

On the Windfall of Friendship: Inoculation Strategies on Social Networks

Dominic Meier¹, Yvonne Anne Oswald¹, Stefan Schmid², Roger Wattenhofer¹

¹Computer Engineering and Networks Laboratory, ETH Zurich, Switzerland

²Institut für Informatik, Technische Universität München, Germany

meierdo@ethz.ch, yoswald@ethz.ch, schmiste@in.tum.de,
wattenhofer@tik.ee.ethz.ch

ABSTRACT

This paper studies a virus inoculation game on social networks. A framework is presented which allows the measuring of the windfall of friendship, i.e., how much players benefit if they care about the welfare of their direct neighbors in the social network graph compared to purely selfish environments. We analyze the corresponding equilibria and show that the computation of the worst and best Nash equilibrium is \mathcal{NP} -hard. Intriguingly, even though the windfall of friendship can never be negative, the social welfare does not increase monotonically with the extent to which players care for each other. While these phenomena are known on an anecdotal level, our framework allows us to quantify these effects analytically.

Categories and Subject Descriptors

F.2.2 [Theory of Computation]: Analysis of Algorithms and Problem Complexity—*Nonnumerical Algorithms and Problems*; G.2.2 [Discrete Mathematics]: Graph Theory

General Terms

Theory, Security, Human Factors, Economics

Keywords

Game Theory, Social Networks, Equilibria, Virus Propagation, Windfall of Friendship

1. INTRODUCTION

Social networks have existed for thousands of years, but it was not until recently that researchers have started to gain scientific insights into phenomena like the *small world property*. The rise of the Internet has enabled people to connect with each other in new ways and to find friends sharing the same interests from all over the planet. A social network on the Internet can manifest itself in various forms. For instance, on *Facebook*, people maintain virtual references

to their friends. The contacts stored on mobile phones or email clients form a social network as well. The analysis of such networks is an interesting endeavor, as they comprise many aspects of our society in general.

A classic tool to model human behavior is *game theory*. It has been a fruitful research field in economics and sociology for many years. Recently, computer scientists have started to use game theory methods to shed light onto the complexities of today's highly decentralized networks. In game theoretic models, one traditionally assumes that people act autonomously and are steered by the desire to maximize their benefits (or utility). Under this assumption, it is possible to quantify the performance loss of a distributed system compared to situations where all participants collaborate perfectly. A widely studied measure which captures this loss of social welfare is the *Price of Anarchy* (PoA). Even though these concepts can lead to important insights in many environments, we believe that in some situations, the underlying assumptions do not reflect reality well enough. One such example are social networks: most likely people act less selfishly towards their friends than towards complete strangers. Such altruistic behavior is not considered in typical game-theoretic models.

This paper aims at combining these two active threads of research: social networks and game theory. We introduce a framework taking into consideration that people may care about the well-being of their friends. In particular, we define the *Windfall of Friendship* (WoF) which captures to what extent the social welfare improves in social networks compared to purely selfish systems. In order to illustrate our techniques we provide a game-theoretic analysis of a *virus inoculation game*.

Social networks are not only attractive to their participants, e.g., it is well-known that the user profiles are an interesting data source for the PR industry to provide tailored advertisements. Moreover, social network graphs can also be exploited for attacks, e.g., email viruses using the users' address books for propagating, worms spreading on mobile phone networks and over the Internet telephony tool Skype have been reported (e.g., [8]).

In this paper, we investigate the propagation of such viruses on social networks. Concretely, we assume that the users have the choice between risking infection and inoculating by buying anti-virus software. As expected, our analysis reveals that the players in this game always benefit from caring about the other participants in the social network rather than being selfish. Intriguingly, however, we find that the Windfall of Friendship does not increase monotonically with

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'08, July 8–12, 2008, Chicago, Illinois, USA.

Copyright 2008 ACM 978-1-60558-169-9/08/07 ...\$5.00.

stronger relationships. Despite of the phenomenon being an “ever-green” in political debates, to the best of our knowledge, this is the first paper to quantify this effect formally.

As another contribution, this paper shows that computing the best and the worst friendship Nash equilibrium is \mathcal{NP} -hard. In addition, simple networks such as complete graphs and stars are considered, and upper and lower bounds are derived. For example, we show that the Windfall of Friendship in a complete graph is at most $4/3$; this is tight in the sense that there are problem instances where the situation can indeed improve this much. Moreover, we show that in star graphs, friendship can help to eliminate undesirable equilibria. Generally, we discover that even in simple graphs the Windfall of Friendship can attain a large spectrum of values, from constant ratios up to $\Theta(n)$, n being the network size, which is asymptotically maximal for general graphs.

The remainder of this paper is organized as follows. In Section 2, we review related work in the area of social networks and game theory. Section 3 formally introduces our model and framework. Our main contributions are presented in Sections 4 and 5 where results for general and special graphs are derived. Finally, we conclude the paper in Section 6.

2. RELATED WORK

Social networks are currently a hot topic not only in social sciences, but also in ethnology, psychology and communication sciences. Computer scientists have become interested in these networks as well, and many social networks exist today on the Internet, e.g., *Facebook*, *LinkedIn*, *MySpace*, *Orkut*, or *Xing*, to name but a few. The famous *small world experiment* [23] conducted by Stanley Milgram 1967 has gained attention by the algorithm community [17] and inspired research on topics such as decentralized search algorithms [18, 22], routing on social networks [9, 17, 21] and the identification of communities [7, 26]. The dynamics of epidemic propagation of information or diseases has been studied from an algorithmic perspective as well [19, 20]. Knowledge on effects of this cascading behavior is useful to understand phenomena as diverse as word-of-mouth effects, the diffusion of innovation, the emergence of bubbles in a financial market or the rise of a political candidate. It can also help to identify sets of influential nodes in networks where marketing is particularly efficient (*viral marketing*). For a good overview on economic aspects of social networks, we refer the reader to [4], which, inter alia, compares random graph theory with game theoretic models for the formation of social networks.

Recently, game theory has also received much attention by computer scientists. Today, many different actors and stake-holders influence the decentralized growth of the Internet. Game theory is a useful tool to gain insights into its socio-economic complexity. Many aspects have been studied, e.g., *routing* [27, 28], *multicast transmissions* [6], or *network creation* [5, 24].

This paper seeks to apply game theory to social networks where players are not completely selfish and autonomous but have friends about whose well-being they care to some extent. We exemplify our mathematical framework with a virus inoculation game on social graphs. There is a large body of literature on the propagation of viruses [3, 10, 15, 16, 29]. Miscellaneous misuse of social networks has been re-

ported, e.g., *email viruses*¹ have used address lists to propagate to the users’ friends. Similar vulnerabilities have been exploited to spread worms on the *mobile phone network* [8] and on the Internet telephony tool *Skype*².

The papers closest to ours are [1, 25]. Our model is inspired by Aspnes et al. [1]. The authors apply a classic game-theoretic analysis and show that selfish systems can be very inefficient, as the Price of Anarchy is $\Theta(n)$, where n is the total number of players. They show that computing the social optimum is \mathcal{NP} -hard and give a reduction to the combinatorial problem *sum-of-squares partition*. They also present a $O(\log^2 n)$ approximation. Moscibroda et al. [25] have extended this model by introducing malicious players in the selfish network. This extension allows us to estimate the robustness of a distributed system to malicious attacks. They also find that in a non-oblivious model, intriguingly, the presence of malicious players may actually *improve* the social welfare. The *Windfall of Malice* has also been studied in the context of congestion games [2] by Babaioff et al. In contrast to these papers, our focus here is on social graphs where players are concerned about their friends’ benefits.

Finally, there is other literature on game theory where players are influenced by their neighbors. In *graphical economics* [12, 14], an undirected graph is given where an edge between two players denotes that free trade is allowed between the two parties, where the absence of such an edge denotes an embargo or an other restricted form of direct trade. The payoff of a player is a function of the actions of the players in its neighborhood only. In contrast to our work, a different equilibrium concept is used and no social aspects are taken into consideration.

3. MODEL

This section introduces our framework. In order to gain insights into the Windfall of Friendship, we study a virus inoculation game on a social graph. We present the model of this game and we show how it can be extended to incorporate social aspects.

3.1 Virus Inoculation Game

The virus inoculation game was introduced by [1]. We are given an undirected network graph $G = (V, E)$ of $n = |V|$ players (or nodes) $p_i \in V$, for $i = 1, \dots, n$, who are connected by a set of edges (or *links*) E . Every player has to decide whether it wants to *inoculate* (e.g., purchase and install anti-virus software) which costs C , or whether it prefers saving money and facing the risk of being infected. We assume that being infected yields a damage cost of L (e.g., a computer is out of work for L days). In other words, an instance I of a game consists of a graph $G = (V, E)$, the inoculation cost C and a damage cost L . We introduce a variable a_i for every player p_i denoting p_i ’s chosen *strategy*. Namely, $a_i = 1$ describes that player p_i is protected whereas for a player p_j willing to take the risk, $a_j = 0$. In the following, we will assume that $a_j \in \{0, 1\}$, that is, we do not allow players to *mix* (i.e., use probabilistic distributions over) their strategies. These choices are summarized by the *strategy profile*, the vector $\vec{a} = (a_1, \dots, a_n)$. After the players have made their decisions, a virus spreads in the

¹E.g., the Outlook worm *Worm.ExploreZip*.

²See <http://news.softpedia.com/news/Skype-Attacked-By-Fast-Spreading-Virus-52039.shtml>.

network. The propagation model is as follows: first, one node p of the network is chosen uniformly at random as a starting point. If this node is inoculated, there is no damage and the process terminates. Otherwise, the virus infects p and all unprotected neighbors of p . The virus now propagates recursively to their unprotected neighbors. Hence, the more insecure players are connected, the more likely they are to be infected. The vulnerable region (set of nodes) in which an insecure player p_i lies is referred to as p_i 's *attack component*.

We only consider a limited region of the parameter space to avoid trivial cases. If the cost C is too large, no player will inoculate, resulting in a totally insecure network and therefore all nodes eventually will be infected. On the other hand, if $C \ll L$, the best strategy for all players is to inoculate. Thus, we will assume that $C \leq L$ and $C > L/n$ in the following.

In our game, a player has the following expected cost:

DEFINITION 3.1 (ACTUAL INDIVIDUAL COST).
The actual individual cost of a player p_i is defined as

$$c_a(i, \vec{a}) = a_i \cdot C + (1 - a_i)L \cdot \frac{k_i}{n}$$

where k_i denotes the size of p_i 's attack component. If p_i is inoculated, k_i stands for the size of the attack component that would result if p_i became insecure. In the following, let $c_a^0(i, \vec{a})$ refer to the actual cost of an insecure and $c_a^1(i, \vec{a})$ to the actual cost of a secure node p_i .

The total *social cost* of a game is defined as the sum of the cost of all participants: $C_a(\vec{a}) = \sum_{p_i \in V} c_a(i, \vec{a})$.

Classic game theory assumes that all players act selfishly, i.e., each player seeks to minimize its individual cost. In order to study the impact of such selfish behavior, the solution concept of a *Nash equilibrium* (NE) is used. A Nash equilibrium is a strategy profile where no selfish player can unilaterally reduce its individual cost given the strategy choices of the other players. We can think of Nash equilibria as the stable strategy profiles of games with selfish players. As stated earlier already, we consider only pure Nash equilibria in this paper, i.e., players cannot use random distributions over their strategies but are bound to decide whether they want to inoculate or not.

In a pure Nash equilibrium, it must hold for each player p_i that given a strategy profile $\vec{a} \forall p_i \in V, \forall a'_i \neq a_i : c_a(i, \vec{a}) \leq c_a(i, (a_1, \dots, a'_i, \dots, a_n))$, implying that player p_i cannot decrease its cost by choosing an alternative strategy a'_i . In order to quantify the performance loss due to selfishness, the (not necessarily unique) Nash equilibria are compared to the optimum situation where all players collaborate. To this end we consider the *Price of Anarchy* (PoA), i.e., the ratio of the social cost of the worst Nash equilibrium divided by the optimal social cost for a problem instance I . More formally, $PoA(I) = \max_{NE} C_{NE}(I) / C_{OPT}(I)$.

3.2 Social Networks

We now introduce our model for social networks. We define a *Friendship Factor* F which captures the extent to which players care about their *friends*, i.e., about the players *adjacent* to them in the social network. More formally, F is the factor by which a player p_i takes the individual cost of its neighbors into account when deciding for a strategy. F can assume any value between 0 and 1. $F = 0$ implies

that the players do not consider their neighbors' cost at all, whereas $F = 1$ implies that a player values the well-being of its neighbors to the same extent as its own. Let $\Gamma(p_i)$ denote the set of neighbors of a player p_i . Moreover, let $\Gamma_{sec}(p_i) \subseteq \Gamma(p_i)$ be the set of inoculated neighbors, and $\Gamma_{\overline{sec}}(p_i) = \Gamma(p_i) \setminus \Gamma_{sec}(p_i)$ the remaining insecure neighbors.

We distinguish between a player's *actual cost* and a player's *perceived cost*. A player's actual individual cost is the expected cost arising for each player defined in Definition 3.1 used to compute a game's social cost. In our social network, the decisions of our players are steered by the players' *perceived cost*.

DEFINITION 3.2 (PERCEIVED INDIVIDUAL COST).
The perceived individual cost of a player p_i is defined as

$$c_p(i, \vec{a}) = c_a(i, \vec{a}) + F \cdot \sum_{p_j \in \Gamma(p_i)} c_a(j, \vec{a}).$$

In the following, we write $c_p^0(i, \vec{a})$ to denote the perceived cost of an insecure node p_i and $c_p^1(i, \vec{a})$ for the perceived cost of an inoculated node.

This definition entails a new notion of equilibrium. We define a *friendship Nash equilibrium* (FNE) as a strategy profile \vec{a} where no player can reduce its *perceived cost* by unilaterally changing its strategy given the strategies of the other players. Formally, $\forall p_i \in V, \forall a'_i \neq a_i : c_p(i, \vec{a}) \leq c_p(i, (a_1, \dots, a'_i, \dots, a_n))$. Given this equilibrium concept, we define the *Windfall of Friendship* Υ .

DEFINITION 3.3 (WINDFALL OF FRIENDSHIP (WOF)).
The *Windfall of Friendship* $\Upsilon(F, I)$ is the ratio of the social cost of the worst Nash equilibrium for I and the social cost of the worst friendship Nash equilibrium for I :

$$\Upsilon(F, I) = \frac{\max_{NE} C_{NE}(I)}{\max_{FNE} C_{FNE}(F, I)}$$

$\Upsilon(F, I) > 1$ implies the existence of a real windfall in the system, whereas $\Upsilon(F, I) < 1$ denotes that the social cost can become *greater* in social graphs than in purely selfish environments.

4. GENERAL ANALYSIS

In this section we characterize friendship Nash equilibria and derive general results on the Windfall of Friendship for the virus propagation game in social networks. It has been shown in [1] that in classic Nash equilibria ($F = 0$), an attack component can never consist of more than Cn/L insecure nodes. A similar characteristic also holds for friendship Nash equilibria. As every player cares about its neighbors, the maximal attack component size in which an insecure player p_i still does not inoculate depends on the number of p_i 's insecure neighbors and the size of their attack components. Therefore, it differs from player to player. We have the following helper lemma.

LEMMA 4.1. *The player p_i will inoculate if and only if the size of its attack component is*

$$k_i > \frac{Cn/L + F \cdot \sum_{p_j \in \Gamma_{\overline{sec}}(p_i)} k_j}{1 + F|\Gamma_{\overline{sec}}(p_i)|},$$

where the k_j s are the attack component sizes of p_i 's insecure neighbors assuming p_i is secure.

PROOF. Player p_i will inoculate if and only if this choice lowers the perceived cost. By Definition 3.2, the perceived individual cost of an inoculated node are

$$c_p^1(i, \vec{a}) = C + F \left(|\Gamma_{sec}(p_i)|C + \sum_{p_j \in \Gamma_{\overline{sec}}(p_i)} L \frac{k_j}{n} \right)$$

and for an insecure node we have

$$c_p^0(i, \vec{a}) = L \frac{k_i}{n} + F \left(|\Gamma_{sec}(p_i)|C + |\Gamma_{\overline{sec}}(p_i)|L \frac{k_i}{n} \right).$$

For p_i to prefer to inoculate it must hold that

$$\begin{aligned} c_p^0(i, \vec{a}) &> c_p^1(i, \vec{a}) \Leftrightarrow \\ L \frac{k_i}{n} + F \cdot |\Gamma_{\overline{sec}}(p_i)|L \frac{k_i}{n} &> C + F \cdot \sum_{p_j \in \Gamma_{\overline{sec}}(p_i)} L \frac{k_j}{n} \Leftrightarrow \\ L \frac{k_i}{n} (1 + F|\Gamma_{\overline{sec}}(p_i)|) &> C + \frac{FL}{n} \cdot \sum_{p_j \in \Gamma_{\overline{sec}}(p_i)} k_j \Leftrightarrow \\ k_i (1 + F|\Gamma_{\overline{sec}}(p_i)|) &> Cn/L + F \cdot \sum_{p_j \in \Gamma_{\overline{sec}}(p_i)} k_j \Leftrightarrow \\ k_i &> \frac{Cn/L + F \cdot \sum_{p_j \in \Gamma_{\overline{sec}}(p_i)} k_j}{1 + F|\Gamma_{\overline{sec}}(p_i)|}. \end{aligned}$$

□

A pivotal question is of course whether social networks where players care about their friends yield better equilibria than selfish environments. The following theorem answers this questions affirmatively: the worst FNE costs never more than the worst NE.

THEOREM 4.2. *For all instances of the virus inoculation game and $0 \leq F \leq 1$, it holds that*

$$1 \leq \Upsilon(F, I) \leq PoA(I).$$

PROOF. The proof idea for $\Upsilon(F, I) \geq 1$ is the following: for an instance I we consider an arbitrary FNE with $F > 0$. Given this equilibrium, we show the existence of a NE with larger social cost. Let α be any (e.g., the worst) FNE in the social model. If α is also a NE in the same instance with $F = 0$ then we are done. Otherwise there is at least one player p_i that prefers to change its strategy. Assume p_i is insecure but favors inoculation. Therefore p_i 's attack component has on the one hand to be of size at least $k_i' > Cn/L$ [1] and on the other hand of size at most $k_i'' = (Cn/L + F \cdot \sum_{p_j \in \Gamma_{\overline{sec}}(p_i)} k_j) / (1 + F|\Gamma_{\overline{sec}}(p_i)|) \leq (Cn/L + F|\Gamma_{\overline{sec}}(p_i)|(k_i'' - 1)) / (1 + F|\Gamma_{\overline{sec}}(p_i)|) \Leftrightarrow k_i'' \leq Cn/L - F|\Gamma_{\overline{sec}}(p_i)|$ (cf Lemma 4.1). This is impossible and yields a contradiction to the assumption that in the selfish network, an additional player wants to inoculate.

It remains to study the case where p_i is secure in the FNE but prefers to be insecure in the NE. Observe that, since every player has the same preference on the attack component's size when $F = 0$, a newly insecure player cannot trigger other players to inoculate. Furthermore, only the players inside p_i 's attack component are affected by this change. The total cost of this attack component increases

by at least

$$\begin{aligned} x &= \underbrace{\frac{k_i}{n}L - C}_{p_i} + \underbrace{\sum_{p_j \in \Gamma_{\overline{sec}}(p_i)} \left(\frac{k_i}{n}L - \frac{k_j}{n}L \right)}_{p_i \text{'s insecure neighbors}} \\ &= \frac{k_i}{n}L - C + \frac{L}{n} (|\Gamma_{\overline{sec}}(p_i)|k_i - \sum_{p_j \in \Gamma_{\overline{sec}}(p_i)} k_j). \end{aligned}$$

Applying Lemma 4.1 guarantees that

$$\sum_{p_j \in \Gamma_{\overline{sec}}(p_i)} k_j \leq \frac{k_i(1 + F|\Gamma_{\overline{sec}}(p_i)|) - Cn/L}{F}.$$

This results in

$$\begin{aligned} x &\geq \frac{L}{n} \left(|\Gamma_{\overline{sec}}(p_i)|k_i - \frac{k_i(1 + F|\Gamma_{\overline{sec}}(p_i)|) - Cn/L}{F} \right) \\ &= \frac{k_iL}{n} \left(1 - \frac{1}{F} \right) - C \left(1 - \frac{1}{F} \right) > 0, \end{aligned}$$

since a player only gives up its protection if $C > \frac{k_iL}{n}$. If more players are unhappy with their situation and become vulnerable, the cost for the NE increases further. In conclusion, there exists a NE for every FNE with $F \geq 0$ for the same instance which is at least as expensive.

The upper bound for the WoF, i.e., $PoA(I) \geq \Upsilon(F, I)$, follows directly from the definitions: while the PoA is the ratio of the NE's social cost divided by the social optimum, $\Upsilon(F, I)$ is the ratio between the cost of the NE and the FNE. As the FNE's cost must be at least as large as the social optimum cost the claim follows. □

REMARK 4.3. *Note that Aspnes et al. [1] proved that the Price of Anarchy never exceeds the size of the network, i.e., $n \geq PoA(I)$. Consequently, the Windfall of Friendship cannot be larger than n due to Theorem 4.2.*

The above result leads to the question of whether the Windfall of Friendship grows monotonically with stronger social ties, i.e., with larger friendship factors F . Intriguingly, this is not the case.

THEOREM 4.4. *For all network with more than seven nodes, there exist game instances where $\Upsilon(F, I)$ does not grow monotonically in F .*

PROOF. We give a counter example for the star graph S_n which has one center node and $n - 1$ leaf nodes. Consider two friendship factors, F_l and F_s where $F_l > F_s$. We show that for the large friendship factor F_l , there exists a FNE, FNE_1 , where only the center node and one leaf node remain insecure. For the same setting but with a small friendship factor F_s , at least two leaf nodes will remain insecure, which will trigger the center node to inoculate, yielding a FNE, FNE_2 , where only the center node is secure.

Consider FNE_1 first. Let c be the insecure center node, let l_1 be the insecure leaf node, and let l_2 be a secure leaf node. In order for FNE_1 to constitute a Nash equilibrium, the following conditions must hold:

$$\text{node } c : \frac{2L}{n} + \frac{2F_l L}{n} < C + \frac{F_l L}{n}$$

$$\text{node } l_1 : \frac{2L}{n} + \frac{2F_l L}{n} < C + \frac{F_l L}{n}$$

$$\text{node } l_2 : C + \frac{2F_l L}{n} < \frac{3L}{n} + \frac{3F_l L}{n}$$

For FNE_2 , let c be the insecure center node, let l_1 be one of the two insecure leaf nodes, and let l_2 be a secure leaf node. In order for the leaf nodes to be happy with their situation but for the center node to prefer to inoculate, it must hold that:

$$\text{node } c : C + \frac{2F_s L}{n} < \frac{3L}{n} + \frac{6F_s L}{n}$$

$$\text{node } l_1 : \frac{3L}{n} + \frac{3F_s L}{n} < C + \frac{2F_s L}{n}$$

$$\text{node } l_2 : C + \frac{3F_s L}{n} < \frac{4L}{n} + \frac{4F_s L}{n}$$

Now choose $C := 5L/(2n) + F_l L/n$. This yields the following conditions: $F_l > F_s + 1/2$, $F_l < F_s + 3/2$, and $F_l < 4F_s + 1/2$. These conditions are easily fulfilled, e.g., with $F_l = 3/4$ and $F_s = 1/8$. Observe that the social cost of the first FNE (for F_l) is $Cost(S_n, \vec{a}_{FNE1}) = (n-2)C + 4L/n$, whereas for the second FNE (for F_s) $Cost(S_n, \vec{a}_{FNE2}) = C + (n-1)L/n$. Thus, $Cost(S_n, \vec{a}_{FNE1}) - Cost(S_n, \vec{a}_{FNE2}) = (n-3)C - (n+3)L/n > 0$ as we have chosen $C > 5L/(2n)$ and as, due to our assumption, $n > 7$. This concludes the proof. \square

Reasoning about best and worst Nash equilibria raises the question of how difficult it is to compute such equilibria. We can generalize the proof given in [1] and show that computing the most economical and the most expensive FNE is hard for any friendship factor.

THEOREM 4.5. *Computing the best and the worst pure FNE is \mathcal{NP} -complete for any $F \in [0, 1]$.*

PROOF. We prove this theorem by a reduction from two \mathcal{NP} -hard problems, VERTEX COVER [13] and INDEPENDENT DOMINATING SET [11]. Concretely, for the decision version of the problem, we show that answering the question whether there exists a FNE costing less than k , or more than k respectively, is at least as hard as solving vertex cover or independent dominating set. Note that verifying whether a proposed solution is correct can be done in polynomial time, hence the problems are indeed in \mathcal{NP} .

Fix some graph $G = (V, E)$ and set $C = 1$ and $L = n/1.5$. We show first that the following two conditions are necessary and sufficient for a FNE: (a) all neighbors of an insecure node are secure, and (b) every inoculated node has at least one insecure neighbor. Due to our assumption that $C > L/n$, condition (b) is satisfied in all FNE. To see that condition (a) holds as well, assume the contrary, i.e., an attack component of size at least two. An insecure node p_i in this attack component bears the cost $\frac{k_i}{n}L + F(|\Gamma_{sec}(p_i)|C + |\Gamma_{\overline{sec}}(p_i)|\frac{k_i}{n}L)$. Changing its strategy reduces its cost by at least $\Delta_i = \frac{k_i}{n}L + F|\Gamma_{\overline{sec}}(p_i)|\frac{k_i}{n}L - C - F|\Gamma_{sec}(p_i)|\frac{k_i-1}{n}L = \frac{k_i}{n}L + F|\Gamma_{\overline{sec}}(p_i)|\frac{1}{n}L - C$. By our assumption that $k_i \geq 2$, and hence $|\Gamma_{\overline{sec}}(p_i)| \geq 1$, it holds that $\Delta_i > 0$, resulting in p_i becoming secure. Hence, condition (a) holds in all FNE as well. For the opposite direction assume that an insecure node wants to change its strategy even though (a) and (b) are true. This is impossible because in this case (b) would be violated. An inoculated node would destroy (a) by adopting another strategy. Thus (a) and (b) are sufficient for a FNE.

We now argue that G has a vertex cover of size k if and only if the virus game has a FNE with k or fewer secure nodes, or equivalently an equilibrium with social cost at most $Ck + (n-k)L/n$, as each insecure node must be in a component of size 1 and contributes exactly L/n expected cost. Given a minimal vertex cover $V' \subseteq V$, observe that installing the software on all nodes in V' satisfies condition (a) because V' is a vertex cover and (b) because V' is minimal. Conversely, if V' is the set of secure nodes in a FNE, then V' is a vertex cover by condition (a) which is minimal by condition (b).

For the worst FNE, we consider an instance of the independent dominating set problem. Given an independent dominating set V' , installing the software on all nodes except the nodes in V' satisfies condition (a) because V' is independent and (b) because V' is a dominating set. Conversely, the insecure nodes in any FNE are independent by condition (a) and dominating by condition (b). This shows that G has an independent dominating set of size at most k if and only if it has a FNE with at least $n-k$ secure nodes. \square

5. WINDFALL FOR SPECIAL GRAPHS

While the last section has presented general results on equilibria in social networks and the Windfall of Friendship, we now present upper and lower bounds on the Windfall of Friendship for concrete topologies, namely the *complete graph* K_n and the *star graph* S_n .

5.1 Complete Graphs

In order to initiate the study of the Windfall of Friendship, we consider a very simple topology, the complete graph K_n where all players are connected to each other. First consider the classic setting where nodes do not care about their neighbors ($F = 0$). We have the following result:

LEMMA 5.1. *In the graph K_n , there are two Nash equilibria with total cost*

$$NE_1: Cost(K_n, \vec{a}_{NE1}) = C(n - \lfloor Cn/L \rfloor + 1) + L/n(\lfloor Cn/L \rfloor - 1)^2,$$

and

$$NE_2: Cost(K_n, \vec{a}_{NE2}) = C(n - \lfloor Cn/L \rfloor) + L/n(\lfloor Cn/L \rfloor)^2.$$

If $\lfloor Cn/L \rfloor - 1 = \lfloor Cn/L \rfloor$, there is only one Nash equilibrium.

PROOF. Let \vec{a} be a NE. Consider an inoculated node p_i and an insecure node p_j , and hence $c_a(i, \vec{a}) = C$ and $c_a(j, \vec{a}) = L\frac{k_j}{n}$, where k_j is the total number of insecure nodes in K_n . In order for p_i to remain inoculated, it must hold that $C \leq (k_j + 1)L/n$, so $k_j \geq \lfloor Cn/L - 1 \rfloor$; for p_j to remain insecure, it holds that $k_j L/n \leq C$, so $k_j \leq \lfloor Cn/L \rfloor$. As the total social cost in K_n is given by $Cost(K_n, \vec{a}) = (n - k_j)C + k_j^2 L/n$, the claim follows. \square

Observe that the equilibrium size of the attack component is roughly twice the size of the attack component of the social optimum, as $Cost(K_n, \vec{a}) = (n - k_j)C + k_j^2 L/n$ is minimized for $k_j = Cn/2L$.

LEMMA 5.2. *In the social optimum for K_n , the size of the attack component is either $\lfloor \frac{1}{2}Cn/L \rfloor$ or $\lceil \frac{1}{2}Cn/L \rceil$, yielding*

a total social cost of

$$\text{Cost}(K_n, \vec{a}_{OPT}) = (n - \lfloor \frac{1}{2}Cn/L \rfloor)C + (\lfloor \frac{1}{2}Cn/L \rfloor)^2 \frac{L}{n}$$

or

$$\text{Cost}(K_n, \vec{a}_{OPT}) = (n - \lceil \frac{1}{2}Cn/L \rceil)C + (\lceil \frac{1}{2}Cn/L \rceil)^2 \frac{L}{n}.$$

In order to compute the Windfall of Friendship, the friendship Nash equilibria in social networks have to be identified.

LEMMA 5.3. *In K_n , there are two friendship Nash equilibria with total cost*

$$\begin{aligned} \text{FNE}_1: \text{Cost}(K_n, \vec{a}_{FNE1}) &= C \left(n - \left\lceil \frac{Cn/L - 1}{1 + F} \right\rceil \right) \\ &\quad + L/n \left(\left\lceil \frac{Cn/L - 1}{1 + F} \right\rceil \right)^2, \end{aligned}$$

and

$$\begin{aligned} \text{FNE}_2: \text{Cost}(K_n, \vec{a}_{FNE2}) &= C \left(n - \left\lfloor \frac{Cn/L + F}{1 + F} \right\rfloor \right) \\ &\quad + L/n \left(\left\lfloor \frac{Cn/L + F}{1 + F} \right\rfloor \right)^2. \end{aligned}$$

If $\lceil (Cn/L - 1)/(1 + F) \rceil = \lfloor (Cn/L + F)/(1 + F) \rfloor$, there is only one FNE.

PROOF. According to Lemma 4.1, in a FNE, a node p_i remains secure if otherwise the component had size at least $k_i = k_j + 1 \geq (Cn/L + Fk_j^2)/(1 + Fk_j)$ where k_j is the number of insecure nodes. This implies that $k_j \geq \lceil (Cn/L - 1)/(1 + F) \rceil$. Dually, for an insecure node p_j it holds that $k_j \leq (Cn/L + F(k_j - 1)^2)/(1 + F(k_j - 1))$ and therefore $k_j \leq \lfloor (Cn/L + F)/(1 + F) \rfloor$. Given these bounds on the total number of insecure nodes in a FNE, the social cost can be obtained by substituting k_j in $\text{Cost}(K_n, \vec{a}) = (n - k_j)C + k_j^2 L/n$. As the difference between the upper and the lower bound for k_j is at most 1, there are at most two equilibria and the claim follows. \square

Given the characteristics of the different equilibria, we have the following theorem.

THEOREM 5.4. *In K_n , the Windfall of Friendship is at most $\Upsilon(F, I) = 4/3$ for an arbitrary network size. This is tight in the sense that there are indeed instances where the worst FNE is a factor 4/3 better than the worst NE.*

PROOF. *Upper Bound.* We first derive the upper bound on $\Upsilon(F, I)$.

$$\begin{aligned} \Upsilon(F, I) &= \frac{\text{Cost}(K_n, \vec{a}_{NE})}{\text{Cost}(K_n, \vec{a}_{FNE})} \\ &\leq \frac{\text{Cost}(K_n, \vec{a}_{NE})}{\text{Cost}(K_n, \vec{a}_{OPT})} \\ &\leq \frac{(n - \lceil Cn/L - 1 \rceil)C + (\lceil Cn/L \rceil)^2 \frac{L}{n}}{(n - \frac{1}{2}Cn/L)C + (\frac{1}{2}Cn/L)^2 \frac{L}{n}} \end{aligned}$$

as the optimal social cost (cf Lemma 5.2) is smaller or equal to the social cost of any FNE. Simplifying this expression yields

$$\Upsilon(F, I) \leq \frac{n(1 - C/L)C + C^2 n/L}{n(1 - \frac{1}{2}C/L)C + \frac{1}{4}C^2 n/L} = \frac{1}{1 - \frac{1}{4}C/L}.$$

This term is maximized for $L = C$, implying that $\Upsilon(F, I) \leq 4/3$, for arbitrary n .

Lower Bound. We now show that the ratio between the equilibria cost reaches 4/3.

There exists exactly one social optimum of cost $Ln/2 + (n/2)^2 L/n = 3nL/4$ for even n and $C = L$ by Lemma 5.2. For $F = 1$ this is also the only friendship Nash equilibrium due to Lemma 5.3. In the selfish game however the Nash equilibrium has fewer inoculated nodes and is of cost nL (see Lemma 5.1). Since these are the only Nash equilibria they constitute the worst equilibria and the ratio is

$$\Upsilon(F, I) = \frac{\text{Cost}(K_n, \vec{a}_{NE})}{\text{Cost}(K_n, \vec{a}_{FNE})} = \frac{nL}{3/4nL} = 4/3.$$

\square

To conclude our analysis of K_n , observe that friendship Nash equilibria always exist in complete graphs, and that in environments where one node at a time is given the chance to change its strategy in a best response manner quickly results in such an equilibrium.

5.2 Star

While the analysis of K_n was simple, it turns out that already slightly more sophisticated graphs are challenging. In the following, we investigate the Windfall of Friendship in star graphs S_n . Note that in S_n , the social welfare is maximized if the center node inoculates and all other nodes do not. The total inoculation cost then is C and the attack components are all of size 1, yielding a total social cost of $\text{Cost}(S_n, \vec{a}_{OPT}) = C + (n - 1)L/n$.

LEMMA 5.5. *In the social optimum of the star graph S_n , only the center node is inoculated. The social cost is*

$$\text{Cost}(S_n, \vec{a}_{OPT}) = C + (n - 1)L/n.$$

The situation where only the center node is inoculated also constitutes a NE. However, there are more Nash equilibria.

LEMMA 5.6. *In the star graph S_n , there are at most three Nash equilibria with total cost*

$$\begin{aligned} \text{NE}_1: \text{Cost}(S_n, \vec{a}_{NE1}) &= C + (n - 1)L/n, \\ \text{NE}_2: \text{Cost}(S_n, \vec{a}_{NE2}) &= C(n - \lceil Cn/L \rceil + 1) \\ &\quad + L/n(\lceil Cn/L \rceil - 1)^2, \end{aligned}$$

and

$$\text{NE}_3: \text{Cost}(S_n, \vec{a}_{NE3}) = C(n - \lfloor Cn/L \rfloor) + L/n(\lfloor Cn/L \rfloor)^2.$$

If $Cn/L \notin \mathbb{N}$, only two equilibria exist.

PROOF. If the center node is the only secure node, changing its strategy costs L but saves only C . When a leaf node becomes secure, its cost changes from L/n to C . These changes are unprofitable, and the social cost of this NE is $\text{Cost}(S_n, \vec{a}_{NE1}) = C + (n - 1)L/n$.

For the other Nash equilibria the center node is not inoculated. Let the number of insecure leaf nodes be n_0 . In order for a secure node to remain secure, it must hold that $C \leq (n_0 + 2)L/n$, and hence $n_0 \geq \lceil Cn/L - 2 \rceil$. For an insecure node to remain insecure, it must hold that $(1 + n_0)L/n \leq C$, thus $n_0 \leq \lfloor Cn/L - 1 \rfloor$. Therefore, we can conclude that there are at most two Nash equilibria, one with $\lceil Cn/L - 1 \rceil$ and one with $\lfloor Cn/L \rfloor$ many insecure nodes. The total social cost follows by substituting n_0 in the total social cost function. Finally, observe that if $Cn/L \in \mathbb{N}$ and $Cn/L > 3$, all three equilibria exist in parallel. \square

Let us consider the social network scenario again.

LEMMA 5.7. *In S_n , there are at most three friendship Nash equilibria with total cost*

$$\begin{aligned} FNE_1: \text{Cost}(S_n, \vec{a}_{FNE1}) &= C + (n-1)L/n, \\ FNE_2: \text{Cost}(S_n, \vec{a}_{FNE2}) &= C(n - \lceil Cn/L - F \rceil + 1) \\ &\quad + L/n(\lceil Cn/L - F \rceil - 1)^2, \end{aligned}$$

and

$$FNE_3: \text{Cost}(S_n, \vec{a}_{FNE3}) = C(n - \lfloor Cn/L - F \rfloor) + L/n(\lfloor Cn/L - F \rfloor)^2.$$

If $Cn/L - F \notin \mathbb{N}$, at most 2 friendship Nash equilibria exist.

PROOF. First, observe that having only an inoculated center node constitutes a FNE. In order for the center node to remain inoculated, it must hold that $C + F(n-1)L\frac{1}{n} \leq nL/n + F(n-1)L\frac{n}{n} = L + F(n-1)L$. All leaf nodes remain insecure as long as $L/n + FC \leq C + FC \Leftrightarrow L/n \leq C$. These conditions are always true, and we have $\text{Cost}(S_n, \vec{a}_{FNE1}) = C + (n-1)L/n$. If the center node is not inoculated, we have n_0 insecure and $n - n_0 - 1$ inoculated leaf nodes. In order for a secure leaf node to remain secure, it is necessary that $C + F\frac{n_0+1}{n}L \leq \frac{n_0+2}{n}L + F\frac{n_0+2}{n}L$, so $n_0 \geq \lceil Cn/L - F \rceil - 2$. For an insecure leaf node, it must hold that $\frac{n_0+1}{n}L + F\frac{n_0+1}{n}L \leq C + F\frac{n_0}{n}L$, so $n_0 \leq \lfloor Cn/L - F \rfloor - 1$. The claim follows by substitution. \square

Note that there are instances where FNE_1 is the only friendship Nash equilibrium. We already made use of this phenomenon in Section 4 to show that $\Upsilon(F, I)$ is not monotonically increasing in F . The next lemma states under which circumstances this is the case.

LEMMA 5.8. *In S_n , there is a unique FNE equivalent to the social optimum if and only if*

$$\lceil Cn/L - F \rceil - \lfloor \frac{1}{2F}(\sqrt{1 - 4F(1 - Cn/L)} - 1) \rfloor - 2 \geq 0$$

PROOF. S_n has only one FNE if every (insecure) leaf node is content with its chosen strategy but the insecure center node would rather inoculate. In order for an insecure leaf node to remain insecure we have $n_0 \leq \lfloor Cn/L - 1 - F \rfloor$ and the insecure center node wants to inoculate if and only if

$$\begin{aligned} &C + F(n - n_0 - 1)C + Fn_0\frac{1}{n}L \\ &< (n_0 + 1)\frac{L}{n} + F(n - n_0 - 1)C + Fn_0\frac{n_0 + 1}{n}L, \end{aligned}$$

which is equivalent to $Fn_0^2 + n_0 + 1 - Cn/L > 0$. This implies that $n_0 \geq \lfloor \frac{1}{2F}(\sqrt{1 - 4F(1 - Cn/L)} - 1) + 1 \rfloor$. Therefore there is only one FNE if and only if there exists an integer n_0 such that $\lfloor \frac{1}{2F}(\sqrt{1 - 4F(1 - Cn/L)} - 1) + 1 \rfloor \leq n_0 \leq \lfloor Cn/L - 1 - F \rfloor$. \square

Given the characterization of the various equilibria, the Windfall of Friendship can be computed.

THEOREM 5.9. *If $\lfloor \frac{1}{2F}(\sqrt{1 - 4F(1 - Cn/L)} - 1) \rfloor + 2 - \lceil Cn/L - F \rceil \leq 0$, the Windfall of Friendship is*

$$\Upsilon(F, I) \geq \frac{(n-2)C + L/n}{C + (n-1)L/n}, \quad \text{else} \quad \Upsilon(F, I) \leq \frac{n+1}{n-3}.$$

PROOF. According to Lemma 5.8, the friendship Nash equilibrium is unique and hence equivalent to the social optimum if $\lfloor Cn/L - F \rceil - \lfloor \frac{1}{2F}(\sqrt{1 - 4F(1 - Cn/L)} - 1) \rfloor - 2 \geq 0$. On the other hand, observe that there always exist expensive Nash equilibria where the center node is not inoculated. Hence, we have

$$\begin{aligned} \Upsilon(F, I) &= \frac{\text{Cost}(S_n, \vec{a}_{NE})}{\text{Cost}(S_n, \vec{a}_{FNE})} = \frac{\text{Cost}(S_n, \vec{a}_{NE})}{\text{Cost}(S_n, \vec{a}_{OPT})} \\ &\geq \frac{(n - \lfloor Cn/L - 1 \rfloor)C + (\lceil Cn/L \rceil - 1)^2 L/n}{C + (n-1)L/n} \\ &\geq \frac{C(n-2) + L/n}{C + (n-1)L/n}. \end{aligned}$$

Otherwise, i.e., if there exist friendship Nash equilibria with an insecure center node, an upper bound for the WoF can be computed

$$\begin{aligned} \Upsilon(F, I) &= \frac{\text{Cost}(S_n, \vec{a}_{NE})}{\text{Cost}(S_n, \vec{a}_{FNE})} \\ &\leq \frac{(n - \lceil Cn/L - 1 \rceil)C + (\lfloor Cn/L \rfloor)^2 L/n}{(n - \lfloor Cn/L - F \rfloor)C + (\lceil Cn/L - 1 - F \rceil)^2 L/n} \\ &\leq \frac{(n+1)C}{nC + FC - 2C(1+F) + (1+F)^2 L/n} \\ &< \frac{(n+1)C}{C(n+F-2(1+F))} < \frac{n+1}{n-3}. \end{aligned}$$

\square

Theorem 5.9 reveals that caring about the cost incurred by friends is particularly helpful to reach more desirable equilibria. In large star networks, the social welfare can be much higher than in Nash equilibria: in particular, the Windfall of Friendship can increase linearly in n , and hence indeed be asymptotically as large as the Price of Anarchy. However, if $\lfloor Cn/L - F \rceil - \lfloor \frac{1}{2F}(\sqrt{1 - 4F(1 - Cn/L)} - 1) \rfloor - 2 \geq 0$ does not hold, social networks are not much better than purely selfish systems: the maximal gain is constant.

Finally observe that in stars friendship Nash equilibria always exist and can be computed efficiently (in linear time) by any best response strategy.

5.3 Discussion

We consider the results derived in this section as a first step. Besides stars and complete graphs, there are many other interesting graph classes, e.g., Kleinberg graphs featuring the small-world property. In this section, we have shown that even simple graphs such as the star graph all possible values of the Windfall of Friendship can be reached, from constant ratios up to ratios linear in n , which is asymptotically maximal for general graphs as well since the Price of Anarchy is bounded by n [1].

6. CONCLUSION

This paper has studied a virus inoculation game on social networks. We have presented a framework which allows us to quantify the effects of caring for direct neighbors on the social network graph. We believe that our work opens many exciting possibilities for future research. First, the analysis of the Windfall of Friendship needs to be continued for other graphs, e.g., random graphs with small world properties. Moreover, the question of the existence of friendship Nash equilibria needs to be addressed on general graphs. Observe

that while there are always Nash equilibria when players do not care about each other, in the social setting, two players in the same attack component may not have the same preferences on inoculation: one player can be discontent in the same component where another player is happy with its strategy. Participants in social networks are often involved in very different kind of interactions, and hence, many other games can be analyzed on these networks with our framework. In this paper the players only care for their direct neighbors. Investigating the effects of considering players over multiple hops (maybe to a lesser extent) or letting the metric distance between two players decide about their degree of friendship, i.e. porting the Virus Inoculation Game into the Euclidean space, are interesting research directions. The virus game itself can also be extended of course. For example, a virus might not infect unlimitedly many insecure players that are adjacent to an already infected player. Or there might be different kinds of viruses and different kinds of nodes, which are resistant to some viruses but not to others. Furthermore, one could also modify the underlying model by introducing malicious players in this social adaption as well.

7. REFERENCES

- [1] J. Aspnes, K. Chang, and A. Yampolskiy. Inoculation Strategies for Victims of Viruses and the Sum-of-Squares Partition Problem. In *Proc. 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2005.
- [2] M. Babaioff, R. Kleinberg, and C. H. Papadimitriou. Congestion Games with Malicious Players. In *Proc. 18th ACM Conference on Electronic Commerce (EC)*, 2007.
- [3] N. T. Bailey. *The Mathematical Theory of Infectious Diseases and its Applications*. Hafner Press, 1975.
- [4] R. Blundell, W. Newey, and T. Persson (Eds). *Advances in Economics and Econometrics (Chapter 1: The Economics of Social Networks)*. 2006.
- [5] A. Fabrikant, A. Luthra, E. Maneva, C. H. Papadimitriou, and S. Shenker. On a Network Creation Game. In *Proc. 22nd Annual Symposium on Principles of Distributed Computing (PODC)*, 2003.
- [6] J. Feigenbaum, C. H. Papadimitriou, and S. Shenker. Sharing the Cost of Multicast Transmissions. *Journal of Computer and System Sciences*, 63(1):21–41, 2001.
- [7] G. W. Flake, S. Lawrence, C. L. Giles, and F. M. Coetzee. Self-Organization and Identification of Web Communities. *Computer*, 35(3):66–71, 2002.
- [8] C. Fleizach, M. Liljenstam, P. Johansson, G. M. Voelker, and A. Mehes. Can You Infect Me Now? Malware Propagation in Mobile Phone Networks. In *Proc. 2007 ACM Workshop on Recurring Malcode (WORM)*, 2007.
- [9] P. Fraigniaud, C. Gavoille, and C. Paul. Eclecticism Shrinks Even Small Worlds. In *Proc. 23rd Annual ACM Symposium on Principles of Distributed Computing (PODC)*, 2004.
- [10] J. C. Frauenthal. *Mathematical Modeling in Epidemiology*. Springer-Verlag, 1980.
- [11] M. R. Garey and D. S. Johnson. *Computers and Intractability : A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
- [12] S. M. Kakade, M. Kearns, and L. E. Ortiz. Graphical Economics. In *Proc. 17th Annual Conference on Learning Theory (COLT)*, 2004.
- [13] R. M. Karp. Reducibility Among Combinatorial Problems. *Complexity of Computer Computations*, pages 85–103, 1972.
- [14] M. Kearns, M. Littman, and S. Singh. Graphical Models for Game Theory. In *Proc. Conference on Uncertainty in Artificial Intelligence (UAI)*, 2001.
- [15] J. O. Kephart and S. R. White. Measuring and Modeling Computer Virus Prevalence. In *Proc. IEEE Symposium on Security and Privacy*, 1993.
- [16] J. O. Kephart, S. R. White, and D. M. Chess. Computers and Epidemiology. *IEEE Spectrum*, 30(5):20–26, 1993.
- [17] J. Kleinberg. The Small-World Phenomenon: An Algorithmic Perspective. In *Proc. 32nd ACM Symposium on Theory of Computing (STOC)*, 2000.
- [18] J. Kleinberg. Complex Networks and Decentralized Search Algorithms. In *Proc. International Congress of Mathematicians (ICM)*, 2006.
- [19] J. Kleinberg. *Algorithmic Game Theory. Chapter 24 : Cascading Behavior in Networks: Algorithmic and Economic Issues*. N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani (Eds), Cambridge University Press, 2007.
- [20] J. Leskovec, L. A. Adamic, and B. A. Huberman. The Dynamics of Viral Marketing. *ACM Transactions on the Web*, 2007.
- [21] D. Liben-Nowell, J. Novak, R. Kumar, P. Raghavan, and A. Tomkins. Geographic Routing in Social Networks. In *Proc. National Academy of Sciences*, number 33, pages 11623–11628. National Acad Sciences, 2005.
- [22] G. S. Manku, M. Naor, and U. Wieder. Know thy Neighbor’s Neighbor: the Power of Lookahead in Randomized P2P Networks. In *Proc. 36th ACM Symposium on Theory of Computing (STOC)*, 2004.
- [23] S. Milgram. The Small World Problem. *Psychology Today*, pages 60–67, 1967.
- [24] T. Moscibroda, S. Schmid, and R. Wattenhofer. On the Topologies Formed by Selfish Peers. In *Proc. 25th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, 2006.
- [25] T. Moscibroda, S. Schmid, and R. Wattenhofer. When Selfish Meets Evil: Byzantine Players in a Virus Inoculation Game. In *Proc. 25th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, 2006.
- [26] M. E. J. Newman and M. Girvan. Finding and Evaluating Community Structure in Networks. *Physical Review E*, 69, 2004.
- [27] T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.
- [28] T. Roughgarden and Éva Tardos. How Bad is Selfish Routing? *Journals of the ACM*, 49(2):236–259, 2002.
- [29] C. Wang, J. C. Knight, and M. C. Elder. On Computer Viral Infection and the Effect of Immunization. In *Proc. 16th Annual Computer Security Applications Conference (ACSAC)*, 2000.