

Congestion Games with Malicious Players

Moshe Babaioff^{a,1} Robert Kleinberg^{b,*,2}
Christos H. Papadimitriou^{c,3}

^a*Microsoft Research - Silicon Valley,
1065 La Avenida,
Mountain View, CA 94043.*

^b*Department of Computer Science,
Cornell University
Ithaca, NY 14853.*

^c*Computer Science Division,
University of California at Berkeley,
Berkeley, CA 94720.*

Abstract

We study the equilibria of non-atomic congestion games in which there are two types of players: rational players, who seek to minimize their own delay, and malicious players, who seek to maximize the average delay experienced by the rational players. We study the existence of pure and mixed Nash equilibria for these games, and we seek to quantify the impact of the malicious players on the equilibrium. One counterintuitive phenomenon which we demonstrate is the “windfall of malice”: paradoxically, when a myopically malicious player gains control of a fraction of the flow, the new equilibrium may be more favorable for the remaining rational players than the previous equilibrium.

Key words: Selfish Routing, Malicious Behavior, Equilibrium, Congestion Games.

* Corresponding author.

Email addresses: moshe@microsoft.com (Moshe Babaioff),
rdk@cs.cornell.edu (Robert Kleinberg), christos@cs.berkeley.edu (Christos H. Papadimitriou).

¹ Research partially supported by NSF ITR Award ANI-0331659.

² Research supported by NSF awards CCF-0643934 and CCF-0729102 and by an NSF Mathematical Sciences Postdoctoral Research Fellowship.

³ Supported by NSF grant CCF-0635319, a MICRO grant, and a gift from Yahoo! Research.

1 Introduction

Game Theory is the study of strategic behavior of rational agents. Described this way, Game Theory sounds a little unrealistic, because much of what is going on in the real world is, intuitively, irrational. The standard game-theoretic response to this line of criticism is that what we intuitively call “irrational agents” are just rational players with very strange utility functions. This is a reasonable retort when it refers to the class of all normal-form games. However, much work in Game Theory, and especially in the interface between Game Theory, Networking, and Computation, is about standard types of games — such as auctions, congestion games, facility location games, network creation games, etc. Many of these are studied with the explicit ambition to apply the results to the real world. Since the utilities of these games cannot be arbitrarily “strange,” the original criticism stands.

In this paper we model one type of what is usually meant by “irrationality” in the above argument, namely *malicious players*. We define malicious players in the context of a particular *symmetric game*; suppose that, in an n -player symmetric game, the utilities of $m < n$ players, henceforth called malicious, change from the common utility shared by all players *to the negative sum of the utilities of the $n - m$ non-malicious players*. We are interested in the effect such change has on the quality as well as nature (pure versus mixed) of the game’s Nash equilibria. We define the *price of malice* to be the relative deterioration of the sum of the utilities of the non-malicious players when the remaining players turn malicious.

That malice has a price is not very surprising; what is somewhat unexpected is that *this price can be negative*, and malicious players may improve system performance, intuitively because their presence may incentivize the other players to forego antisocial selfish behavior. For example, consider a version of the prisoner’s dilemma with three players and three strategies: collaborate, defect, and *inspect*. When a player inspects, her own utility is negative, but that of any defecting player deteriorates as well. It is easy to see that the numbers can be set in such a way that, if any of the three players turns malicious (and therefore inspects, sacrificing her own well-being in order to hurt the others), then the other players end up collaborating at equilibrium, for a net increase in their sum of utilities — in fact, a relative increase that can be arbitrarily high.

We study the effects of malicious agents on non-atomic congestion games (Roughgarden and Tardos, 2002). In such games a continuum of players, comprising a flow of some specified value v , choose routes in a network with a single source and sink whose edges have load-dependent delays. It is known that such a game has a pure Nash equilibrium with equal delays for all; the quality of this equilibrium (compared to the “social optimum” minimum delay flow) has been studied extensively (Roughgarden and Tardos, 2002; Rough-

garden, 2005). But suppose instead that some fraction of the flow becomes controlled by a *malicious player* whose utility is the total delay by each of the other players. Does this new game have a pure Nash equilibrium? And, if it does, how does it compare with the equilibrium without a malicious player? We define the *price of malice* to be the limit of this deterioration per unit of malicious flow, as the portion of flow controlled by the malicious player goes to zero.

If flows are allowed to contain cycles (and therefore a malicious player can make loads arbitrarily high by going around in circles indefinitely), then it is easy to construct situations in which a malicious player can wreak havoc on a network (strictly speaking, such situations do not have a Nash equilibrium, and so we cannot speak of a price of malice). We show (Theorem 6) that even in acyclic networks the price of malice can be significant; upper bounding the price of malice by the network parameters (such as the number of edges and the “relative slope” of the delays) is an important open problem.

Perhaps the heaviest price of malice may be the fact that the presence of a malicious player *upsets the Nash equilibrium regime of congestion games*. Ordinarily, congestion games are known to always have a pure Nash equilibrium. In contrast, in Section 4 we notice that, in the presence of a malicious player, pure Nash equilibria may not exist. However, we prove two compensating results: First, there is always a “semi-pure” Nash equilibrium, in which only the malicious player mixes strategies. Second, if the delays are weakly concave (and in particular if they are linear) then pure Nash equilibria exist. The existence proofs rely on Kakutani’s Fixed Point Theorem and Prokhorov’s Theorem, two powerful results that we have not seen before applied in the context of congestion games.

1.1 Related work

Several recent papers have considered agents that are not acting rationally, and while the general philosophical direction of our work is somewhat similar to these works, there are still significant differences between our work and the papers we describe below.

Two papers (Morgan et al., 2003; Brandt et al., 2007) consider auctions with agents that derive utility from the disutility of others, and present similar results. Both papers derive symmetric Bayes Nash equilibria for spiteful agents in first-price and second-price sealed bid auctions. A spiteful agent’s value for an outcome is a convex combination of his own original profit and the total loss of the other agents (taken with coefficient α , the *spite* coefficient). The papers consider the equilibrium when *all* agents are spiteful with the same coefficient (unlike in our model in which only a small fraction of the flow is controlled by a malicious player, and this player is purely malicious, i.e. spiteful with coefficient 1). Interestingly they show that the revenue equivalence between

second-price and first-price auctions breaks down with spiteful agents, with second-price outperforming first-price.

Eliasz (2002) considers the problem of implementation when some agents are faulty, playing an arbitrary strategy, possibly in a malicious way. The paper presents a solution concept to allow implementation in case that up to k agents are faulty, but neither their identity nor their exact number are known. Unlike our model (which assumes that a malicious player will choose a strategy that maximizes the combined discontent of the non-malicious players, given the strategic choices of all other players) the non-rational agents in their model can play *arbitrarily* and the paper focuses on implementation issues, while we focus on quantifying the implications of malice on given systems.

Closest to our work is the paper by Karakostas and Viglas (2003) (henceforth KV) which studies equilibria for network congestion games with malicious users. The KV model of malicious behavior corresponds to a continuum of infinitesimal malicious players, collectively controlling the malicious flow. We allow for a more powerful malicious behavior by allowing coordination (modeled as a single myopic malicious agent controlling all the malicious flow). We show that coordination indeed leads to a different equilibrium concept for some networks, but not when latency functions are concave. The difference between the two equilibrium concepts, for general networks, arises from the fact that a single malicious player can use mixed strategies that are unavailable to a continuum of uncoordinated malicious players. The focus of the KV paper is on generalizing the notion of Price of Anarchy to the case that there is also a malicious flow in the system, by comparing the case that the good users are controlled by a single entity, to the case that they are behaving selfishly. One of the main open problems in this paper is the connection between the social cost at an equilibrium point with and without malicious users. Our work addresses this issue by defining the notion of Price of Malice and studying it.

Our study of the Price of Malice has strong thematic similarities to (Mosciro et al., 2006) which aims to study the implications of malicious behavior on systems consisting of selfish agents. The paper presents a concept of Price of Malice and says “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.”. Yet, there are many differences between this paper and ours. Important differences exist in the definition of equilibrium in the presence of malice and the definition of the Price of Malice. First, selfish players in their game are extremely risk averse and basically each one perceives the malicious agents as if they are all attacking him or her. Second, the definition of the Price of Malice is very different, as they look at the ratio between two different worst-case ratios (the price of anarchy with b malicious agents and with 0 malicious agents) even though those worst-case ratios may arise on different problem instances. Instead of this type of indirect comparison, we directly compare the outcome of games with a mixture of rational and

malicious agents to the outcome with only rational agents.

2 The Price of Malice

2.1 Definitions

2.1.1 Non-atomic congestion games

The following definitions are standard from the theory of congestion games, e.g. (Rosenthal, 1973), and readers familiar with this material are encouraged to proceed to Section 2.1.2. We use \mathbb{R}_+ to denote the set of non-negative real numbers.

Definition 1 (congestion game) A symmetric non-atomic congestion game (henceforth, simply called a “congestion game”) is specified by an ordered quadruple $\mathcal{G} = (E, \vec{\ell}, \Pi, v)$, whose components are called:

- the edge set E , a finite set;
- the vector of latency functions $\vec{\ell}$, a function from \mathbb{R}_+ to the vector space \mathbb{R}^E : for $e \in E$, the e -th component of $\vec{\ell}$ is denoted by ℓ_e , and is a non-decreasing function from \mathbb{R}_+ to \mathbb{R}_+ ;
- the path set Π , a subset of 2^E ; and
- the flow value v , a non-negative real number which we will sometimes denote by $v(\mathcal{G})$.

If E is the edge set of a (directed or undirected) graph G , and Π is a set of paths in G , then we call \mathcal{G} a network congestion game. We will use the terminology “edge” and “path” in describing abstract congestion games, though in the general case we do not expect E to be interpreted as a set of edges of a graph nor P as a set of paths in a graph.

Definition 2 (flow) A flow in a congestion game \mathcal{G} is a function f from Π to \mathbb{R}_+ . (One interprets $f(P)$ as the amount of flow using path P .) The flow value is $v(f) = \sum_{P \in \Pi} f(P)$. The set of all flows in \mathcal{G} is denoted by $F(\mathcal{G})$. The set of all flows whose flow value is equal to some number w is denoted by $F(\mathcal{G}, w)$.

Note that the set $F(\mathcal{G}, w)$ is a compact, convex set (in fact, a convex polytope) in \mathbb{R}^Π . $F(\mathcal{G})$ inherits the topology from $\mathbb{R}^{|E|}$. We will abbreviate $F(\mathcal{G}, w)$ to F when \mathcal{G}, w are understood from context.

Definition 3 (cost) If f is a flow, the load on an edge $e \in E$ is

$$x_e(f) = \sum_{P \in \Pi | e \in P} f(P).$$

The delay on a path $P \in \Pi$ is

$$L(P) = \sum_{e \in P} \ell_e(x_e(f)).$$

The cost of f is

$$C(f) = \sum_{P \in \Pi} f(P)L(P) = \sum_{e \in E} x_e(f)\ell_e(x_e(f)).$$

Definition 4 (Nash flow) If f is a flow, the set of best responses to f is the set $\arg \min_{P \in \Pi} L(P)$. A flow f in a congestion game \mathcal{G} is a Nash flow if $v(f) = v(\mathcal{G})$ and every path $P \in \Pi$ which satisfies $f(P) > 0$ is a best response to f . The Nash cost and Nash delay of \mathcal{G} , denoted by $C(\mathcal{G})$ and $D(\mathcal{G})$, are the quantities $C(f)$ and $D(f) = C(f)/v(f)$, respectively, where f is any Nash flow of \mathcal{G} . We will see in Proposition 1 below that $C(\mathcal{G})$ and $D(\mathcal{G})$ do not depend on the choice of the Nash flow f .

Definition 5 (potential function) For a congestion game \mathcal{G} , the potential function $\Phi_{\mathcal{G}}$ (denoted simply by Φ when the game \mathcal{G} is understood from context) is a real-valued function on $F(\mathcal{G})$ defined by

$$\Phi_{\mathcal{G}}(f) = \sum_{e \in E} \int_0^{x_e(f)} \ell_e(y) dy.$$

The following standard facts about the potential function will be useful to us.

Proposition 1 (Roughgarden and Tardos (2002)) The potential function $\Phi = \Phi_{\mathcal{G}}$ is a convex function on $F(\mathcal{G})$. It is strictly convex if all of the latency functions ℓ_e are strictly increasing. For a flow f of value $v(\mathcal{G})$, the following are equivalent:

- (1) f is a local minimum of Φ .
- (2) f is a global minimum of Φ .
- (3) f is a Nash flow.

Moreover, for any two Nash flows f, \tilde{f} we have $C(f) = C(\tilde{f})$, and furthermore the Nash delay $D(f)$ is equal to the delay $L(P)$ on any path $P \in \Pi$ satisfying $f(P) > 0$.

2.1.2 Congestion games with malicious players

Definition 6 (malicious player) A congestion game with a malicious player is specified by a congestion game \mathcal{G} together with a real number $w(\mathcal{G})$ satisfying $0 \leq w(\mathcal{G}) \leq v(\mathcal{G})$. We interpret $w(\mathcal{G})$ as the amount of flow controlled by the malicious player.

When the malicious player routes its flow using a particular (possibly randomized) flow g , the remaining flow (controlled by the rational players) is, in

effect, participating in a modified congestion game whose latency functions have been changed to reflect the load imposed by the malicious player. We now define this notion precisely.

Definition 7 (induced game) Let $\mathcal{G} = (E, \vec{\ell}, \Pi, v)$ be a congestion game with a malicious player, $w = w(\mathcal{G})$, and γ a probability measure on $F(\mathcal{G}, w)$. The induced latency function ℓ_e^γ on an edge e is defined by

$$\ell_e^\gamma(x) = \mathbf{E}(\ell_e(x + x_e(g))),$$

where g is a random sample from the distribution γ . The induced game \mathcal{G}^γ is the congestion game $(E, \vec{\ell}^\gamma, \Pi, v - w)$. If f is a flow in \mathcal{G} , the induced cost $C^\gamma(f)$ is the cost of f in the induced game \mathcal{G}^γ . When γ is a point mass concentrated on a single flow $g \in F(\mathcal{G}, w)$, we will use the notation \mathcal{G}^g (resp. C^g, ℓ_e^g) to mean the same thing as \mathcal{G}^γ (resp. C^γ, ℓ_e^γ).

Definition 8 (malicious best response) If f is a flow in \mathcal{G} , the set of malicious best responses to f is the set

$$MBR(f) = \arg \max_{g \in F(\mathcal{G}, w)} C^g(f).$$

A probability measure on $F(\mathcal{G}, w)$ is a malicious best response to f if it is supported on the set $MBR(f)$.

The definition requires the malicious player to send its entire flow of size w . Note that for any flow f , as the latency function are monotonically non-decreasing, the more flow the malicious player sends the higher his utility. Thus, even if we were to allow the malicious player to send less flow (or even no flow), he would always choose to send the entire amount of flow he controls.

We are now in a position to define the equilibria of a congestion game with a malicious player. Intuitively, a pair of flows (f, g) — with f representing the rational players and g representing the malicious player — is an equilibrium if none of the rational players can unilaterally improve their delay by switching to a different path, and if the malicious player can not inflict greater damage on the rational players by shifting from g to some other flow. In order to guarantee the existence of equilibria, it is necessary to allow the malicious player to use a mixed strategy, i.e. to sample a random flow from $F(\mathcal{G}, w)$. Thus an equilibrium is actually a pair (f, γ) where f is a flow of value $v - w$ and γ is a distribution on the set of flows of value w .

Definition 9 (equilibrium) If \mathcal{G} is a congestion game with a malicious player and $w = w(\mathcal{G})$ is the amount of malicious flow, then an equilibrium of \mathcal{G} is an ordered pair (f, γ) such that f is a Nash flow in the induced game \mathcal{G}^γ , and γ is a malicious best response to f . An equilibrium is pure if γ is a point mass concentrated on a single flow $g \in F(\mathcal{G}, w)$.

Definition 10 (Nash delay) Let \mathcal{G} be a congestion game with a malicious player, $w = w(\mathcal{G})$, and \mathcal{E} the set of equilibria of \mathcal{G} . The Nash delay $D(\mathcal{G}, w)$

is defined to be the supremum of the set $\{C^\gamma(f)/v(f) \mid (f, \gamma) \in \mathcal{E}\}$.

Note that, in order for $D(\mathcal{G}, w)$ to be well-defined, it must be the case that the set of equilibria of \mathcal{G} (with w units of malicious flow) is nonempty. We will see in Section 4 that this is indeed the case.

For a congestion game \mathcal{G} with flow value $v = v(\mathcal{G})$, the price of malice measures the rate at which the Nash delay deteriorates as a small fraction of the flow comes under the control of a malicious player.

Definition 11 (price of malice) *The price of malice, $\text{POM}(\mathcal{G})$, is defined by*

$$\text{POM}(\mathcal{G}) = \lim_{\varepsilon \rightarrow 0^+} \frac{D(\mathcal{G}, \varepsilon v) - D(\mathcal{G})}{\varepsilon D(\mathcal{G})} = \frac{1}{D(\mathcal{G})} \cdot \frac{d}{d\varepsilon} (D(\mathcal{G}, v\varepsilon))_{\varepsilon=0} \quad (1)$$

when the limit exists.⁴

Note that the price of malice quantifies the first order effect of a small fraction of malicious flow. Clearly one can change the definition to capture lower order effects (for example an $O(\varepsilon^2)$ increase in relative delay).

A counterintuitive phenomenon which we will explore later in this paper is the *windfall of malice*, whereby the presence of a malicious player in the game actually improves the delay experienced by the rational players. We say that a game exhibits windfall of malice if it has a negative price of malice.

2.2 A differential criterion for equilibrium

It is useful to relate the definition of a malicious best response given above (Definition 8) to a criterion which is based on the derivatives of the latency functions, and which says that the malicious player's flow should be distributed on paths which maximize the marginal cost (to the rational players) per unit of flow. Throughout this section we assume that \mathcal{G} is a congestion game with differentiable latency functions.

Definition 12 (differential MBR) *Let \mathcal{G} be a congestion game with differentiable latency functions. Consider any two flows $f, g \in F(\mathcal{G})$. We say that g is a differential malicious best response (DMBR) to f if for every two paths $P, P' \in \Pi$ such that $g(P) > 0$, we have*

$$\sum_{e \in P} x_e(f) \ell'_e(x_e(f) + x_e(g)) \geq \sum_{e \in P'} x_e(f) \ell'_e(x_e(f) + x_e(g)).$$

This definition is closely related to the definition of malicious best response implied by equation (9) in (Karakostas and Viglas, 2003). Indeed, we will see that being a DMBR to f is always a necessary condition for being a malicious best response to f , and that when the latency functions are concave it is also

⁴ We look at the right side derivative as the flow value must be non-negative.

a sufficient condition. Thus our definition of malicious best response (hence also our definition of equilibrium) is equivalent to the definition given by Karakostas and Viglas (2003) in the special case when latency functions are concave.

Lemma 2 *Every malicious best response to f is a DMBR to f .*

Proof: Let g be a malicious best response to f , and let $P, P' \in \Pi$ be two paths such that $g(P) > 0$. For $t \geq 0$, consider the flow $g^{(t)}$ defined by

$$g^{(t)}(Q) = \begin{cases} g(Q) & \text{if } Q \neq P, P' \\ g(Q) - t & \text{if } Q = P \\ g(Q) + t & \text{if } Q = P' \end{cases}$$

If $w = v(g)$ then $g^{(t)} \in F(\mathcal{G}, w)$ for $t \in [0, g(P)]$. Since $g \in \arg \max_{h \in F(\mathcal{G}, w)} C^h(f)$ we have

$$\frac{d}{dt} (C^{g^{(t)}}(f))_{t=0} \leq 0.$$

The left side is equal to $\sum_{e \in P'} x_e(f) \ell'_e(x_e(f) + x_e(g)) - \sum_{e \in P} x_e(f) \ell'_e(x_e(f) + x_e(g))$. \square

Lemma 3 *If g is a DMBR to f , then for every flow h of value $v(g)$,*

$$\sum_{e \in E} x_e(f) \ell'_e(x_e(f) + x_e(g)) [x_e(g) - x_e(h)] \geq 0. \quad (2)$$

Proof: For any path P , define $B(P)$ to be the sum

$$B(P) = \sum_{e \in P} x_e(f) \ell'_e(x_e(f) + x_e(g)).$$

The left side of (2) is equal to $\sum_{P \in \Pi} [g(P) - h(P)] B(P)$. Hence (2) is equivalent to

$$\sum_{P \in \Pi} g(P) B(P) \geq \sum_{P \in \Pi} h(P) B(P). \quad (3)$$

If $M = \max_{P \in \Pi} B(P)$ then by the definition of a DMBR, we have $B(P) = M$ for every path P such that $g(P) > 0$; hence the left side of (3) is equal to $v(g) \cdot M$. Similarly the right side is bounded above by $v(g) \cdot M$. \square

Theorem 4 *Assume that \mathcal{G} is a congestion game with malicious players and for every edge e , ℓ_e is a differentiable, weakly concave function. Then a flow g is a DMBR to f if and only if g is a malicious best response to f .*

Proof: By Lemma 2 every malicious best response to f is a DMBR to f , so we are left to show that every DMBR to f is a malicious best response to f .

For an edge e , let λ_e be the function

$$\lambda_e(x) = \ell_e(x_e(f) + x_e(g)) + \ell'_e(x_e(f) + x_e(g))(x - x_e(f) - x_e(g)).$$

This is a linear function of x which satisfies

$$\begin{aligned}\lambda_e(x_e(f) + x_e(g)) &= \ell_e(x_e(f) + x_e(g)) \\ \lambda'_e(x_e(f) + x_e(g)) &= \ell'_e(x_e(f) + x_e(g)).\end{aligned}$$

Since ℓ_e is concave and λ_e is a linear function whose value and first derivative agree with those of ℓ_e at $x_e(f) + x_e(g)$, we may conclude that $\lambda_e(x) \geq \ell_e(x)$ for all x . Now suppose that g is a DMBR to f , and h is any flow of value $v(g)$. By Lemma 3 we have

$$\begin{aligned}\sum_{e \in E} x_e(f) \ell'_e(x_e(f) + x_e(g)) [x_e(h) - x_e(g)] &\leq 0 \\ \sum_{e \in E} x_e(f) [\lambda_e(x_e(f) + x_e(h)) - \lambda_e(x_e(f) + x_e(g))] &\leq 0 \\ \sum_{e \in E} x_e(f) \lambda_e(x_e(f) + x_e(h)) &\leq \sum_{e \in E} x_e(f) \lambda_e(x_e(f) + x_e(g)).\end{aligned}$$

Now using the fact that $\lambda_e(x) \geq \ell_e(x)$ for all x , with equality when $x = x_e(f) + x_e(g)$, we obtain

$$\sum_{e \in E} x_e(f) \ell_e(x_e(f) + x_e(h)) \leq \sum_{e \in E} x_e(f) \ell_e(x_e(f) + x_e(g)).$$

As h was an arbitrary flow of value $v(g)$, this confirms that g is a malicious best response to f . \square

Our definition of malicious best response may be regarded as the appropriate definition for modeling a single (myopically) malicious player controlling w units of flow, while the definition of differential malicious best response models a continuum of infinitesimal malicious players, collectively controlling w units of flow. (Definition 12 is tantamount to asserting that one cannot increase $C^g(f)$ by rerouting an infinitesimal amount of flow.) Since all malicious players experience the same payoff, it is plausible that w units of flow controlled by a continuum of such players will behave identically to the same amount of flow controlled by a single malicious player. Indeed, Lemma 4 shows that this is exactly what happens when the latency functions are concave. Interestingly, this is not what happens in general when latency functions can be non-concave (see Example 1). The reason is that a single malicious player has the power to play a mixed strategy, while this can never happen with a continuum of malicious players unless we allow them to correlate their random choices. (Even if each of the infinitesimal players uses a mixed strategy, if their random choices are independent then the law of large numbers ensures that their combined flow is equal to a single element of $F(\mathcal{G}, w)$ with probability 1.)

We next show that in any equilibrium with small enough malicious flow, the malicious flow uses only paths that maximize the per-unit cost at the Nash equilibrium.

Definition 13 (differentially malicious flow) *Let \mathcal{G} be a congestion game with differentiable latency functions. Consider a Nash flow $f_N \in F(\mathcal{G})$. A path $P^* \in \Pi$ is a differentially malicious path w.r.t. f_N if for every $P \in \Pi$*

$$\sum_{e \in P^*} x_e(f_N) \ell'_e(x_e(f_N)) \geq \sum_{e \in P} x_e(f_N) \ell'_e(x_e(f_N))$$

A flow g is differentially malicious w.r.t. f_N if for any $P \in \Pi$ such that $g(P) > 0$, P is a differentially malicious path w.r.t. f_N .

Proposition 5 *Let \mathcal{G} be a congestion game with a malicious player, with continuously differentiable latency functions. For any Nash flow $f_N \in F(\mathcal{G})$ there is an open set $U \subseteq F(\mathcal{G})$ containing f_N , such that for any $f \in U$, it holds that every $g \in \text{MBR}(f)$ (of value $v(f_N) - v(f)$) is also differentially malicious w.r.t. f_N .*

Proof: For any path $P \in \Pi$, define a two-variable function $h_P : F(\mathcal{G}) \times F(\mathcal{G}) \rightarrow \mathbb{R}$ as follows:

$$h_P(f, g) = \sum_{e \in P} x_e(f) \ell'_e(x_e(f) + x_e(g)).$$

Observe that h_P is a continuous function because ℓ_e is continuously differentiable for every edge e . Observe also that g is a DMBR to f if and only if $\{P : g(P) > 0\} \subseteq \arg \max_{P \in \Pi} h_P(f, g)$.

Now let Π_{mal} denote the set of all paths which are differentially malicious w.r.t. f_N , i.e. $\Pi_{mal} = \arg \max_{P \in \Pi} h_P(f_N, 0)$. If $\Pi_{mal} = \Pi$ then every path is differentially malicious and the proposition follows trivially. Otherwise, the two-variable function $h(f, g)$ defined by

$$h(f, g) = \min\{h_{P^*}(f, g) - h_P(f, g) \mid P^* \in \Pi_{mal}, P \notin \Pi_{mal}\}$$

is continuous and satisfies $h(f_N, 0) > 0$, by the definition of Π_{mal} . Therefore the set $W = \{(f, g) : h(f, g) > 0\}$ is an open neighborhood of $(f_N, 0)$ in $F(\mathcal{G}) \times F(\mathcal{G})$. Let $W_1 \times W_2$ be an open subset of W such that $f_N \in W_1$, $0 \in W_2$. Without loss of generality (replacing W_1, W_2 with smaller open neighborhoods of $f_N, 0$ if necessary) we may assume that for some real number $\delta > 0$,

$$\begin{aligned} W_1 &\subseteq \{f \in F(\mathcal{G}) \mid v(f) > v(\mathcal{G}) - \delta\} \\ W_2 &= \{g \in F(\mathcal{G}) \mid v(g) < \delta\} \end{aligned}$$

We claim that $U = W_1$ satisfies the conclusion of the proposition. For any $f \in W_1$, if g is a malicious best response to f of value $v(\mathcal{G}) - v(f)$, then $v(g) < \delta$ hence $g \in W_2$. Thus $(f, g) \in W_1 \times W_2 \subseteq W$, which implies $h(f, g) > 0$. By the definition of $h(f, g)$, this implies that $\arg \max_{P \in \Pi} h_P(f, g) \subseteq \Pi_{mal}$. Recalling that g is a malicious best response (and hence, by Lemma 2, a DMBR) to f , we see that every path P with $g(P) > 0$ is an element of $\arg \max_{P \in \Pi} h_P(f, g)$, hence every such P belongs to Π_{mal} . \square

2.3 Lower bound on the price of malice

In this section we construct network congestion games with a large price of malice. Intuitively, the price of malice can be large for at least two reasons:

- (1) The network contains some edges whose latency functions grow very rapidly, so that a small amount of additional flow can have a very large impact on the delay.
- (2) The network contains a very long path, so that the malicious player can send its flow on this path and thereby influence many of the paths being used by the rational players.

We capture the first property using the notion of *relative slope* of the latency functions, which has also been used elsewhere in the literature on selfish routing, e.g. (Fischer et al., 2006). We capture the second property by building a congestion game with a unique equilibrium, namely a pure equilibrium in which the rational players use many disjoint short paths and the malicious player uses a single long path that intersects all of the short paths.

Definition 14 *Let $\ell : [0, 1] \rightarrow \mathbb{R}_+$ be a continuous non-decreasing function that is continuously differentiable. The relative slope of ℓ is defined to be the number*

$$d = \sup_{x \in [0, 1]} \frac{x\ell'(x)}{\ell(x)}.$$

Note that the relative slope of ℓ is actually the maximal value of the elasticity of ℓ in its domain $[0, 1]$.

Theorem 6 *Let $\ell : [0, 1] \rightarrow \mathbb{R}_+$ be a continuous non-decreasing function that is continuously differentiable, and let d be the relative slope of ℓ . For any m there exists a network congestion game with $O(m)$ edges, such that the latency function of each edge is either ℓ or 0, and such that the price of malice is $d(m - 1)$.*

Proof: The continuous function $\frac{x\ell'(x)}{\ell(x)}$ achieves its supremum, d , at some point of the interval $[0, 1]$ because $[0, 1]$ is compact. Let X_0 be a point where the supremum is achieved. The network is illustrated in Figure 1(a) for $m = 5$. In this network congestion game the flow value is mX_0 . The network has m parallel paths of length 3, all have the same latency function $\ell(x)$ on the middle edge. All other edges have a constant latency 0. Backward edges enable the malicious flow to travel all the edges with non-zero latency functions (a path of length $2m + 1$).

Figure 1(b) illustrates the path that the malicious flow of size εmX_0 takes. As the latency functions are the same on every one of the m paths, in equilibrium the rational flow of size $(1 - \varepsilon)mX_0$ will be split equally on the m paths, thus in equilibrium on each path there is a rational flow of size $(1 - \varepsilon)X_0$ (in case that $\varepsilon = 0$ this means a flow of X_0). This is illustrated




Fig. 1. A network congestion game with a large price of malice. (a) The congestion game. (b) The malicious player’s equilibrium strategy. (c) The rational players’ equilibrium strategy.

in Figure 1(c). The total flow on each of the middle edges of the m paths is $(1 - \varepsilon)X_0 + \varepsilon m X_0 = X_0 + \varepsilon X_0(m - 1)$.

The latency with ε units of malicious flow is $\ell(X_0 + \varepsilon X_0(m - 1))$, and the latency with no malicious flow is $\ell(X_0)$. Using the fact that $\lim_{\varepsilon \rightarrow 0} \frac{\ell(x + a\varepsilon) - \ell(x)}{\varepsilon} = a \cdot \ell'(x)$, for $a = X_0(m - 1)$ we obtain that the price of malice is

$$\begin{aligned} \text{POM}(\mathcal{G}) &= \lim_{\varepsilon \rightarrow 0} \frac{\ell(X_0 + \varepsilon X_0(m - 1)) - \ell(X_0)}{\varepsilon \ell(X_0)} \\ &= \frac{X_0(m - 1)\ell'(X_0)}{\ell(X_0)} = (m - 1) \cdot d \end{aligned}$$

□

3 The Windfall of Malice

In this section we show that there exists a network with “windfall of malice”, that is, replacing some of the rational flow with malicious flow causes the delay of the rational flow to *decrease*. Moreover, for this network the windfall is unbounded although the network has bounded size. It increases linearly with the relative slope of the latency functions. At first look it seems surprising that there can be a *decrease* in the latency experienced by the rational agents, as the malicious agent is trying to *maximize* the latency of the rational agents. But

this phenomenon is not too different from the well-known Braess' paradox, which gives an example of a network for which an *increase* in the latency function on an edge *improves* the Nash delay. While in the Braess' paradox network there is no windfall of malice, we are able to construct a network that is based on that network that does have a windfall. In the network that we construct, the malicious agent, by myopically trying to do as much harm as possible, increases the latency on every possible edge, and by doing so it causes the rational agents to take alternative routes that are less harmful to the other rational agents.

Proposition 7 *There exists a network congestion game for which the price of malice is negative. Moreover, for any d there is a constant size network congestion game with price of malice $-d/9$, in which the latency functions are homogeneous polynomials of degree d .*

Proof: We construct a network congestion game \mathcal{G} with flow value 1 and network with source s and target t as presented in Figure 2(a). The latency function on each edge is presented in the graph. (Some of the edges have a constant latency of either 0 or 1, and we just write the constant near the appropriate edge).

Figure 2(b) presents the path that the malicious flow of value ε takes (the path $(s u m n d t)$). As for this path the malicious flow goes on every edge with non-constant latency, it is clear that this flow is always a malicious best response, independent of the rational flow. This implies that there is a unique Nash delay in the induced game.

Figure 2(c) presents the rational flow. Adjacent to each edge we mark the value of flow on the edge, this value is a function of ε , the size of the malicious flow. The symmetry in the induced latency functions ensures that the rational flow at any equilibrium must be symmetric as shown in the figure (e.g. the flow must be the same for the edge $(s u)$ and the edge $(d t)$). Interestingly we are able to calculate the price of malice without explicitly calculating the equilibrium flow when $\varepsilon > 0$. We first note that without any malicious flow ($\varepsilon = 0$) the unique Nash flow is $a(0) = 1, b(0) = 1/2$. We next move to consider the case that $\varepsilon > 0$.

Flow conservation implies that

$$2a(\varepsilon) - 2b(\varepsilon) = 1 - \varepsilon \tag{4}$$

For $\varepsilon > 0$ it must be the case that there is a positive flow on the edge $(s d)$ ($a(\varepsilon) - 2b(\varepsilon) > 0$) as if $a(\varepsilon) - 2b(\varepsilon) = 0$ this implies that $a(\varepsilon) = 1 - \varepsilon$ and the rational flow on the sub-path $(s u m d)$ or the sub-path $(s u n d)$ has delay larger than 1, while the edge $(s d)$ has a delay of 1, which is a contradiction to the flow being an equilibrium. This implies that the delay on the sub-path $(s u m d)$ (and the sub-path $(s u n d)$) must be equal 1, or equivalently:

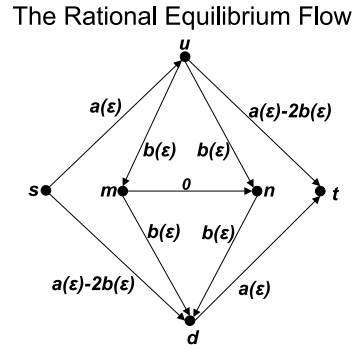
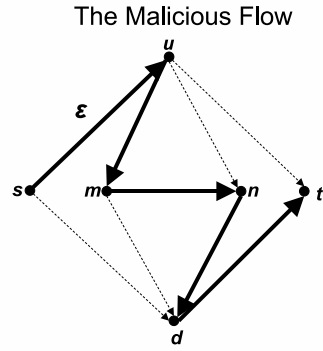
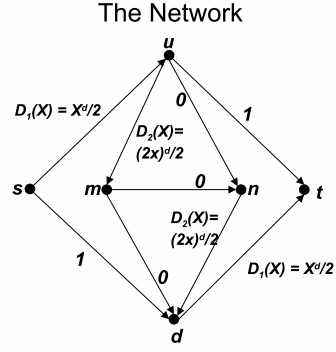


Fig. 2. A network congestion game with a negative price of malice. (a) The congestion game. (b) The malicious player's equilibrium strategy. (c) The rational players' equilibrium strategy.

$$D_1(a(\varepsilon) + \varepsilon) + D_2(b(\varepsilon) + \varepsilon) = 1 \tag{5}$$

We denote $y(\varepsilon) = a(\varepsilon) + \varepsilon$. Equation 4 implies that $b(\varepsilon) = y(\varepsilon) - 1/2 - \varepsilon/2$. Now Equation 5 is equivalent to

$$D_1(y(\varepsilon)) + D_2(y(\varepsilon) - 1/2 + \varepsilon/2) = 1$$

Taking the derivative with respect to ε we derive that

$$y'(\varepsilon) \cdot D'_1(y(\varepsilon)) + (y'(\varepsilon) + 1/2) \cdot D'_2(y(\varepsilon) - 1/2 + \varepsilon/2) = 0$$

For $\varepsilon = 0$ we conclude that

$$y'(0) \cdot D'_1(y(0)) + (y'(0) + 1/2) \cdot D'_2(y(0) - 1/2) = 0$$

Solving for $y'(0)$ we get

$$y'(0) = -\frac{1}{2} \cdot \frac{D'_2(y(0) - 1/2)}{D'_1(y(0)) + D'_2(y(0) - 1/2)} \quad (6)$$

The equilibrium delay in this game is $D(\mathcal{G}, \varepsilon) = 1 + D_1(y(\varepsilon))$ and it holds that $\frac{d}{d\varepsilon}(D(\mathcal{G}, \varepsilon))_{\varepsilon=0} = (y'(\varepsilon)D'_1(y(\varepsilon)))_{\varepsilon=0} = y'(0)D'_1(y(0))$

We are now in a position to calculate the price of malice for this game.

$$\text{POM}(\mathcal{G}) = \frac{1}{D(\mathcal{G})} \cdot \frac{d}{d\varepsilon}(D(\mathcal{G}, \varepsilon))_{\varepsilon=0} = \frac{y'(0)D'_1(y(0))}{1 + D_1(y(0))} \quad (7)$$

Recall that $y(\varepsilon) = a(\varepsilon) + \varepsilon$ and that $a(0) = 1$, thus $y(0) = 1$. Additionally, $D_1(x) = x^d/2$ and $D_2(x) = (2x)^d/2$ thus $D'_1(x) = d \cdot x^{d-1}/2$ and $D'_2(x) = d \cdot (2x)^{d-1}$. This means that $D'_1(y(0)) = D'_1(1) = d/2$ and $D'_2(y(0) - 1/2) = D'_2(1/2) = d$, which can be used in Equation 6 to derive that

$$y'(0) = -\frac{1}{2} \cdot \frac{d}{d/2 + d} = -\frac{1}{3}$$

As $D'_1(y(0)) = d/2$ and $D_1(y(0)) = 1/2$, by Equation 7 we derive that

$$\text{POM}(\mathcal{G}) = \frac{y'(0)D'_1(y(0))}{1 + D_1(y(0))} = \frac{-1/3 \cdot d/2}{1 + 1/2} = -d/9$$

and this concludes the proof of the proposition. \square

We note that although from a qualitative point of view the windfall of malice is tightly related to Braess' Paradox, the two are very different quantitatively. While Roughgarden (2001) has shown that the severity of Braess' Paradox (i.e., the ratio of Nash cost in a network to Nash cost in any sub-network) is bounded by a constant depending only on the network size (i.e., independent of the relative slope of the latency functions) we have shown that the windfall of malice can grow unboundedly large in a constant size graph, if the latency functions are polynomials of unbounded degree.

4 The Existence of Equilibria

Congestion games *without* malicious players have pure Nash equilibria because they are potential games: the potential function Φ defined in Definition 5 decreases whenever a player shifts from one path to another one with lower delay, hence any flow which minimizes Φ must be a pure Nash equilibrium of the congestion game. But congestion games *with* malicious players are not potential games, and as such there is no guarantee that they will have pure Nash equilibria. In fact, a simple example illustrates that a pure Nash equilibrium may not exist even for network congestion games played on a pair of parallel links with continuous latency functions.

Example 1 Consider a network congestion game in a graph consisting of a source and sink joined by two parallel edges e, e' whose latency functions are $\ell_e(x) = \ell_{e'}(x) = x^2$. Let $v = 2$ and $w = 1$, so that the rational players control 1 unit of flow and the malicious player also controls 1 unit of flow. We claim that this game has no pure Nash equilibrium. To prove it, assume by contradiction that (f, g) is a pair of flows constituting a Nash equilibrium. Let $a = f(\{e\}), b = g(\{e\})$. Then $f(\{e'\}) = 1 - a$ and $g(\{e'\}) = 1 - b$, and

$$C^g(f) = a(a + b)^2 + (1 - a)(2 - a - b)^2. \quad (8)$$

One consequence of (8) is that $C^g(f)$ is a strictly convex function of the parameter b , so its maximum is achieved when $b = 0$ or $b = 1$ (or both) but not when $0 < b < 1$. Since we are assuming g is a malicious best response to f , it must be the case that $b = 0$ or $b = 1$. Assume without loss of generality that $b = 0$. Then the induced game \mathcal{G}^g has latency functions $\ell_e^g(x) = x^2, \ell_{e'}^g(x) = (1 + x)^2$. Since we are assuming f is a Nash flow for \mathcal{G}^g , we find that $a = 1$. But then the malicious best response to f is $b = 1$, contradicting our earlier assumption that $b = 0$.

At an intuitive level, the reason why the game constructed in this example has no pure Nash equilibrium is similar to the reason why there is no pure Nash equilibrium in the game “matching pennies”. The strict convexity of the latency functions gives the malicious player an incentive to make the load on e, e' as unbalanced as possible, while the rational players have an incentive to make the load on e, e' as balanced as possible; no distribution of flow can simultaneously satisfy the objectives of both types of players.

In light of Example 1, we devote the rest of this section to proving two theorems: first, congestion games with malicious players have pure Nash equilibria as long as the latency functions are continuous and *weakly concave*⁵; second, congestion games with malicious players always have equilibria in the sense of Definition 9.

⁵ In particular, as a special case, pure Nash equilibria always exist when the latency functions are linear.

Theorem 8 *If \mathcal{G} is a congestion game with a malicious player, and for every edge e , ℓ_e is a continuous, weakly concave function, then there exists a pure equilibrium of \mathcal{G} .*

Given Theorem 4, which ensures that our definition of equilibrium is equivalent to the Karakostas-Viglas definition in the case of concave latency functions, it is possible to deduce this theorem from Theorem 1 of (Karakostas and Viglas, 2003). (Actually, our Theorem 8 makes slightly weaker hypotheses about the latency functions, but the proof technique used by Karakostas and Viglas (2003) implies our theorem without much difficulty.) In the interest of making this paper self-contained, we present a simple alternative proof below.

Proof: Let $w = w(\mathcal{G})$. For a flow $g \in F(\mathcal{G}, w)$, let Φ^g denote the potential function of the game \mathcal{G}^g . Recall from Proposition 1 that a flow f is a Nash flow of \mathcal{G}^g if and only if f is a minimizer of Φ^g . Thus a pair $(f, g) \in F(\mathcal{G}, v - w) \times F(\mathcal{G}, w)$ is a pure equilibrium of \mathcal{G} if and only if f is a minimizer of $\Phi^g(f)$ and g is a maximizer of $C^g(f)$. In other words, a pure equilibrium of \mathcal{G} is equivalent to a pure equilibrium of the two-player normal form in which the strategy sets of the two players are $F(\mathcal{G}, v - w)$ and $F(\mathcal{G}, w)$, respectively, and their payoff functions are $-\Phi^g(f)$ and $C^g(f)$, respectively.

Now let us recall the following easy consequence of Kakutani's Fixed Point Theorem. (See, for example, Proposition 20.3 of (Osborne and Rubinstein, 1994).)

Proposition 9 *A normal form game with finitely many players has a pure Nash equilibrium provided that*

- *Each player's strategy set is a nonempty compact convex subset of a Euclidean space.*
- *For each i , the payoff function of player i is continuous and is a weakly concave function of player i 's strategy.*

The first condition is satisfied because the sets $F(\mathcal{G}, v - w), F(\mathcal{G}, w)$ are nonempty convex polytopes. To verify the second condition, first recall that Φ^g is a continuous and weakly convex function, so $-\Phi^g$ is continuous and weakly concave. Finally, recall that

$$C^g(f) = \sum_{e \in E} x_e(f) \ell_e(x_e(f) + x_e(g)).$$

The function $x_e(g)$ is a linear function of g , and ℓ_e is continuous and weakly concave, so $\ell_e(x_e(f) + x_e(g))$ is a continuous, weakly concave function of g . For fixed f , $C^g(f)$ is a non-negative linear combination of such functions, so it is also continuous and weakly concave, as desired. \square

For congestion games with general latency functions, the existence of an equilibrium does not follow easily from any of the known equilibrium exist-

tence theorems for games with a continuum of players, e.g. (Schmeidler, 1973; Mas-Colell, 1984). This is because our definition of congestion games with malicious players combines features of nonatomic games (the rational players) and atomic games (the malicious player), and also because our equilibrium concept partly requires a pure strategy (the rational flow) and partly a mixed strategy (the malicious flow). Accordingly, we devote the remainder of this section to a self-contained proof of the existence of equilibria in congestion games with malicious players.

Proposition 10 *If \mathcal{G} is a congestion game with a malicious player, and all the latency functions ℓ_e are continuous and strictly increasing, then \mathcal{G} has an equilibrium.*

Proof: We use the same two-player game $\tilde{\mathcal{G}}$ introduced in the proof of Theorem 8. The strategy sets $F(\mathcal{G}, v - w)$ and $F(\mathcal{G}, w)$ are compact Hausdorff topological spaces, and the payoff functions $-\Phi^g(f)$ and $C^g(f)$ are continuous, so the existence theorem for mixed Nash equilibria of games with compact Hausdorff strategy sets (Glicksberg, 1952) ensures that there exist Borel probability measures β_0, γ_0 on $F(\mathcal{G}, v - w)$ and $F(w)$, respectively, such that

$$\beta_0 \in \arg \min_{\beta} \Phi^{\gamma_0}(\beta) \tag{9}$$

$$\gamma_0 \in \arg \max_{\gamma} C^{\gamma}(\beta_0). \tag{10}$$

(Here $\Phi^{\gamma}(\beta)$ and $C^{\gamma}(\beta)$ denote the expected values of $\Phi^g(f)$ and $C^g(f)$ when f, g are sampled independently at random from distributions β, γ , respectively.) The only reason that (β_0, γ_0) may not constitute an equilibrium of \mathcal{G} is that our definition of equilibrium requires the rational players to use a pure strategy, not a mixed strategy. In other words, we require the distribution β_0 to be a point mass concentrated at a single flow $f_0 \in F(\mathcal{G}, v - w)$.

Let f_0 denote the flow $f_0(P) = \mathbf{E}_{f \leftarrow \beta_0} [f(P)]$. Our assumption that the latency functions are strictly increasing implies that the function Φ^{γ_0} is strictly convex, so by Jensen's inequality,

$$\Phi^{\gamma_0}(f_0) \leq \Phi^{\gamma_0}(\beta_0), \tag{11}$$

with equality if and only if the distribution β_0 is a point mass concentrated at f_0 . The left and right sides of (11) are in fact equal, by (9). Consequently β_0 is a point mass concentrated at f_0 . By (9) and (10), we may now conclude that (f_0, γ_0) is an equilibrium of \mathcal{G} . \square

Theorem 11 *Every congestion game with a malicious player and continuous latency functions has an equilibrium.*

Proof: We have seen that the theorem holds when the latency functions are strictly increasing, so the idea of the proof is to approximate an arbitrary congestion game $\mathcal{G} = (E, \vec{\ell}, \Pi, v)$ by games with strictly increasing latency

functions. For every positive integer n , let $\ell_e^{(n)}$ denote the latency function $\ell_e^{(n)}(x) = \ell_e(x) + x/n$, and let $\mathcal{G}^{(n)}$ denote the congestion game $(E, \vec{\ell}^{(n)}, \Pi, v)$. Proposition 10 ensures the existence of an equilibrium (f_n, γ_n) for $\mathcal{G}^{(n)}$. We next argue that this sequence of equilibria has a convergent subsequence, under a suitable definition of convergence.

For a separable compact metric space X , we may topologize the set $\Delta(X)$ of Borel probability measures on X using the *weak topology*, in which a sequence μ_1, μ_2, \dots converges to a probability measure μ if and only if $\int f d\mu_n \rightarrow \int f d\mu$ for every bounded continuous function f on X . The space $\Delta(X)$ is compact in the weak topology by Prokhorov's Theorem (Billingsley, 1999). Since both $F(\mathcal{G}, v-w)$ and $F(\mathcal{G}, w)$ are separable compact metric spaces, we conclude that the space $F(\mathcal{G}, v-w) \times \Delta(F(\mathcal{G}, w))$ is compact and therefore the sequence (f_n, γ_n) has a convergent subsequence. Replacing the sequence $\mathcal{G}^{(1)}, \mathcal{G}^{(2)}, \dots$ with a proper subsequence if necessary, we may assume from now on that we have a sequence of games $\mathcal{G}^{(n)} = (E, \vec{\ell}^{(n)}, \Pi, v)$ with equilibria (f_n, γ_n) such that $\ell_e^{(n)}(x) = \ell_e(x) + \alpha_n x$ for some sequence of constants $\alpha_1, \alpha_2, \dots$ converging to zero, and such that the sequence $(f_1, \gamma_1), (f_2, \gamma_2), \dots$ converges to a point $(f, \gamma) \in F(\mathcal{G}, v-w) \times \Delta(F(\mathcal{G}, w))$. We must now prove that (f, γ) is an equilibrium of \mathcal{G} .

In Lemma 13 below, we prove that

$$\Phi^{\gamma_n}(f_n) \rightarrow \Phi^\gamma(f) \tag{12}$$

$$C^{\gamma_n}(f_n) \rightarrow C^\gamma(f). \tag{13}$$

Assuming (12)-(13) for now, consider any $f' \in F(\mathcal{G}, v-w)$ and $\gamma' \in \Delta(F(\mathcal{G}, w))$. The function $g \mapsto \Phi^g(f')$ is a bounded continuous function of $g \in F(\mathcal{G}, w)$; by the definition of weak convergence this implies $\Phi^{\gamma_n}(f') \rightarrow \Phi^{\gamma'}(f')$. Combining this with (12) we obtain

$$\Phi^\gamma(f) - \Phi^{\gamma'}(f') = \lim_{n \rightarrow \infty} (\Phi^{\gamma_n}(f_n) - \Phi^{\gamma_n}(f')) \leq 0,$$

hence f is a best response to γ . The functions $g \mapsto C^g(f_n)$, for $n = 1, 2, \dots$, are a sequence of uniformly bounded measurable functions of $g \in F(\mathcal{G}, w)$, and $\lim_{n \rightarrow \infty} C^g(f_n) = C^g(f)$ for all g . By Lebesgue's dominated convergence theorem, $C^{\gamma'}(f_n) \rightarrow C^{\gamma'}(f)$. Combining this with (13) we obtain

$$C^\gamma(f) - C^{\gamma'}(f) = \lim_{n \rightarrow \infty} (C^{\gamma_n}(f_n) - C^{\gamma'}(f_n)) \geq 0,$$

hence γ is a best response to f . Thus (f, γ) is an equilibrium as claimed. \square

It remains to supply the proof of the step which was omitted in the proof of Theorem 11. This step is established in Lemma 13 below, but we will begin with a technical lemma which aids in the proof of Lemma 13.

Lemma 12 *Let X, Y be compact metric spaces, and let $\Delta(Y)$ denote the space of Borel probability measures on Y , endowed with the weak topology. If $F :$*

$X \times Y \rightarrow \mathbb{R}$ is a continuous function, then the function $F^\circ : X \times \mathbf{\Delta}(Y) \rightarrow \mathbb{R}$ defined by

$$F^\circ(x, \nu) = \int_Y F(x, y) d\nu(y)$$

is continuous.

While Lemma 12 can be derived from more powerful results in measure theory, we include the simple proof here in order to make our exposition more self-contained.

Proof: Let $\{(x_n, \nu_n)\}_{n=1}^\infty$ denote any sequence which converges to a limit point (x, ν) in $X \times \mathbf{\Delta}(Y)$. Since $\mathbf{\Delta}(Y)$ is metrizable (Billingsley, 1999), the space $X \times \mathbf{\Delta}(Y)$ is also metrizable, and to verify continuity of F° it suffices to check that $F^\circ(x_n, \nu_n) \rightarrow F^\circ(x, \nu)$ as $n \rightarrow \infty$. We bound $|F^\circ(x, \nu) - F^\circ(x_n, \nu_n)|$ from above as:

$$\begin{aligned} |F^\circ(x, \nu) - F^\circ(x_n, \nu_n)| &= \left| \int_Y F(x, y) d\nu(y) - \int_Y F(x_n, y) d\nu_n(y) \right| \\ &\leq \left| \int_Y F(x, y) d\nu(y) - \int_Y F(x, y) d\nu_n(y) \right| \\ &\quad + \left| \int_Y F(x, y) d\nu_n(y) - \int_Y F(x_n, y) d\nu_n(y) \right| \\ &\leq \left| \int_Y F(x, y) d\nu(y) - \int_Y F(x, y) d\nu_n(y) \right| \\ &\quad + \int_Y |F(x, y) - F(x_n, y)| d\nu_n(y). \end{aligned} \quad (14)$$

The first term on the right side of (14) converges to zero because $\nu_n \rightarrow \nu$ in the weak topology. The second term converges to zero by Lebesgue's dominated convergence theorem, as the functions $F(x_n, y)$ are uniformly bounded above by a constant function and $F(x, y)$ is the pointwise limit of $F(x_n, y)$. \square

Lemma 13 *If $(f_1, \gamma_1), (f_2, \gamma_2), \dots$ converges to a point (f, γ) in the space $F(\mathcal{G}, v - w) \times \mathbf{\Delta}(F(\mathcal{G}, w))$, then $\Phi^{\gamma_n}(f_n) \rightarrow \Phi^\gamma(f)$ and $C^{\gamma_n}(f_n) \rightarrow C^\gamma(f)$.*

Proof: Define real-valued functions A, B on the set $F(\mathcal{G}, v - w) \times F(\mathcal{G}, w)$ as follows.

$$\begin{aligned} A(h, g) &= \sum_{e \in E} \int_0^{x_e(h)} \ell_e(y + x_e(g)) dy \\ B(h, g) &= \sum_{e \in E} x_e(h) \ell_e(x_e(h) + x_e(g)). \end{aligned}$$

These are continuous functions, and using the notation of Lemma 12 we have

$$\begin{aligned} \Phi^\nu(h) &= A^\circ(h, \nu) \\ C^\nu(h) &= B^\circ(h, \nu) \end{aligned}$$

for every $h \in F(\mathcal{G}, v - w)$ and $\nu \in \Delta(F(\mathcal{G}, w))$. The lemma now follows by a direct application of Lemma 12. \square

5 Conclusions and Open Problems

This paper raises many more questions than it answers. We believe that our definition of malice can be productive in many other contexts; but even if one focuses on congestion games, as we did, there are many open problems to consider.

- We have only derived lower bounds on the price of malice, as well as on its windfall. What are the right upper and lower bounds on the price and windfall of malice, in terms of the max path length and the relative slope (See Definition 14) of the latency functions?
- Given that the presence of malicious players can affect networks in totally different ways, ranging, as we have seen, from disastrous to beneficial, it becomes imperative to understand the circumstances under which these conditions prevail. That is, we are interested in the *characterization* problem of networks for which, say, there is a positive windfall of malice; similarly for a positive price of malice.
- In the same spirit as the characterization problem, it would be equally interesting to be able to determine algorithmically the price of malice for individual networks. This brings up the following suite of problems: Given a network with a fraction of malicious flow, find a semi-pure Nash equilibrium, as guaranteed by Theorem 11. Or, given such a network with weakly concave (or even linear) delays, find a pure Nash equilibrium (Theorem 8). Are these problems PPAD-complete (our proof establishes that they are in the class PPAD (Papadimitriou, 1994)), or is there an alternative algorithmic way of establishing existence? Since uniqueness of equilibria is no longer guaranteed, perhaps the most useful problem is to find an equilibrium with the largest (or smallest) possible price of malice; this problem may even be NP-complete.

6 Acknowledgements

We thank Tim Roughgarden for helpful conversations about this work. We thank Nicolás E. Stier-Moses for pointing us to (Karakostas and Viglas, 2003).

References

Billingsley, P., 1999. Convergence of Probability Measures. John Wiley.

- Brandt, F., Sandholm, T., Shoham, Y., 2007. Spiteful bidding in sealed-bid auctions. In: Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI).
- Eliasz, K., July 2002. Fault tolerant implementation. *Review of Economic Studies* 69 (3), 589–610.
- Fischer, S., Räcke, H., Vöcking, B., 2006. Fast convergence to wardrop equilibria by adaptive sampling methods. In: Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC 2006). pp. 653–662.
- Glicksberg, I. L., 1952. A further generalization of the Kakutani fixed point theorem, with application to Nash equilibrium points. *Proc. American Math. Soc.* 3 (1), 170–174.
- Karakostas, G., Viglas, A., 2003. Equilibria for networks with malicious users. In: Ibaraki, T., Katoh, N., Ono, H. (Eds.), ISAAC. Vol. 2906 of Lecture Notes in Computer Science. Springer, pp. 696–704.
- Mas-Colell, A., December 1984. On a theorem of Schmeidler. *Journal of Mathematical Economics* 13 (3), 201–206.
- Morgan, J., Steiglitz, K., Reis, G., 2003. The spite motive and equilibrium behavior in auctions. *Contributions to Economic Analysis & Policy* 2 (1), 1102–1102.
- Moscibroda, T., Schmid, S., Wattenhofer, R., July 2006. When Selfish Meets Evil: Byzantine Players in a Virus Inoculation Game. In: 25th Annual Symposium on Principles of Distributed Computing (PODC), Denver, Colorado, USA.
- Osborne, M. J., Rubinstein, A., 1994. A course in game theory. MIT Press.
- Papadimitriou, C. H., 1994. On the complexity of the parity argument and other inefficient proofs of existence. *J. Comput. Syst. Sci.* 48 (3), 498–532.
- Rosenthal, R. W., 1973. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2, 65–67.
- Roughgarden, T., 2001. Designing networks for selfish users is hard. In: FOCS. pp. 472–481.
- Roughgarden, T., 2005. *Selfish Routing and the Price of Anarchy*. MIT Press.
- Roughgarden, T., Tardos, E., 2002. How bad is selfish routing? *J. ACM* 49 (2), 236–259.
- Schmeidler, D., April 1973. Equilibrium points of non-atomic games. *Journal of Statistical Physics* 7 (4), 195–200.