

The Price of Malice: A Game-Theoretic Framework for Malicious Behavior in Distributed Systems

Thomas Moscibroda, Stefan Schmid, and Roger Wattenhofer

Abstract. In recent years, game theory has provided insights into the behavior of distributed systems by modeling the players as utility-maximizing agents. In particular, it has been shown that selfishness causes many systems to perform in a globally sub-optimal fashion. Such systems are said to have a large *price of anarchy*. In this article, we extend this field of research by allowing some players to be malicious rather than selfish. What, we ask, is the impact of malicious players on the system consisting of otherwise selfish players? In particular, we introduce the *price of malice* as a measure that captures how much the system's efficiency degrades in the presence of malicious players, compared to a purely selfish environment. As a specific example, we analyze the price of malice of a game that models the containment of the spread of viruses. In this game, each player or node can choose whether to install antivirus software. Then, a virus starts from a random node and recursively infects all neighboring nodes that are not inoculated. We establish various results about this game. For instance, we quantify how much the presence of malicious players can deteriorate or—in case of highly risk-averse selfish players—improve the social welfare of the distributed system.

I. Introduction

The introduction of microeconomic models and game theory in computer science has led to insights into today's large-scale distributed systems. Many aspects of peer-to-peer (p2p) networks or the Internet, which typically connect different utility-optimizing stakeholders or agents, have been studied from a game-theoretic point of view over the last several years. However, it is well known that the noncooperation challenge in distributed systems is not restricted

to selfishness. Many p2p networks and distributed systems are faced with the problem of *malicious* participants or adversaries who try—independently of their own cost—to degrade the utility of the entire system, to attack correctness of certain computations, or to cause endless changes that render the system unstable. In some sense, these malicious adversaries may be seen as acting selfishly, but with a cost function that is the inverse of the system's social welfare.

This article introduces a mathematical framework that seeks to combine two fruitful threads of distributed systems research: game theory and fault-tolerance. We consider a system of selfish individuals whose only goal is to optimize their own benefit, and add malicious players who attack the system in order to deteriorate its overall performance. We ask, what is the impact of the malicious players on a selfish system's efficiency? We exemplify our theory by giving an analysis of a virus inoculation game.

To study the impact of malicious players on a given system formally, we introduce the *price of malice* of selfish systems. The price of malice is a ratio that expresses how much the presence of malicious players deteriorates the social welfare of a system consisting of selfish players. More technically, the price of malice is the ratio between the social welfare or performance achieved by a selfish system containing a number of *malicious* players and the social welfare achieved by an entirely selfish society.

It is interesting to compare the price of malice with the notion of the *price of anarchy* [Koutsoupias and Papadimitriou 99]. The price of anarchy captures the maximum degradation of a system due to selfish behavior of its users or participants. That is, the price of anarchy relates the social welfare generated by players acting in an egoistic manner to an optimal solution obtained by perfectly collaborating participants. In comparison, the price of malice's reference point is not a socially optimal welfare, but the welfare achieved by an entirely selfish system.

The price of anarchy and the price of malice are therefore two orthogonal measures that describe inherent properties of distributed socioeconomic systems. Specifically, a system may have a small price of anarchy, but a large price of malice, and vice versa. The fact that a system has a large price of anarchy indicates that it is necessary to design mechanisms (such as taxes or payment schemes) that force players to collaborate more effectively. However, it is more difficult to improve or repair systems having a large price of malice, since malicious players do not respond to (financial) incentives. Often, the only solution is to defend the system against malicious intruders, or at least to ensure that the number of malicious players in the system remains small. By introducing a model that formally comprises the notions of *malicious Nash equilibria*, the *malicious price of anarchy*, and the *price of malice*, we are able to analyze what happens in

selfish systems if the aim of one or more players is to hinder the system in its operation or to bog down its performance as much as possible.

The price of malice crucially depends on the amount of information the selfish players have about the presence and behavior of the malicious players, and how they respond to this information. In other words, the utility function that eventually defines the selfish players' reaction depends on how they *subjectively* perceive and judge the threat of malicious players. Hence, the utility of selfish players is computed using the *perceived expected cost* rather than the unknown actual cost. For example, it can be shown that in case of risk-averse players, the presence of malicious players may actually *improve* the social welfare. Specifically, there are situations in which—in view of the risk caused by malicious players—selfish players become more willing to collaborate, thereby improving the social welfare compared to a system without malicious players. To the best of our knowledge, this is the first framework that allows for an *analytical quantification* of this interesting and counterintuitive phenomenon, which we call the *fear factor*.¹ Potentially, this gives rise to new research questions in many areas including distributed systems, economics, politics, and sociology.

In this article, we investigate a concrete example in which selfish and malicious players interact. In this simple game, we consider a network of players (or equivalently, nodes or peers) such that each player can choose between paying for inoculation and risking infection by a virus. After the nodes have made their choices, a virus starts at some random node and propagates recursively to all neighboring nodes that are not inoculated.

Besides studying the price of malice of this game, we also consider the question of *stability*. Particularly, we investigate how many malicious players suffice to prevent the system from stabilizing.

2. Related Work

Security and robustness of distributed systems against malicious faults are of prime importance and have been an active field of research for many years. For example, in [Castro and Liskov 99], a malicious-fault-tolerant distributed file system has been proposed. Possibility and impossibility results on the malicious consensus problem have been achieved in a variety of models and settings [Dolev 82, Lamport et al. 82, Shostak et al. 80]. In addition, the distributed computing community has come up with results and solutions for a wide variety of other problems with malicious faults. Examples are *clock synchronization*

¹Recently, this phenomenon has also been referred to as the *windfall of malice*.

[Welch and Lynch 88], *broadcast* [Koo 04, Srikant and Toueg 87], and *quorum systems* [Malkhi and Reiter 98]. All of the above works assume that nonmalicious players (or processes) are benevolent and attempt to reach a common goal.

Malicious behavior is also subject to intensive research in cryptography. For instance, there is a large body of work in the area of *secure multiparty computation* [Yao 82]. In this context, the two interesting papers [Halpern and Teague 04] and [Abraham et al. 06] need to be mentioned, which consider self-interested players, e.g., in secret sharing problems. In the latter, it is shown that Nash equilibria exist for secret sharing whereby no member of a coalition of size up to k nodes can do better even if the entire coalition defects, provided that players prefer to get the secret information than not to get it.

This article strives for combining fault-tolerance research with game theory. In this respect, our work is related to the notions of *fault-tolerant implementation* introduced in [Eliaz 02] and of *BAR fault tolerance* introduced in [Aiyer et al. 05] (see also [Li et al. 06, Clement et al. 08]). In [Eliaz 02], implementation problems are investigated with k faulty players in the population, but neither their number nor their identity is known. A planner's objective then is to design an equilibrium whereby the nonfaulty players act according to his rules. In [Aiyer et al. 05], the authors describe an asynchronous state machine replication protocol that tolerates **B**yzantine, **A**ltruistic, and **R**ational (BAR) behavior. Interestingly, they find that the presence of Byzantine players can simplify the design of protocols if players are risk-averse. In contrast to our work, rather than evaluating the degradation due to noncooperative behavior, a concrete protocol for a *cooperative backup service* is provided.

To the best of our knowledge, the first paper to study equilibria with a malicious player is [Karakostas and Viglas 07]. The authors consider a routing application whereby a single malicious player uses his flow through the network in an effort to cause the maximum possible damage. In order to evaluate the impact of such malicious behavior, a coordination ratio is introduced that compares the social costs of the worst *Wardrop equilibrium* to the social costs of the best *minimax saddle point*. A different basing point is used, whereby the benevolent coordinator cannot influence the malicious player: in the "social optimum" only the costs of the players that can be coordinated are minimized, while the malicious player seeks to maximize costs.

There exists other work on game-theoretic systems in which not every participating agent acts in a rational or selfish way. In the *Stackelberg theory* [Roughgarden 01], for instance, the model consists of a set of selfish players, but a certain fraction of the entire population is *controlled by a global leader*. The leader's goal is to devise a strategy that induces an optimal or near-optimal so-called Stackelberg equilibrium.

There has been a large number of new results since the initial publication of this work in [Moscibroda et al. 06]. The work [Babaioff et al. 09] analyzes the so-called *windfall of malice* (which is related to our fear factor) in the context of *nonatomic congestion games*. The players in this work are less risk-averse, and a different solution concept is employed: while malicious players cannot hide their actual strategies, they are bound to route a certain amount of traffic that may lead to a windfall of malice.

The authors conjecture that there is no windfall of malice in their congestion game model in networks in which the set of paths is a *matroid*, and also find that the absence of a windfall of malice is related to the absence of a *generalized Braess's paradox*.

Further results following [Moscibroda et al. 06] include [Roth 08, Diaz et al. 09, Chakrabarty et al. 09, Gradwohl and Reingold 08, Gabarro et al. 08, Eidenbenz et al. 07, Lelarge and Bolot 09, Fultz and Grossklags 09, Chen and Kempe 09, Li et al. 08, Cohen et al. 08, Meier et al. 08]. The recent work [Roth 08] studies the price of malice in linear congestion games, making weaker assumptions on the behavior of rational and malicious players, and providing a no-regret analysis. The authors of [Diaz et al. 09] study where fear in mediation can improve the outcome, and compare virus inoculation and congestion games. Malicious players in a load-balancing game have been investigated in [Chakrabarty et al. 09]. The authors of [Gradwohl and Reingold 08] start with the observation that a Nash equilibrium does not provide any guarantees for a selfish player in the presence of irrational participants, but they can prove that large games are naturally fault-tolerant.

In [Gabarro et al. 08], game theory is used to analyze the effect of a number of service failures during the execution of a grid orchestration. The authors of [Eidenbenz et al. 07] investigate a scenario with selfish players that are influenced by a malicious mechanism designer rather than a benevolent one. It is shown that sometimes, a malicious mechanism designer can worsen the outcomes at no cost. In order to cope with the problem of insufficient self-protection against viruses in distributed systems, [Lelarge and Bolot 09] proposes an insurance as a powerful incentive mechanism that pushes agents to invest in self-protection.

The reference [Fultz and Grossklags 09] explores economic security incentives and studies the trade-off between protection and self-insurance. Moreover, note that there also exists work on auctions with agents that derive utility from the disutility of others [Brandt et al. 07, Morgan et al. 03]. Both papers derive symmetric Bayesian Nash equilibria for spiteful agents in first-price and second-price sealed bid auctions. The authors show that the revenue equivalence between second-price and first-price auctions breaks down with spiteful agents, with second-price outperforming first-price. The reference [Chen and Kempe 09]

studies auctions whose bidders are embedded in a social or economic network. Bidders who do not win the auction themselves might derive utility from the auction, namely, when a friend wins; on the other hand, when an enemy or competitor wins, a bidder might derive negative utility. As far as applications are concerned, at OSDI 2008, the peer-to-peer application FlightPath [Li et al. 08] was presented, which is robust to selfish and malicious behavior. Finally, there is an interest in selfish behavior in self-stabilization [Cohen et al. 08], and in the effects of altruism in addition to selfishness [Hoefer and Skopalik 09, Meier et al. 08]; in particular, [Meier et al. 08] studies the same game-theoretic framework as we introduce here, but considers a social setting.

Of course, coordination and collaboration problems involving malicious players also exist outside game theory. For example, consider the *collaborative filtering* problem studied in [Awerbuch et al. 05] (see also [Awerbuch et al. 04]), where there are malicious players among the n players participating in a reputation system like eBay. The goal of the honest players is to find a good object, and they can use a shared billboard to collaborate. The dilemma of an honest player is how to balance the desire to reduce her cost by taking advantage of the reports posted by honest peers against the fear of being exploited by adopting reports posted by malicious players.

Finally, there exist many virus propagation models in the literature. While traditional epidemiological models characterize infection in terms of birth rate and death rate of the virus [Bailey 75], more recently, models have been proposed for all kinds of graphs, including Internet-like power-law graphs [Pastor-Satorras and Vespignani 02]. In particular, our game-theoretic virus propagation model is based on [Aspnes et al. 05], whose authors model the containment of the spread of viruses in general graphs. They characterize equilibria in selfish environments and also give an approximation algorithm for the centralized, nonselfish case.

3. Framework

We present our model in two steps. First, we discuss the virus inoculation game derived from [Aspnes et al. 05]. Subsequently, we introduce our framework of malicious game theory including the definition of the price of malice.

3.1. Virus Inoculation Game

Similarly to [Aspnes et al. 05], we model the virus inoculation game as a scenario with n strategic players each of whom corresponds to a node in an undirected

grid $G[r, c]$ of r rows and c columns.² Henceforth, we will refer to the upper left corner of the grid as $G[0, 0]$, i.e., indices start with 0.

Each node i has two choices: either do nothing and risk infection by a virus, or inoculate itself by installing antivirus software. For a node, installing the antivirus software has the obvious advantage that it becomes immune against infection. On the other hand, the process of installing the software entails a cost in terms of money and/or time. Hence, a strategic player may or may not opt for inoculation depending on which choice maximizes his own utility.

The nodes' choices can be summarized by a strategy profile $\vec{a} \in \{0, 1\}^n$, where $a_i = 1$ signifies that node i installs the antivirus software, and $a_i = 0$ that it does not install it. We call nodes i with $a_i = 1$ *secure*, and denote the set of secure nodes by $I_{\vec{a}}$. After the nodes have made their choices, the adversary picks some node uniformly at random as a starting point for infection. Infection then propagates on the network graph and infects all nonsecure nodes that are in the same nonsecure connected component as the starting point of infection. Technically, we associate an *attack graph* $G_{\vec{a}} = G \setminus I_{\vec{a}}$ with \vec{a} . It is essentially the network graph in which all secure nodes and their incident edges are removed.

In this article, we consider the following costs: installing antivirus software on a selfish node entails an inoculation cost of 1 at this node. If a selfish node does not inoculate and becomes infected, it suffers a loss equal to L . Therefore, the cost of a selfish node i can be summarized as follows:

$$\text{cost}_i(\vec{a}) = a_i + (1 - a_i) \cdot L \cdot \frac{k_i}{n},$$

where k_i/n is the probability that node i is infected, conditioned on the event that it does not install the antivirus software. Thereby, k_i is the size of the connected component containing i in $G_{\vec{a}}$. Finally, the *social cost* of a strategy profile \vec{a} is the sum of all individual costs, i.e., $\text{Cost}(\vec{a}) = \sum_{j \in S} \text{cost}_j(\vec{a})$, where S denotes the set of all selfish players. When the strategy profile \vec{a} is clear from the context, we sometimes use abbreviations cost_i and Cost to denote individual cost and social cost, respectively.

3.2. Malicious Game Theory

In order to understand the impact of malicious players on the selfish system, we extend the virus inoculation game with malicious players. Formally, there are n nodes in the network. Of these n nodes, b are *malicious nodes* that do not strive to minimize their own costs. Instead, the goal of these malicious nodes is to deteriorate the overall system performance as much as possible, i.e., to

²Our results can be generalized to other highly regular, low-dimensional graphs such as the two-dimensional torus, i.e., a grid that wraps around at the boundaries.

maximize the resulting social cost of the solution. The remaining $s := n - b$ nodes are *selfish* and aim at maximizing their own utility. We denote the set of malicious and selfish players by B and S , respectively. Then $b := |B|$, $s := |S|$, and $n = s + b$.

While selfish nodes behave as discussed in Section 3.1, we assume that the malicious nodes pursue the following strategy: they claim to be inoculated (i.e., they proclaim their strategy to be $a_i = 1$), but actually they are not. In order to emphasize that malicious nodes are only seemingly secure, we denote the set of really inoculated and secure selfish nodes by $I_{\vec{a}}^{\text{self}}$. The attack graph resulting from strategy profile \vec{a} is then $G_{\vec{a}} = G - I_{\vec{a}}^{\text{self}}$. This is the network graph without secure, selfish nodes, but including all malicious nodes. We can therefore define the individual cost incurred at a selfish node $i \in S$ as follows.

Definition 3.1. (Actual individual cost.) We define the (actual) individual cost $\text{cost}_i(\vec{a})$ of a node $i \in S$ as

$$\text{cost}_i(\vec{a}) := a_i + (1 - a_i) \cdot L \cdot \frac{k_i}{n},$$

where k_i is the size of the connected component of node i in the attack graph $G_{\vec{a}}$.

Notice that in spite of its being equivalent to the corresponding definition in Section 3.1, we call this cost *actual individual cost*. This is to emphasize the fact that selfish players may not know about the existence of malicious players, and therefore, they are unable to compute their actual individual cost. Even if they are aware of the malicious players' existence, they might not know the malicious players' exact locations or strategies. In other words, with the addition of malicious players, selfish nodes no longer have *perfect knowledge* about the network and its nodes' choices.

In case of imperfect information, a node might deal with its uncertainty in different ways. For example, a node might be risk-averse and act in a conservative manner. These observations imply that before the location and strategies of malicious players are revealed (i.e., before the virus infection occurs), a selfish player i experiences a *perceived individual cost* $\widehat{\text{cost}}_i(\vec{a})$. This perceived cost can differ from the *actual individual cost* $\text{cost}_i(\vec{a})$ that a node eventually has to pay.

Definition 3.2. (Perceived individual cost.) Consider a selfish game with malicious players in which selfish players have imperfect knowledge about the existence, location, or strategy of malicious players. In this case, the perceived individual cost $\widehat{\text{cost}}_i(\vec{a})$ of a selfish player i captures the cost expected by player i given his knowledge about the malicious players. This cost depends on the underlying model.

The strategic decisions of selfish players can be based only on the *perceived cost* (not on their actual individual costs), since the actual individual cost can be computed only once the locations and strategies of malicious players are revealed. In this article, we will study the following two basic models.

Definition 3.3. (Oblivious model.) In the *oblivious model*, selfish players are not aware of the existence of malicious players. That is, selfish players assume that all other players in the system are selfish as well.

Definition 3.4. (Nonoblivious model.) In the *nonoblivious model*, selfish players know about the existence of malicious players. Specifically, we assume that every selfish player knows b , the number of malicious players in the system, but he does not know about these players' exact locations or strategies. Moreover, we assume that selfish players are highly risk averse in the sense that they aim at minimizing their maximal individual cost. Let \mathcal{D} be the set of possible distributions of malicious players among all players. A selfish player i experiences a perceived individual cost of

$$\widehat{\text{cost}}_i(\vec{a}) := \max_{d \in \mathcal{D}} \{\text{cost}_i(\vec{a}, d)\},$$

where $\text{cost}_i(\vec{a}, d)$ denotes the actual costs of i if the malicious players are distributed according to $d \in \mathcal{D}$.

In the virus inoculation game, and in an oblivious model, the perceived cost is typically smaller than the actual cost: a node $i \in S$ does not take into consideration the malicious nodes that may increase the size of i 's attack component. In the nonoblivious risk-averse model, on the other hand, a node actually overestimates its expected actual cost by considering a worst-case scenario: a selfish player assumes that the malicious nodes are—from his individual point of view—distributed in a worst-case fashion among all players. Therefore, the perceived individual cost may be larger than the actual cost.

Since our goal is to understand the impact of malicious behavior on a system of selfish players, the cost of malicious players is not included in the social cost. If it were, it would in general be easy for malicious players to arbitrarily deteriorate the social welfare of a system by simply increasing their own costs as much as possible. Moreover, since malicious players are malicious anyway, there is no particular reason why the overall system should care about these players' costs.

The total *social cost* $\text{Cost}(\vec{a})$ of a strategy is defined as the sum of the (actual) individual costs of all selfish players. Since each node in the same connected component of $G_{\vec{a}}$ has the same probability of infection, the l_i selfish nodes in the i th attack component face a loss of $l_i \cdot (Lk_i/n)$ if the component is infected.

Definition 3.5. (Social cost.) The social cost is given by the sum of the actual individual costs of selfish players:

$$\text{Cost}(\vec{a}) = \sum_{j \in \mathcal{S}} \text{cost}_j(\vec{a}) = |I_{\vec{a}}^{\text{self}}| + \frac{L}{n} \sum_{i=1}^l k_i l_i,$$

where on the right-hand side, the first term is the inoculation cost and the second term is the infection cost, and where k_1, k_2, \dots, k_l are the sizes of the components in $G_{\vec{a}}$, and l_1, l_2, \dots, l_l are the sizes of the same components without counting the malicious nodes. We refer to the cost due to inoculation as the *inoculation cost* $\text{Cost}_{\text{inoc}}$, and to the cost due to the virus infections as the *infection cost* $\text{Cost}_{\text{infect}}$.

The social cost of a setting in which all nodes perfectly collaborate, i.e., in which there are neither selfish nor malicious nodes, is called the *social optimum*.

Definition 3.6. (Optimal social cost.) The optimal social cost Cost_{opt} is the sum of all the players' actual individual costs in case of perfect collaboration.

Recall that the Nash equilibrium describes a situation in which no selfish node has an incentive to unilaterally change its strategy. In the following, we extend the definition of a Nash equilibrium to incorporate malicious nodes. The *malicious Nash equilibrium* (MNE) describes a configuration in which no selfish player can reduce his *perceived* cost by changing his strategy, given that the strategies of all other players are fixed.³

Definition 3.7. (Malicious Nash equilibrium (MNE).) Let $\vec{a}[i|x]$ be the strategy vector that is identical to \vec{a} except for the i th component a_i , which is replaced by x . In a malicious Nash equilibrium, no selfish player $i \in \mathcal{S}$ has an incentive to change his strategy if the strategies of all other (selfish and malicious) players are fixed, i.e.,

$$\forall i \in \mathcal{S} : \widehat{\text{cost}}_i(\vec{a}) \leq \widehat{\text{cost}}_i(\vec{a}[i|a'_i]),$$

for every possible strategy a'_i .

While the malicious Nash equilibrium must be defined by the *perceived* individual costs, the resulting social cost is determined by the *actual* costs. After all, it is the actual individual costs that players will eventually have to pay. In the

³Notice that we do not define the malicious Nash equilibrium with *actual* individual costs, because they are not known to the players.

following, we will refer to the social cost of the *worst malicious Nash equilibrium* of a problem instance I as $\text{Cost}_{\text{mNe}}(I, b)$.

It is well known that selfish and malicious players often interact in a manner that yields suboptimal solutions. The degree of degradation resulting from selfish and malicious players compared to the social optimum is captured by the *price of malicious anarchy*.

Definition 3.8. (Price of malicious anarchy.) The price of malicious anarchy captures how much worse a malicious Nash equilibrium can be compared to a collaborative optimal solution. More formally, in a scenario with b malicious players, the price of malicious anarchy $\text{PoMA}(b)$ is the ratio between the worst-case social cost of a malicious Nash equilibrium divided by the minimal social cost, i.e., for all problem instances I ,

$$\text{PoMA}(I, b) = \frac{\max_{\text{mNe}} \text{Cost}_{\text{mNe}}(I, b)}{\text{Cost}_{\text{opt}}(I)}.$$

Note that in the absence of malicious players, i.e., if the system consists of selfish players only, the price of malicious anarchy is equivalent to the well-known price of anarchy (PoA) studied in classical game theory. Specifically, $\text{PoA} = \text{PoMA}(0)$.

With these definitions, we are ready to define the *price of malice*, which describes the degree of suboptimality resulting from malicious players in an otherwise selfish system. A high price of malice indicates that an economic system is particularly vulnerable to malicious attacks. On the other hand, if the price of malice is low, the system consisting of selfish players is stable enough to tolerate malicious participants. Clearly, the degree of degradation may depend on the number of malicious players in the game. Hence, the price of malice is a function of b .

Definition 3.9. (Price of malice.) The price of malice captures the ratio between the worst malicious Nash equilibrium with b malicious players and the price of anarchy in a purely selfish system. Formally, for problem instance I ,

$$\text{PoM}(I, b) = \frac{\text{PoMA}(I, b)}{\text{PoMA}(I, 0)}.$$

As will be discussed in Section 4.4, we may also speak of the inverse of the price of malice as the game's *fear factor* $\Psi(b)$. That is, a game's fear factor is given by $\Psi(b) := 1/\text{PoM}(b)$.

4. Virus Game Analysis

In order to derive results for the price of malice in various models, we have to establish structural properties of Nash equilibria and the social optimum in the virus inoculation game. We begin with a simple characterization of Nash equilibria if there are no malicious nodes. The following lemma is derived from the analogous lemma in [Aspnes et al. 05].

Lemma 4.1. *In a pure Nash equilibrium \vec{a} ,*

- (a) *every component in the attack graph $G_{\vec{a}}$ has a size of at most n/L ;*
- (b) *inserting any secure node into $G_{\vec{a}}$ yields a component size of at least n/L .*

Lemma 4.1 implies that if $L \geq n$, all nodes will inoculate in the Nash equilibrium. Therefore, for the rest of this article, we assume that $L < n$.

4.1. Social Optimum

If the inoculation strategies of the individual nodes are planned by a benevolent centralized coordinator, the welfare of the system is maximized. In the following, we will derive an asymptotically tight bound on the cost of this social optimum. Throughout this section, perceived costs equal actual costs because when studying the social optimum, we do not consider malicious players, i.e., $b = 0$ and therefore $s = n$.

Theorem 4.2. *The optimal social cost if all players in S act altruistically is $\text{Cost}_{\text{opt}} \in \Theta(s^{2/3}L^{1/3})$. More specifically,*

$$\frac{1}{3}\sqrt{\pi} \cdot s^{2/3}L^{1/3} \leq \text{Cost}_{\text{opt}} \leq 4s^{2/3}L^{1/3}.$$

Proof. We prove the upper and lower bounds in turn.

Lower bound: If all nodes collaborate to achieve the optimal solution, then $l_i = k_i$, and hence the social cost is given by

$$\text{Cost} = |I_{\vec{a}}| + \frac{L}{n} \sum_{i=1}^l k_i^2,$$

where $|I_{\vec{a}}|$ is the number of inoculated nodes, and the k_i are the sizes of the l components in the attack graph. This sum is minimized when all k_i are of equal

size, say size K . While each secure node has a cost of 1, every other node has an expected cost of $L \cdot K/n$. Hence, setting $\gamma := |I_{\vec{a}}|$ and because $s = n$, the optimal social cost can be bounded as

$$\text{Cost}_{\text{opt}} \geq \gamma + (s - \gamma) \left(\frac{LK}{s} \right). \quad (4.1)$$

A relationship between γ and K follows from a simple geometric argument: if a component in the attack graph is of size K , the number of inoculated nodes at the component's border must be at least $2\pi\sqrt{K}/\pi = 2\sqrt{\pi K}$ (circumference of a disk with area K). Since the total number of such components is at least $\frac{s-\gamma}{K}$ and since each inoculated node can be on the border of at most two components, γ can be expressed as

$$\gamma \geq \frac{s-\gamma}{K} \cdot 2\sqrt{\pi K} \cdot \frac{1}{2} = (s-\gamma) \sqrt{\frac{\pi}{K}}.$$

By solving this inequality for γ , it follows that $\gamma \geq s \cdot \sqrt{\pi/K} / (1 + \sqrt{\pi/K})$. On the other hand, it can be observed that in the optimal solution, for $s > L$, no node is inoculated if all its four neighbors are inoculated. From this, it can be derived that in an optimal solution, $\gamma \leq \frac{s}{2}$. Plugging these two bounds into inequality (4.1), we see that the optimal social cost is at least

$$\text{Cost}_{\text{opt}} \geq s \cdot \frac{\sqrt{\pi/K}}{1 + \sqrt{\pi/K}} + \frac{LK}{2}.$$

The first term of the above expression is monotonically decreasing in K in the range $0, \dots, s$, whereas the second one is monotonically increasing. Therefore, taking the minimum of the two terms for a specific K yields a lower bound on Cost_{opt} . In setting

$$K := \frac{2}{3} \sqrt{\pi} \cdot \left(\frac{s}{L} \right)^{2/3},$$

the second term yields $\frac{1}{3} \sqrt{\pi} \cdot s^{2/3} L^{1/3}$. The first term evaluates to

$$\frac{\sqrt{3/2} \cdot \sqrt[4]{\pi}}{1 + \sqrt{3/2} \cdot \sqrt[4]{\pi}} s^{2/3} L^{1/3} > \frac{1}{3} \sqrt{\pi} \cdot s^{2/3} L^{1/3}.$$

Consequently, we obtain the following lower bound on the cost of the social optimum:

$$\text{Cost}_{\text{opt}} \geq \frac{1}{3} \sqrt{\pi} \cdot s^{2/3} L^{1/3} \in \Omega(s^{2/3} L^{1/3}).$$

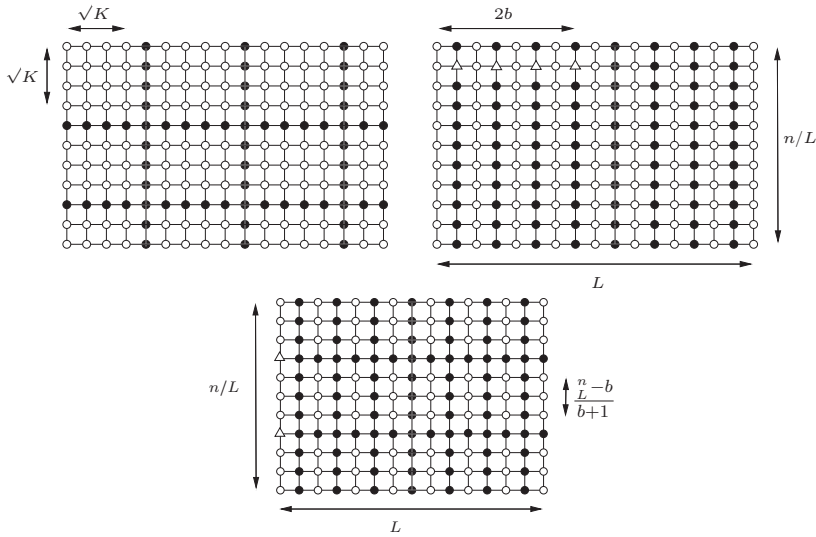


Figure 1. Top left: Upper bound for social optimum. White nodes are insecure; black nodes are secure. Top right: Malicious Nash equilibrium for $G[n/L, L]$ for the oblivious model. Insecure malicious nodes are denoted by white triangles. They are located in a way that may yield an attack component of size $(b + 1)n/L + b$. Bottom: Example with large social cost for the nonoblivious, risk-averse model.

Upper bound: Having established a lower bound on the optimal social cost, we now explicitly construct a solution that is asymptotically optimal and proves the tightness of the above lower bound. Given an arbitrary grid $G[r, c]$, we inoculate the nodes as follows. Let $K := (s/L)^{2/3}$. We secure all nodes in the columns $G[\cdot, i\sqrt{K}]$ for $i \in \{1, \dots, \lfloor c/(\sqrt{K} + 1) \rfloor\}$ and rows $G[i\sqrt{K}, \cdot]$ for $i \in \{1, \dots, \lfloor r/(\sqrt{K} + 1) \rfloor\}$. Consequently, all attack components are of size at most $\sqrt{K} \times \sqrt{K} = K$ as illustrated in Figure 1 (top left). Hence, the total infection cost is at most

$$L \cdot (s - |I_{\vec{a}}|) \frac{K}{s} < LK = s^{2/3} L^{1/3}.$$

It remains to bound the inoculation cost. In an ideal setting where the components perfectly fit into $G[r, c]$ without leftovers, for each component of size K in the attack graph there are exactly $2\sqrt{K} + 1$ inoculated nodes. Let X denote the number of components. Then $X \cdot (K + 2\sqrt{K} + 1) = s$, and therefore, on plugging in the definition of K , we obtain $X = s/[(s/L)^{2/3} + 2(s/L)^{1/3} + 1]$. The

number of inoculated nodes γ is at most

$$\begin{aligned}\gamma &\leq X \cdot (2\sqrt{K} + 1) \leq \frac{s(2\sqrt{K} + 1)}{\left(\frac{s}{L}\right)^{2/3} + 2\left(\frac{s}{L}\right)^{1/3} + 1} \\ &< s^{1/3}L^{2/3} \cdot \left(2\left(\frac{s}{L}\right)^{1/3} + 1\right) = 2s^{2/3}L^{1/3} + s^{1/3}L^{2/3} \\ &\leq 3s^{2/3}L^{1/3}.\end{aligned}$$

Combining the infection and inoculation costs, we can bound the optimal social cost by

$$\text{Cost}_{\text{opt}} < s^{2/3}L^{1/3} + 3s^{2/3}L^{1/3} = 4s^{2/3}L^{1/3}. \quad \square$$

4.2. Price of Anarchy

The price of anarchy compares the social cost of the worst Nash equilibrium (without malicious nodes) to the minimal social cost. In the upcoming section, we will first compute Cost_{Ne} , which is the maximal cost of any Nash equilibrium. Together with the bound for the social optimum in Section 4.1, the price of anarchy will follow.

Lemma 4.3. *The social cost of the worst Nash equilibrium is $\text{Cost}_{\text{Ne}} = \Theta(s)$.*

Proof. First, we show that $\text{Cost}_{\text{Ne}} = \Omega(s)$. Consider a grid $G[s/L, L]$ consisting of an even number L of columns of size s/L . Assume that columns $G[\cdot, 2i]$ for $i \in \{0, 1, \dots, L/2 - 1\}$ consist of insecure nodes only, while all nodes in the remaining columns are secure. According to Lemma 4.1, this situation constitutes a Nash equilibrium. Observe that every second column is inoculated, engendering an inoculation cost of $s/2$. Moreover, with probability $1/2$, the virus starts at an insecure node, yielding infection cost $s/L \cdot L$. The social cost is therefore $\text{Cost}_{\text{Ne}} = s/2 + 1/2 \cdot s/L \cdot L = s$.

It remains to show that $O(s)$ is an upper bound for any Nash equilibrium. Since at most each of the $s = n$ nodes can be inoculated, the inoculation cost cannot exceed s . By Lemma 4.1, we also know that the infected component's size is at most s/L , entailing a total infection cost of at most s as well. Hence, $\text{Cost}_{\text{Ne}} \leq 2s$, and the claim follows. \square

By Theorem 4.2 and Lemma 4.3, we get the following result.

Theorem 4.4. *For the price of anarchy (PoA), we have*

$$\frac{1}{4} \cdot \left(\frac{s}{L}\right)^{1/3} \leq \text{PoA} \leq \frac{6}{\sqrt{\pi}} \cdot \left(\frac{s}{L}\right)^{1/3}.$$

Proof. As for the upper bound, we have

$$\text{PoA} = \frac{\text{Cost}_{\text{Ne}}}{\text{Cost}_{\text{opt}}} \leq \frac{2s}{\frac{1}{3}\sqrt{\pi} \cdot s^{2/3}L^{1/3}} \leq \frac{6s^{1/3}}{\sqrt{\pi} \cdot L^{1/3}},$$

and as for the lower bound, we have $\text{PoA} \geq s/(4 \cdot s^{2/3}L^{1/3})$. □

4.3. Oblivious Model

We begin our study of the price of malice with the oblivious model in which players are clueless about the existence of malicious players in the system (cf. Section 3). As a consequence, it follows that since nodes underestimate the attack components' sizes, the nodes' perceived individual costs are smaller than the actual individual costs. It turns out that in the presence of malicious nodes, the social costs increase in the number of malicious nodes.

Lemma 4.5. *In the oblivious model, the social cost is at least $\text{Cost}_{\text{mNe}} \in \Omega(s + \frac{nb^2}{L})$ for $b < \frac{L}{2} - 1$, and $\text{Cost}_{\text{mNe}} \in \Omega(sL)$ otherwise.*

Proof. Consider again a grid $G[n/L, L]$ with n/L rows and L columns, where every second column consists of secure nodes only. For simplicity, let L be even. Suppose that in each of the first b secure columns there is one malicious node; see Figure 1 (top right). In case $b \geq \frac{L}{2} - 1$, every secure column that separates two insecure columns contains one malicious node. The remaining malicious nodes can be placed at arbitrary places in the secure columns. Because selfish nodes are not aware of the existence of malicious nodes in the network, the perceived cost is $\widehat{\text{cost}}_i = 1$ for inoculated nodes, and $\widehat{\text{cost}}_i = \frac{n/L}{n} \cdot L = 1$ for the other selfish nodes. Hence, the situation constitutes a *malicious Nash equilibrium*.

For computing the social costs of this malicious Nash equilibrium, we distinguish two cases, depending on whether the number of malicious nodes is smaller than $L/2 - 1$. For the first case, assume that $b \geq L/2 - 1$. Because there is at least one malicious node in every secure column that separates two insecure columns, all selfish and malicious players form one large attack component. Consequently, each insecure selfish node $i \in S$ is infected with probability 1 and therefore $\text{Cost}_{\text{mNe}} \geq s \cdot L$.

For the second case, assume that $b < L/2 - 1$. Each of the first secure columns contains exactly one malicious node. Since L is even, there are $s/2 - b$ secure nodes, and hence the inoculation cost is $s/2 - b$. With probability $((b+1)n/L + b)/n$, the infection starts at an insecure or a malicious node of an attack component of size $(b+1) \cdot n/L$, yielding a cost of $(b+1) \cdot n/L \cdot L = n(b+1)$. Moreover, with probability $(s/2 - (b+1)n/L)/n$, an insecure column of size n/L

is hit. Thus, for $b < L/2 - 1$, we get the following lower bound on the social cost:

$$\begin{aligned} \text{Cost}_{\text{mNe}} &= \left(\frac{s}{2} - b\right) + \frac{\frac{(b+1)n}{L} + b}{n} \cdot n(b+1) + \frac{\frac{s}{2} - (b+1)\frac{n}{L}}{n} \cdot \frac{n}{L} \cdot L \\ &= s + \frac{nb^2}{L} + \frac{nb}{L} + b^2 \in \Omega\left(s + \frac{nb^2}{L}\right). \end{aligned} \quad \square$$

Lemma 4.6. *In the oblivious model, the social cost is at most*

$$\text{Cost}_{\text{mNe}} \in O\left(\min\left\{sL, s + \frac{b^2n}{L}\right\}\right).$$

Proof. Since at most every selfish node can be inoculated, it is clear that $\text{Cost}_{\text{inoc}} = O(s)$. It remains to study the infection cost. The infection cost of a node in some component i is L times the probability of this component being hit by the virus, i.e., $L \cdot k_i/n$. Hence, the total infection cost is given by

$$\text{Cost}_{\text{infec}} = \sum_i l_i \cdot \frac{k_i}{n} \cdot L = \frac{L}{n} \sum_i l_i \cdot k_i,$$

where k_i is the size of the attack components (including malicious nodes), and l_i is the number of selfish nodes in this component. In order to bound $\text{Cost}_{\text{infec}}$ from above, let S_{byz} denote the set of components in the attack graph that contain at least one malicious node, and let $\overline{S_{\text{byz}}}$ be the remaining components. We can rewrite the equation above as

$$\text{Cost}_{\text{infec}} = \frac{L}{n} \cdot \left[\sum_{i \in S_{\text{byz}}} l_i \cdot k_i + \sum_{i \in \overline{S_{\text{byz}}}} l_i \cdot k_i \right],$$

that is, we consider the infection cost of components with at least one malicious node separately from the remaining “malicious player-free” components. In the following, let

$$\text{Cost}_{\text{infec}}^{\text{byz}} := \frac{L}{n} \sum_{i \in S_{\text{byz}}} l_i k_i, \quad \text{Cost}_{\text{infec}}^{\overline{\text{byz}}} := \frac{L}{n} \sum_{i \in \overline{S_{\text{byz}}}} l_i k_i.$$

We have to prove that neither $\text{Cost}_{\text{infec}}^{\text{byz}}$ nor $\text{Cost}_{\text{infec}}^{\overline{\text{byz}}}$ exceeds $O(s + \frac{b^2n}{L})$.

As we have shown in the proof of Lemma 4.3 in Section 4.2, the total infection cost of a network consisting only of selfish nodes cannot exceed s . Because in our case nodes are oblivious about the existence of malicious nodes, attack

components without malicious nodes behave as they would in an entirely selfish environment. Therefore, $\text{Cost}_{\text{infect}}^{\text{byz}} \in O(s)$.

It remains to compute the infection cost of those attack components that include at least one malicious node. Let b_i be the number of malicious nodes in the i th component in S_{byz} , and note that $\sum_i b_i = b$. By Lemma 4.1, we know that in the absence of malicious nodes, the size of an attack component is at most $k_i \leq n/L$. Therefore, one malicious node can increase a component by at most n/L nodes plus itself. From this it follows that the size of an attack component i is bounded by

$$k_i \leq (b_i + 1) \cdot \frac{n}{L} + b_i \quad \text{and} \quad l_i \leq (b_i + 1) \cdot \frac{n}{L}.$$

Using this relationship between b_i and the size of the attack component, we can bound $\text{Cost}_{\text{infect}}^{\text{byz}}$ as

$$\begin{aligned} \text{Cost}_{\text{infect}}^{\text{byz}} &= \frac{L}{n} \sum_{i \in S_{\text{byz}}} l_i \cdot k_i \\ &\leq \frac{L}{n} \sum_{i \in S_{\text{byz}}} \left[(b_i + 1) \cdot \frac{n}{L} \cdot \left((b_i + 1) \cdot \frac{n}{L} + b_i \right) \right] \\ &= \sum_{i \in S_{\text{byz}}} \left[(b_i + 1)^2 \frac{n}{L} + b_i(b_i + 1) \right] \\ &< \sum_{i \in S_{\text{byz}}} \left[(b_i + 1)^2 \left(\frac{n}{L} + 1 \right) \right] \\ &= \left(\frac{n}{L} + 1 \right) \cdot \sum_{i \in S_{\text{byz}}} (b_i + 1)^2. \end{aligned}$$

Given the constraint that $b_i \geq 1$ for every b_i , and because $\sum_i b_i = b$, the above convex function assumes its maximum for a single positive $b_i = b$. Consequently,

$$\text{Cost}_{\text{infect}}^{\text{byz}} \leq \left(\frac{n}{L} + 1 \right) \cdot \sum_{i \in S_{\text{byz}}} (b_i + 1)^2 \leq \left(\frac{n}{L} + 1 \right) \cdot (b + 1)^2 \in O\left(\frac{b^2 n}{L}\right).$$

On the other hand, it is clear that at most every selfish node can be infected, and hence $\text{Cost}_{\text{infect}}^{\text{byz}} + \text{Cost}_{\text{infect}}^{\text{byz}} \leq sL$. The proof is concluded by adding the upper bounds for $\text{Cost}_{\text{inoc}}$, $\text{Cost}_{\text{infect}}^{\text{byz}}$, and $\text{Cost}_{\text{infect}}^{\text{byz}}$. \square

Combining Lemmas 4.5 and 4.6 leads to the following theorem, which captures the social cost in the virus inoculation game in the presence of b malicious players among selfish, oblivious nodes.

Theorem 4.7. *The social cost in a malicious Nash equilibrium with b malicious nodes in the oblivious model is $\text{Cost}_{\text{mNe}} \in \Theta(s + \frac{b^2 n}{L})$ for $b < \frac{L}{2} - 1$, and $\text{Cost}_{\text{mNe}} \in \Theta(sL)$ otherwise.*

Proof. In both cases, the lower bound follows from Lemma 4.5. As for the upper bound, note that for $b < L/2 - 1$, and due to $L \leq n = s + b$, we see that $b < (s + b)/2$ and therefore $b < s$. Then the term $s + b^2 n/L$ asymptotically cannot exceed the term sL , and therefore the claim follows. As for the second case, note that for $b \geq \frac{L}{2} - 1$, the term sL is asymptotically less than or equal to $s + b^2 n/L$. \square

Finally, we can derive tight bounds on the price of malicious anarchy and the price of malice by bringing together the results of Theorems 4.2, 4.4, and 4.7.

Theorem 4.8. *In the virus inoculation game with b malicious nodes among selfish oblivious nodes, the price of malicious anarchy and the price of malice are*

$$\text{PoMA}(b) \in \Theta\left(\left(\frac{s}{L}\right)^{1/3} \left(1 + \frac{b^2}{L} + \frac{b^3}{sL}\right)\right)$$

and

$$\text{PoM}(b) \in \Theta\left(1 + \frac{b^2}{L} + \frac{b^3}{sL}\right)$$

for $b < \frac{L}{2} - 1$. Otherwise,

$$\text{PoMA}(b) \in \Theta\left(s^{1/3} L^{2/3}\right) \quad \text{and} \quad \text{PoM}(b) \in \Theta(L).$$

Proof. Consider the case $b < \frac{L}{2} - 1$. For the price of malicious anarchy, we have

$$\text{PoMA}(b) = \frac{\text{Cost}_{\text{mNe}}}{\text{Cost}_{\text{opt}}} = \frac{\Theta(s + \frac{b^2(b+s)}{L})}{\Theta(s^{2/3} L^{1/3})} \in \Theta\left(\left(\frac{s}{L}\right)^{1/3} \cdot \left(1 + \frac{b^2}{L} + \frac{b^3}{sL}\right)\right).$$

From this, the price of malice is computed as follows:

$$\text{PoM}(b) = \frac{\text{PoMA}(b)}{\text{PoA}} \in \Theta\left(1 + \frac{b^2}{L} + \frac{b^3}{sL}\right).$$

The case $b \geq \frac{L}{2} - 1$ follows along the same lines by plugging in the corresponding expressions of Theorem 4.7. \square

Our results on the price of malice in the oblivious case support the intuition that in the absence of knowledge about the existence of malicious players, the quality of the global solution (i.e., the resulting social cost) deteriorates as the number of malicious players increases. In the next section, we will show that the situation may change as soon as selfish players are *aware* of the existence of malicious players.

4.4. Nonoblivious Model

Having studied the oblivious model, we now turn our attention to the nonoblivious case, in which selfish players are aware of the existence of malicious players. If selfish nodes knew about the exact locations of malicious nodes, they would be able to compute their optimal choice exactly. If selfish nodes know only the *number* of malicious nodes in the system, however, the optimal strategy of a player becomes more complex, and the impact on the social cost more interesting. Specifically, it turns out that in this nonoblivious case, the “fear factor” may actually encourage players to act less selfishly and cooperate. Put differently, there may be settings in which the existence of malicious players helps to improve the global social cost, rendering the price of malice less than 1.

4.4.1. The Fear Factor. Intuitively, in the presence of malicious players, nodes may be more willing to pay for inoculation. However, with our framework, we find the interesting phenomenon that the selfish players’ awareness of the existence of malicious players may lead to an *improvement* of the overall system behavior, i.e., the *social welfare*.

In the following, we show the existence of such a *fear factor*, which describes the gain of the overall social efficiency in a selfish system if selfish players are afraid of malicious individuals among them. The fear factor is determined by the ratio between the social cost of the worst malicious Nash equilibrium with b malicious players and the worst Nash equilibrium in a purely selfish system. Technically, we can define the fear factor Ψ as the inverse of the price of malice, i.e.,

$$\Psi(b) := \frac{1}{\text{PoM}(b)}.$$

In other words, the fear factor Ψ quantifies how much the threat of a common enemy can unite selfish individuals, and to what degree the global social performance is improved.

To see that the fear factor can be positive already on a simple topology, consider the following example. We are given a 1-dimensional chain of n nodes, arranged from left to right and numbered $1, \dots, n$, where n is an integer multiple

of $1 + n/L$ and $1 + n/(bL)$. For the fear factor, we have

$$\Psi(b) = \frac{1}{\text{PoM}(b)} = \frac{\text{PoA}}{\text{PoMA}(b)}.$$

Thus $\Psi(b)$ is at least the cost of some Nash equilibrium divided by the cost of the worst malicious equilibrium. Consider the Nash equilibrium for which nodes $i \cdot (1 + n/L)$ are inoculated, for $i = 1, \dots, n/(1 + n/L)$. This equilibrium bears expected costs $n/(1 + n/L) + (n - n/(1 + n/L))/n \cdot n = n$. On the other hand, in the worst malicious equilibrium, the b malicious players form one large attack component: nodes $i \cdot (1 + n/(bL))$ for $i = 1, \dots, n/(1 + n/(bL))$ are inoculated, where nodes $j \cdot (1 + n/(bL))$ for $i = 1, \dots, b$ are malicious. This yields a cost of at most

$$\begin{aligned} & \left(\frac{n}{1 + \frac{n}{bL}} - b \right) + \frac{(b+1) \cdot (\frac{n}{bL} + 1) - 1}{n} \cdot L \cdot (b+1) \cdot \frac{n}{bL} \\ & + \frac{n - (b+1) \cdot (\frac{n}{bL} + 1)}{n} \cdot \frac{n}{bL} \cdot L \\ & = \left(\frac{n}{1 + \frac{n}{bL}} - b \right) + \left((b+1) \cdot \left(\frac{n}{bL} + 1 \right) - 1 \right) \cdot \frac{b+1}{b} \\ & + \frac{n - (b+1) \cdot (\frac{n}{bL} + 1)}{b} < n, \end{aligned}$$

for, e.g., $L = b = 5$. Since the cost of the Nash equilibrium is n and the cost of the worst malicious equilibrium is strictly less than n , this implies $\Psi(b) > 1$.

The existence of a fear factor has been documented in various economic and social models. By combining a game-theoretic framework with the classical notion of malicious players from distributed computing and cryptography, our model allows for an analytical quantification of a system's fear factor Ψ from a computational point of view.

In the virus inoculation game, the fear factor may be both negative and positive. What is interesting to note, however, is that this fear factor Ψ cannot be arbitrarily large, regardless of the number of malicious players b in the system. Instead, as we will show in the next section, the price of malice can never drop below the constant $\frac{\sqrt{\pi}}{48}$, and hence the fear factor is bounded above by $\Psi \leq \frac{48}{\sqrt{\pi}}$ (cf. Section 4.4.2). That is, the social welfare or efficiency gained due to the fear factor cannot exceed a factor of $\Psi \leq \frac{48}{\sqrt{\pi}}$.

4.4.2. Price of Malice. We now derive bounds on the price of malice. By doing so, we also derive an upper bound on the fear factor in our game. Observe that in the nonoblivious case, every selfish node inoculates if $b \geq \frac{n}{L}$, implying a social cost of s . If $b < \frac{n}{L}$, the resulting social costs are bounded, by the following lemma.

Lemma 4.9. For $b < \frac{n}{2L}$, the social cost in a malicious Nash equilibrium in case of nonoblivious, risk-averse players with b malicious nodes is at least

$$\text{Cost}_{\text{mNe}} \geq \frac{s}{2} + \frac{bL}{4}.$$

For all values of b , we have $\text{Cost}_{\text{mNe}} \geq \frac{s}{2}$.

Proof. We start with the more interesting case $b < \frac{n}{2L}$. Consider a grid with L columns each containing n/L nodes. All nodes in columns $2i + 1$ for $i = 0, 1, \dots, \frac{L}{2} - 1$ and all nodes in rows $j \cdot ((n/L - b)/(b + 1))$ for $j = 1, 2, \dots, b$ are inoculated. That is, as illustrated in Figure 1 (top right), each component of insecure selfish nodes is of size $(n/L - b)/(b + 1)$.

First, we show that this configuration constitutes a malicious Nash equilibrium in the risk-averse, nonoblivious case with b malicious nodes. Consider an *insecure node* in some column i . If all b secure nodes in this column are malicious, the size of the resulting attack component is $(n/L - b)/(b + 1) \cdot (b + 1) + b = n/L$. Hence i 's perceived infection cost is

$$\widehat{\text{cost}}_i = L \cdot \frac{(n/L - b)/(b + 1) \cdot (b + 1) + b}{n} = 1,$$

which equals the cost of inoculation. Next, consider an *inoculated selfish node* i and distinguish two cases. In the first case, i separates two components consisting of insecure selfish players, and a change of i 's strategy would merge two components of size $(n/L - b)/(b + 1)$ into a single connected component of insecure selfish nodes. Every malicious node can connect another component of size $(n/L - b)/(b + 1)$ (and itself) to the component containing i . Therefore, the size of the resulting attack component can be as large as

$$\left(2 \cdot \frac{\frac{n}{L} - b}{b + 1} + 1\right) + \left(b \cdot \frac{\frac{n}{L} - b}{b + 1} + b\right) = \frac{b + 2}{b + 1} \left(\frac{n}{L} - b\right) + b + 1 > \frac{n}{L} + \frac{1}{b + 1}.$$

The perceived cost of i without inoculation is therefore

$$\widehat{\text{cost}}_i > L \cdot \frac{\frac{n}{L} + \frac{1}{b + 1}}{n} = 1 + \frac{L}{n(b + 1)} > 1.$$

In the second case, we consider a “crossing” node i that is located at the crossing of a secure row and column. Consider the column to the right (or to the left) of i . If all inoculated nodes in this column are malicious, the entire column plus node i becomes one large attack component. Hence, the perceived cost of i is

$$\widehat{\text{cost}}_i > L \cdot \frac{\frac{n}{L} + 1}{n} > 1.$$

In other words, no selfish node has an incentive to change its strategy, and the situation in Figure 1 (bottom) constitutes a malicious Nash equilibrium. In the sequel, we bound the social cost of this equilibrium from below under the assumption that all b malicious nodes are in column 1. Note that our construction guarantees that this is always possible if $b < \frac{n}{2L}$.

We start with the sum of the infection costs $\text{Cost}_{\text{infect}}^0$ of insecure nodes in column 0. The number of insecure selfish nodes in this component is $\frac{n}{L} - b$. Hence the expected sum of infection costs is

$$\text{Cost}_{\text{infect}}^0 = \left(\frac{n}{L} - b\right) \cdot \frac{\frac{n}{L} - b + b}{n} \cdot L = \frac{n}{L} - b.$$

Let μ be the number of insecure nodes in columns 3, 5, etc. The sum of the infection costs $\text{Cost}_{\text{infect}}^r$ of the remaining attack components (each being of size $(n/L - b)/(b + 1)$) is

$$\text{Cost}_{\text{infect}}^r = \mu \cdot \frac{\frac{n}{L} - b}{n(b + 1)} \cdot L > \mu \cdot \left(\frac{1}{b + 1} - \frac{L}{n}\right).$$

Because the number of insecure nodes in these small attack components is $\mu = \frac{L-1}{2} \cdot (\frac{n}{L} - b)$, it follows that

$$\begin{aligned} \text{Cost}_{\text{infect}}^r &> \frac{L-1}{2} \cdot \left(\frac{n}{L} - b\right) \cdot \left(\frac{1}{b + 1} - \frac{L}{n}\right) \\ &> \frac{1}{2(b + 1)} \left(n - \frac{n}{L} - bL + b\right) - \frac{L}{2}. \end{aligned}$$

Finally, we also need to calculate the total inoculation cost of this topology. Clearly, all $s/2$ nodes in even columns are secure. (Recall that column and row indices start with 0.) Furthermore, b nodes in each odd column (except for the first column) are also inoculated. Hence, the total inoculation cost $\text{Cost}_{\text{inoc}}$ becomes

$$\text{Cost}_{\text{inoc}} = \frac{s}{2} + \frac{bL}{2} - b = \frac{s}{2} + b \left(\frac{L}{2} - 1\right).$$

Combining the various costs, we see that the social cost of the malicious Nash equilibrium is

$$\begin{aligned} \text{Cost}_{\text{mNe}}(b) &\geq \frac{s}{2} + b \left(\frac{L}{2} - 1\right) + \frac{n}{L} - b + \frac{1}{2(b + 1)} \left(n - \frac{n}{L} - bL + b\right) - \frac{L}{2} \\ &\geq \frac{s}{2} + \frac{bL}{4} \end{aligned}$$

for $b \leq \frac{n}{2L}$ and $b \geq 3$.

Finally, note that if $b \geq \frac{n}{2L}$, at least half of the selfish nodes inoculate and hence $\text{Cost}_{\text{mNe}}(b) \geq s/2$. \square

With this lower bound on the social cost of a malicious Nash equilibrium, we can now derive the price of malicious anarchy as well as the price of malice for the nonoblivious, risk-averse model.

Theorem 4.10. *In the nonoblivious, risk-averse model with b malicious nodes, the price of malicious anarchy is at least*

$$\text{PoMA}(b) \geq \frac{1}{8} \left(\left(\frac{s}{L} \right)^{1/3} + \frac{b}{2} \left(\frac{L}{s} \right)^{2/3} \right)$$

for $b < \frac{n}{2L}$. For all b , we have $\text{PoMA}(b) \geq \frac{1}{8} \left(\frac{s}{L} \right)^{1/3}$.

Proof. Lemma 4.9 gives us a lower bound on the social cost of a malicious Nash equilibrium in the nonoblivious, risk-averse model with b malicious nodes. On the other hand, we have seen in Lemma 4.2 that the optimal social cost is at most $4s^{2/3}L^{1/3}$. Hence

$$\text{PoMA}(b) \geq \frac{\frac{s}{2} + \frac{bL}{4}}{4s^{2/3}L^{1/3}} = \frac{1}{8} \left(\frac{s^{1/3}}{L^{1/3}} + \frac{bL^{2/3}}{2s^{2/3}} \right).$$

The second lower bound follows analogously. □

Theorem 4.11. *In the nonoblivious, risk-averse model with b malicious nodes, the price of malice is*

$$\text{PoM}(b) \geq \frac{\sqrt{\pi}}{48} \left(1 + \frac{bL}{2s} \right)$$

for $b < \frac{n}{2L}$. For all b , we have $\text{PoM}(b) \geq \frac{\sqrt{\pi}}{48}$.

Proof. In order to derive the price of malice, we can apply our bound from Theorem 4.10 and the upper bound on the price of anarchy established in Theorem 4.4. Specifically,

$$\text{PoM}(b) = \frac{\text{PoMA}(b)}{\text{PoA}} \geq \frac{\frac{1}{8} \left(\left(\frac{s}{L} \right)^{1/3} + \frac{b}{2} \left(\frac{L}{s} \right)^{2/3} \right)}{\frac{6s^{1/3}}{\sqrt{\pi} \cdot L^{1/3}}}.$$

The theorem then follows from arithmetic simplifications. Again, the second lower bound follows in an analogous way. □

As mentioned in Section 4.4.1, the above bound also implies that the fear factor cannot be arbitrarily large. Instead, Theorem 4.11 shows that the fear factor is bounded above by a constant, specifically $\Psi \leq 48/\sqrt{\pi}$.

5. Stability Considerations

In the previous section, we studied the degradation of the social welfare in a selfish system caused by malicious players. However, besides trying to reduce the optimality of certain outcomes of games, malicious players might also attack the *stability* of a system. In this section, we therefore continue our studies by capturing the amount of instability that can be caused by malicious players in an otherwise selfish system. Particularly, we are interested in the question, how many malicious players suffice in order to keep the system from stabilizing?

In the following, we generalize the model of Section 4 to *arbitrary* network graphs. We assume that the malicious players aim at destabilizing the system by repeatedly announcing that they have changed from the insecure to the secure state and back in a worst-case fashion. Thereby, we consider an oblivious model whereby selfish nodes are not aware of the stability attack. We use the following definitions.

Definition 5.1. (*b*-Stable/*b*-unstable.) We call a game *b*-stable if *b* malicious players cannot prevent the system from reaching a Nash equilibrium. Similarly, a game is called *b*-unstable if *b* malicious players are sufficient to prevent a Nash equilibrium from ever being reached in the presence of oblivious selfish players.

For the virus inoculation game, the following stability properties can be shown.

Theorem 5.2.

- (i) *Generally, the virus inoculation game is not 1-stable.*
- (ii) *For certain restricted classes of network graphs, the virus inoculation game is 1-stable.*
- (iii) *The virus inoculation game is always 2-unstable.*

Proof. *Claim (i):* This claim already holds in simple graphs. Assume that n/L is an integer and that $L > 1$, and consider a one-dimensional chain of nodes $\{0, 1, \dots, n-1\}$. Let the nodes $i \cdot n/L$ be secure, for $i \in \{0, 1, \dots, L-1\}$. By Lemma 4.1, this situation constitutes a Nash equilibrium.

Now assume that node n/L is malicious, and that it changes to the insecure state. Then all other nodes $j \in \{1, 2, 3, \dots, n/L-1, n/L+1, \dots, 2n/L-1\}$ have an incentive to inoculate. However, once such a node j has become secure, node n/L can return to the secure state, yielding components of size smaller than

n/L . Consequently, j is bound to become insecure again. These changes can be repeated forever.

Claim (ii): Interestingly, there are robust graphs in which no single node can destabilize the system. To see this, consider a complete graph, in which each node is connected to all other nodes. From Lemma 4.1, it follows that in this network, all Nash equilibria have just a single attack component. Let \mathcal{C} denote the set of nodes of this component, and let $\bar{\mathcal{C}} := V \setminus \mathcal{C}$ be the set of the remaining (secure) nodes. Also by Lemma 4.1, we have that in any Nash equilibrium, the size of \mathcal{C} is either n/L or $n/L - 1$.

Moreover, observe that independently of which node is malicious and of how the malicious node acts, a situation will eventually be reached with the two components as described above. However, the system having converged to such a state, there exist only four possibilities: the malicious node belongs either to the node set \mathcal{C} or to the node set $\bar{\mathcal{C}}$, and either $|\mathcal{C}| = n/L$ or $|\mathcal{C}| = n/L - 1$. It is run of the mill to verify that in all cases, a malicious node can enforce at most one additional change.

Claim (iii): We use the fact that in the virus inoculation game, a pure Nash equilibrium always exists, and that in the absence of malicious nodes, selfish nodes stabilize quickly [Aspnes et al. 05]. Assume that the malicious nodes first act like selfish nodes until such a classic Nash equilibrium is reached. Now consider an arbitrary secure node $u_1 \in V$, and assume that it is malicious. If u_1 becomes insecure, then according to Lemma 4.1, an attack component \mathcal{C} emerges that consists of n/L or more nodes. If $|\mathcal{C}| > n/L$, at least one node v in \mathcal{C} has an incentive to change to a secure state. Let \mathcal{C}' be the component of v when u_1 is secure, but not v . Assume that after v has changed, u_1 becomes secure again. There are two possibilities. If $|\mathcal{C}'| < n/L$, v will return to the insecure state, and the changes can be repeated forever with only one malicious node. If $|\mathcal{C}'| = n/L$, a second malicious (previously insecure) node u_2 in \mathcal{C}' can force v to become insecure again.

Finally, if $|\mathcal{C}| = n/L$, nodes are indifferent between becoming secure or not. Of course, however, another malicious node on the edge of \mathcal{C} can cause endless changes also in this case. \square

6. Conclusion

This article has initiated the study of distributed systems consisting of both selfish and malicious players. Using our models, we have derived bounds on the *price of malice* in oblivious and nonoblivious systems. Moreover, we have quantified

and bounded above the *fear factor*, which is the *gain* in system efficiency arising from the increased willingness of selfish individuals to cooperate caused by malicious players.

Several questions are left for future research. For example: What is the price of malice in a virus inoculation game on other topologies, e.g., on a hypercubic Pastry network? And what is the price of malice of other games, e.g., of a *caching game* [Chun et al. 04]? It seems that in certain selfish routing games in which a single node can attract a large amount of traffic by announcing short distances to all other nodes, the result is a larger price of malice than in congestion games, for example. Another direction for future work is to study the *impact of knowledge* on the resulting fear factor in nonoblivious models. Specifically, one could assume that players are aware not only of the existence of malicious players, but also of their approximate whereabouts or their statistical distribution. Intuitively, such additional knowledge should decrease the selfish players' incentives for collaboration and thus lower the fear factor.

Our game-theoretic framework can be applied to potentially many economic and social systems. For instance, recently, the framework developed in this article has also been used in the context of social networks to study the effect of selfish players that exhibit “altruistic” behavior toward their friends [Meier et al. 08]: the paper shows that while friendship is always beneficial compared to a purely selfish setting, there are situations in which stronger social ties yield a lower social welfare.

Acknowledgments. A preliminary version of this work has been published as “When Selfish Meets Evil: Byzantine Players in a Virus Inoculation Game” at the 25th ACM Annual Symposium on Principles of Distributed Computing (PODC) [Moscibroda et al. 06].

References

- [Abraham et al. 06] Ittai Abraham, Danny Dolev, Rica Gonen, and Joe Halpern. “Distributed Computing Meets Game Theory: Robust Mechanisms for Rational Secret Sharing and Multiparty Computation.” In *Proceedings of the 25th Annual ACM Symposium on Principles of Distributed Computing*, pp. 53–62. New York: ACM Press, 2006.
- [Aiyer et al. 05] Amitanand Aiyer, Lorenzo Alvisi, Allen Clement, Mike Dahlin, Jean-Philippe Martin, and Carl Porth. “BAR Fault Tolerance for Cooperative Services.” In *Proceedings of the 20th ACM Symposium on Operating Systems Principles*, pp. 45–58. New York: ACM Press, 2005.
- [Aspnes et al. 05] James Aspnes, Kevin Chang, and Aleksandr Yampolskiy. “Inoculation Strategies for Victims of Viruses and the Sum-of-Squares Partition Problem.”

- In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 43–52. Philadelphia: SIAM, 2005.
- [Awerbuch et al. 04] Baruch Awerbuch, Boaz Patt-Shamir, David Peleg, and Mark Tuttle. “Collaboration of Untrusting Peers with Changing Interests.” In *Proceedings of the 5th ACM Conference on Electronic Commerce*, pp. 112–119. New York: ACM Press, 2004.
- [Awerbuch et al. 05] Baruch Awerbuch, Boaz Patt-Shamir, David Peleg, and Mark Tuttle. “Adaptive Collaboration in Peer-to-Peer Systems.” In *Proceedings of the 25th IEEE International Conference on Distributed Computing Systems*, pp. 71–80. Washington, DC: IEEE Computer Society, 2005.
- [Babaioff et al. 09] Moshe Babaioff, Robert Kleinberg, and Christos H. Papadimitriou. “Congestion Games with Malicious Players.” *Games and Economic Behavior* 67:1 (2009), 22–35.
- [Bailey 75] N. T. Bailey. *The Mathematical Theory of Infectious Diseases and Its Applications*. New York: Hafner Press, 1975.
- [Brandt et al. 07] Felix Brandt, Tuomas Sandholm, and Yoav Shoham. “Spiteful Bidding in Sealed-Bid Auctions.” In *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, pp. 1207–1214. San Francisco: Morgan Kaufmann Publishers Inc., 2007.
- [Castro and Liskov 99] Miguel Castro and Barbara Liskov. “Practical Byzantine Fault Tolerance.” In *Proceedings of the Third Symposium on Operating Systems Design and Implementation*, pp. 173–186. Berkeley, CA: USENIX Association, 1999.
- [Chakrabarty et al. 09] Deeparnab Chakrabarty, Chinmay Karandey, and Ashish Sangwan. “The Effect of Malice on Social Optimum in Linear Load Balancing Games.” Preprint arXiv:0910.2655v2, 2009.
- [Chen and Kempe 09] Po-An Chen and David Kempe. “Bayesian Auctions with Friends and Foes.” In *Proceedings of the 2nd International Symposium on Algorithmic Game Theory*, pp. 335–346. Berlin: Springer-Verlag, 2009.
- [Chun et al. 04] Byung-Gon Chun, Kamalika Chaudhuri, Hoeteck Wee, Marco Barreno, Christos H. Papadimitriou, and John Kubiawicz. “Selfish Caching in Distributed Systems: A Game-Theoretic Analysis.” In *Proceedings of the 23rd Annual ACM Symposium on Principles of Distributed Computing*, pp. 21–30. New York: ACM Press, 2004.
- [Clement et al. 08] A. Clement, H. Li, J. Napper, J. Martin, L. Alvisi, and M. Dahlin. “BAR Primer.” In *Proceedings of the 38th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, pp. 287–296. Los Alamitos, CA: IEEE Computer Society, 2008.
- [Cohen et al. 08] Johanne Cohen, Anurag Dasgupta, Sukumar Ghosh, and Sébastien Tixeuil. “An Exercise in Selfish Stabilization.” *ACM Trans. Auton. Adapt. Syst.* 3:4 (2008), Article No. 15.
- [Diaz et al. 09] Josep Diaz, Dieter Mitsche, Navin Rustagi, and Jared Saia. “On the Power of Mediators.” In *Internet and Network Economics: 5th International Workshop, WINE 2009, Rome, Italy, December 14–18, 2009, Proceedings*, Lecture Notes in Computer Science 5929, pp. 455–462. Berlin: Springer-Verlag, 2009.

- [Dolev 82] D. Dolev. “The Byzantine Generals Strike Again.” *Journal of Algorithms* 3:1 (1982), 14–30.
- [Eidenbenz et al. 07] Raphael Eidenbenz, Yvonne Anne Oswald, Stefan Schmid, and Roger Wattenhofer. “Manipulation in Games.” In *Algorithms and Computation: 18th International Symposium, ISAAC 2007, Sendai, Japan, December 17–19, 2007, Proceedings*, Lecture Notes in Computer Science 4835, pp. 365–376. Berlin: Springer-Verlag, 2007.
- [Eliaz 02] Kfir Eliaz. “Fault Tolerant Implementation.” *Review of Economic Studies* 69 (2002), 589–610.
- [Fultz and Grossklags 09] Neal Fultz and Jens Grossklags. “Blue versus Red: Towards a Model of Distributed Security Attacks.” In *Financial Cryptography and Data Security: 13th International Conference, FC 2009, Accra Beach, Barbados, February 23–26, 2009, Revised Selected Papers*, Lecture Notes in Computer Science 5628, pp. 167–183. Berlin: Springer-Verlag, 2009.
- [Gabarro et al. 08] Joaquim Gabarro, Alina Garcia, Maria Serna, Peter Kilpatrick, and Alan Stewart. “Analysing Orchestrations Using Risk Profiles and Angel-Daemon Games.” In *Grid Computing: Achievements and Prospects*, edited by Sergei Gortatch, Paraskevi Fragopoulou, and Thierry Priol, pp. 121–132. New York: Springer, 2008.
- [Gradwohl and Reingold 08] Ronen Gradwohl and Omer Reingold. “Fault Tolerance in Large Games.” In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pp. 274–283. New York: ACM Press, 2008.
- [Halpern and Teague 04] Joseph Halpern and Vanessa Teague. “Rational Secret Sharing and Multiparty Computation.” In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing*, pp. 623–632. New York: ACM Press, 2004.
- [Hoefler and Skopalik 09] Martin Hoefler and Alexander Skopalik. “Altruism in Atomic Congestion Games.” In *Algorithms—ESA 2009: 17th Annual European Symposium, Copenhagen, Denmark, September 7–9, 2009, Proceedings*, Lecture Notes in Computer Science 5757, pp. 179–189. Berlin: Springer-Verlag, 2009.
- [Karakostas and Viglas 07] George Karakostas and Anastasios Viglas. “Equilibria for Networks with Malicious Users.” *Mathematical Programming A* 110:3 (2007), 591–613.
- [Koo 04] Chiu-Yuen Koo. “Broadcast in Radio Networks Tolerating Byzantine Adversarial Behavior.” In *Proceedings of the 23rd Annual ACM Symposium on the Principles of Distributed Computing*, pp. 275–282. New York: ACM Press, 2004.
- [Koutsoupias and Papadimitriou 99] E. Koutsoupias and C. H. Papadimitriou. “Worst-Case Equilibria.” In *STACS 99: 16th Annual Symposium on Theoretical Aspects of Computer Science, Tier, Germany, March 1999, Proceedings*, Lecture Notes in Computer Science 1563, pp. 404–413. Berlin: Springer-Verlag, 1999.
- [Lamport et al. 82] L. Lamport, R. Shostak, and M. Pease. “The Byzantine Generals Problem.” *ACM Trans. Program. Lang. Syst.* 4:3 (1982), 382–401.
- [Lelarge and Bolot 09] Marc Lelarge and Jean Bolot. “Economic Incentives to Increase Security in the Internet: The Case for Insurance.” In *Proceedings of the 28th IEEE*

- International Conference on Computer Communications*, pp. 1494–1502. Los Alamitos, CA: IEEE Press, 2009.
- [Li et al. 06] Harry C. Li, Allen Clement, Edmund L. Wong, Jeff Napper, Indrajit Roy, Lorenzo Alvisi, and Michael Dahlin. “BAR Gossip.” In *Proceedings of the 7th Symposium on Operating Systems Design and Implementation*, pp. 191–204. Berkeley, CA: USENIX Association, 2006.
- [Li et al. 08] H. Li, A. Clement, M. Marchetti, M. Kapritsos, L. Robinson, L. Alvisi, and M. Dahlin. “FlightPath: Obedience vs Choice in Cooperative Services.” In *Proceedings of the 8th Symposium on Operating Systems Design and Implementation*, pp. 355–368. Berkeley, CA: USENIX Association, 2008.
- [Malkhi and Reiter 98] D. Malkhi and M. Reiter. “Byzantine Quorum Systems.” *Journal of Distributed Computing* 11:4 (1998), 203–213.
- [Meier et al. 08] Dominic Meier, Yvonne Anne Oswald, Stefan Schmid, and Roger Wattenhofer. “On the Windfall of Friendship: Inoculation Strategies on Social Networks.” In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pp. 294–301. New York: ACM Press, 2008.
- [Morgan et al. 03] John Morgan, Ken Steiglitz, and George Reis. “The Spite Motive and Equilibrium Behavior in Auctions.” *Contributions to Economic Analysis & Policy* 2:1 (2003), Article No. 5.
- [Moscibroda et al. 06] Thomas Moscibroda, Stefan Schmid, and Roger Wattenhofer. “When Selfish Meets Evil: Byzantine Players in a Virus Inoculation Game.” In *Proceedings of the 25th Annual ACM Symposium on Principles of Distributed Computing*, pp. 35–44. New York: ACM Press, 2006.
- [Pastor-Satorras and Vespignani 02] R. Pastor-Satorras and A. Vespignani. “Immunization of Complex Networks.” *Physical Review Letters* 65 (2002), 036104.
- [Roth 08] A. Roth. “The Price of Malice in Linear Congestion Games.” In *Internet and Network Economics: 4th International Workshop, WINE 2008, Shanghai, China, December 17–20, 2008. Proceedings*, Lecture Notes in Computer Science 5385, pp. 118–125. Berlin: Springer-Verlag, 2008.
- [Roughgarden 01] Tim Roughgarden. “Stackelberg Scheduling Strategies.” In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, pp. 104–113. New York: ACM Press, 2001.
- [Shostak et al. 80] R. Shostak, M. Pease, and L. Lamport. “Reaching Agreement in the Presence of Faults.” *Journal ACM* 27:2 (1980), 228–234.
- [Srikant and Toueg 87] T. K. Srikant and S. Toueg. “Simulating Authenticated Broadcasts to Derive Simple Fault-Tolerant Algorithms.” *Journal of Distributed Computing* 2:2 (1987), 80–94.
- [Welch and Lynch 88] J. L. Welch and N. Lynch. “A New Fault-Tolerant for Clock-Synchronization.” *Information and Communication* 77 (1988), 1–36.
- [Yao 82] Andrew C. Yao. “Protocols for Secure Computations.” In *Proceedings of the 23rd IEEE Symposium on Foundations of Computer Science*, pp. 160–164. Washington, DC: IEEE Computer Society, 1982.

Thomas Moscibroda, Distributed Systems Research Group, Microsoft Research, Redmond, WA (moscitho@microsoft.com)

Stefan Schmid, Deutsche Telekom Laboratories, Technical University Berlin, Berlin, Germany (stefan@net.t-labs.tu-berlin.de)

Roger Wattenhofer, Computer Engineering and Networks Laboratory, ETH Zurich, Zurich, Switzerland (wattenhofer@tik.ee.ethz.ch)

Received June 22, 2009; accepted December 13, 2009.

