intervals in which they have occurred, and then run Rejecto on an augmented graph constructed from the bucketed requests and rejections in each interval. This enables Rejecto to detect compromised accounts in post-compromise intervals.

Project 3: Auctions with Network Externalities

This work is presented in [3]. There are many goods and services for which the utility of an individual consumer increases with the number of consumers in their social network using the same good or service. This phenomenon is called *positive externalities* in the economics literature. There have been extensive studies on various settings of positive externalities in both the economics and computer science communities. Such a model of externality is motivated by several factors:

- The physical effect of the number of buyers on the quality of the good. For example choosing a phone plan over other competing brands depends on the number of users each network has.
- An indirect effect gives rise to consumption externalities. For example the amount of software available in different operating systems is a function of the number of people using them.
- The availability and quality of post-purchase services depend on the size of the community using the good.

Problem Statement

Most of the literature so far has focused on the cardinality based utility model, which assumes the externality leads to an additive increase in value depending on the number of users who obtain service. This implicitly assumes agents are identical in terms of how much they use the service. In several scenarios, the extent to which agents use the service is itself a function of their intrinsic value for the good. This motivates us to introduce the *value based* externality model.

As an example to illustrate the usefulness of such a model, suppose an agent is deciding to adopt a social network. The agent’s type is her signal s_i . This signal stands for how much she plans to use social networks. Furthermore, the agent receives externalities if her friend j also uses the same social network. The amount of externality received by i from j is determined by how much j plans to use the same network (call this $g_{ij}(s_j)$). Therefore, under the value based utility model, the agent’s utility depends linearly on her friends’ private information about how much they use the social network.

More formally, agents demand one unit of the item, which is available in unlimited supply. Given the agents’ intrinsic types $\{s_i\}$, suppose the set of agents winning the item is W . Then, the valuation for agent $i \in W$ is

$$v_i(s_i, W) = h_i(s_i) + \sum_{j \in W, j \neq i} g_{ij}(s_j)$$

On the other hand, for $i \notin W$, we have $v_i(s_i, W) = 0$. Since we consider positive externalities, we assume the functions h_i and g_{ij} are non-negative and non-decreasing.

We consider the Bayesian auction design setting, where the agents’ intrinsic types are assumed to be drawn from a distribution known to the auctioneer. The goal of the auctioneer is to design

an incentive compatible (truthful) and individually rational mechanism that optimizes expected revenue, where the expectation is over the distribution of types, as well as the randomness introduced by the mechanism. Our solution concept is ex-post, meaning that even when agents know the signals of the other agents, they report their signals truthfully, and the price charged to them is never more than their value for the product.

Results and Methodology

In our work, we presented a characterization of truthful mechanisms in the above setting. Using this characterization, we showed that when the intrinsic types of the agents are drawn from certain natural classes of distributions, the optimal auction (which agents to allocate the goods to and at what price) can be efficiently computed. Our characterization allows us to perform a standard transform from values to virtual values that is widely used in mechanism design. However, unlike standard mechanism design, the presence of values in the externality function causes some mathematical complications in algorithm design. Despite this difficulty, we show that the winning set can be computed by a clever adaptation of the densest subgraph algorithm in order to compute the minimal densest subgraph, and this allows us to design the optimal efficient computable mechanism. We further showed that the optimal ex-post incentive compatible and individually rational mechanism is always deterministic, meaning that it finds one winning set and one collection of prices, which is quite surprising in this space of problems. The details of the analysis and model are presented in [3].

References

- [1] Q. Cao, M. Sirivianos, X. Yang, and K. Munagala. Combating friend spam using social rejections. In *35th IEEE International Conference on Distributed Computing Systems, ICDCS 2015, Columbus, OH, USA, June 29 - July 2*, pages 235–244, 2015.
- [2] A. Goel, K. Munagala, A. Sharma, and H. Zhang. A note on modeling retweet cascades on twitter. In *Algorithms and Models for the Web Graph - 12th International Workshop, WAW 2015, Eindhoven, The Netherlands, December 10-11*, pages 119–131, 2015.
- [3] K. Munagala and X. Xu. Value-based network externalities and optimal auction design. In *Web and Internet Economics - 10th International Conference, WINE 2014, Beijing, China, December 14-17*, pages 147–160, 2014.