

Near Rationality and Competitive Equilibria in Networked Systems

Nicolas Christin Jens Grossklags John Chuang
christin@sims.berkeley.edu jensg@sims.berkeley.edu chuang@sims.berkeley.edu

School of Information Management and Systems
University of California, Berkeley
102 South Hall
Berkeley, CA 94720-4600

ABSTRACT

A growing body of literature in networked systems research relies on game theory and mechanism design to model and address the potential lack of cooperation between self-interested users. Most game-theoretic models applied to system research only describe competitive equilibria in terms of pure Nash equilibria, that is, a situation where the strategy of each user is deterministic, and is her best response to the strategies of all the other users. However, the assumptions necessary for a pure Nash equilibrium to hold may be too stringent for practical systems. Using three case studies on network formation, computer security, and TCP congestion control, we outline the limits of game-theoretic models relying on Nash equilibria, and we argue that considering competitive equilibria of a more general form helps in assessing the accuracy of a game-theoretic model, and can even help in reconciling predictions from game-theoretic models with empirically observed behavior.

Categories and Subject Descriptors

C.2 [Computer Systems Organization]: Computer-Communication Networks

General Terms

Design, Economics, Theory

Keywords

Distributed systems, Game theory, Competitive equilibria, Modeling

1. INTRODUCTION

Empirical evidence of phenomena such as free-riding in peer-to-peer systems [1] or unfairness in ad-hoc networks [19] challenges the traditional system design assumption that all users of a network are able and willing to cooperate for the greater good of the community. Hence, system architects have become increasingly interested

in considering network participants as selfish [30] or competing [29] entities. For instance, in an effort to discourage free-riding, some deployed peer-to-peer systems such as KaZaA or BitTorrent [9] rely on simple incentive mechanisms. More generally, as summarized in [13, 26, 30], a number of recent research efforts have been applying concepts from game theory and mechanism design to networked systems in an effort to align the incentives of each (self-interested) user with the goal of maximizing the overall system performance.

A cornerstone of game theory and mechanism design is the notion of competitive equilibrium, which is used to predict user behavior and infer the outcome of a competitive game. As discussed in [26], the concept of *Nash equilibrium* is predominantly used in system research to characterize user behavior. Assuming each user obtains a utility dependent on the strategy she adopts, a Nash equilibrium is defined as a set of strategies from which no user willing to maximize her own utility has any incentive to deviate [25].

While Nash equilibria are a very powerful tool for predicting outcomes in competitive environments, their application to system design generally relies on a few assumptions, notably, that (1) each participant is infallible (i.e., perfectly rational), and that (2) each user has perfect knowledge of the structure of the game, including strategies available to every other participant and their associated utilities. There seems to be a class of problems for which these assumptions may be too restrictive, for instance, characterizing competitive equilibria in systems where participants have limited knowledge of the state of the rest of the network.

As a practical example of the potential limits of a game theoretical analysis of a networked system solely based on Nash equilibria, one can argue that, in the case of a peer-to-peer file-sharing system that does not provide incentives for users to share, the unique Nash equilibrium leads to the “tragedy of the commons [18],” that is, a situation where users do not share anything to minimize the cost they incur, thereby leading the entire system to collapse. The mere fact that, in practice, some users are sharing files, even in peer-to-peer systems that do not rely on incentive mechanisms, hints that a Nash equilibrium is not actually reached.

The argument that Nash equilibria may be too restrictive to characterize networked systems is not entirely new. In [14], Friedman and Shenker notably argued that equilibria resulting from learning behaviors were a tool better suited for characterizing how network users share network resources than Nash equilibria. In this paper, we take a quite different stand, by advocating to consider competitive equilibria of a slightly more general form than pure Nash equilibria. More precisely, we argue that using simple extensions of pure Nash equilibria (1) helps to assess the robustness of a

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

SIGCOMM'04 Workshops, Aug. 30+Sept. 3, 2004, Portland, Oregon, USA.
Copyright 2004 ACM 1-58113-942-X/04/0008 ...\$5.00.

game-theoretic model to small deviations from expected behavior, thereby providing insight in the accuracy of the model, and (2) may help reconcile empirical observations with analytical modeling.

We illustrate our point by presenting three case studies, on network formation, security, and TCP congestion control, where outcomes predicted by Nash equilibria are not entirely correlated by empirical observations. In each case study, we investigate if and how more general forms of competitive equilibria can be used to better describe observed behavior, or if, on the other hand, the model needs to be refined.

The remainder of this paper is organized as follows. In Section 2, we provide some background by formally discussing the concept of Nash equilibrium and its extensions or potential alternatives. In Section 3, we present our case studies. Finally, in Section 4, we discuss our findings, outline a possible agenda for future research, and draw conclusions from our observations.

2. BACKGROUND

We consider strategic interactions (called *games*) of the following simple form: the individual decision-makers (also called *players*) of a game simultaneously choose actions that are derived from their available strategies. The players will receive payoffs that depend on the combination of the actions chosen by each player.

More precisely, consider a set $N = \{1, \dots, n\}$ of players. Denote as S_i the set of *pure* (i.e., deterministic) strategies available to player i , and denote as s_i an arbitrary member of i 's strategy set. A probability distribution over pure strategies is called a *mixed* strategy σ_i . Accordingly, the set of mixed strategies for each player, Σ_i , contains the set of pure strategies, S_i , as degenerate cases. Each player's randomization is statistically independent of those of the other players. Then, u_i represents player i 's payoff (or *utility*) function: $u_i(\sigma_i, \sigma_{-i})$ is the payoff to player i given her strategy (σ_i) and the other players' strategies (summarized as σ_{-i}). An n -player game can then be described as $G = \{N; \Sigma_i, \Sigma_{-i}; u_i, u_{-i}\}$.

Players are in a Nash equilibrium if a change in strategies by any one of them would lead that player to obtain a lower utility than if she remained with her current strategy [25]. Formally, we can define a Nash equilibrium as follows:

DEFINITION 1. A vector of mixed strategies $\sigma^* = (\sigma_1^*, \dots, \sigma_n^*) \in \Sigma$ comprises a mixed-strategy Nash equilibrium of a game G if, for all $i \in N$ and for all $\sigma_i \in \Sigma_i$, $u_i(\sigma_i, \sigma_{-i}^*) - u_i(\sigma_i^*, \sigma_{-i}^*) \leq 0$.

A pure-strategy Nash equilibrium is a vector of pure strategies, $s^* \in S$, that satisfies the equivalent condition.

The main advantage of the concept of Nash equilibrium resides in its simplicity. However, because Nash equilibria rely on very stringent assumptions on the capabilities and objectives of each player, they can predict counter-intuitive or unrealistic outcomes. Thus, the economics community has provided an increasing number of refinements to strengthen the concept of Nash equilibrium (e.g., perfect vs. proper equilibria). Similarly, some have investigated how to weaken the rational choice assumptions on which the Nash equilibrium concept is built: a rational player is expected to demonstrate error-free decision-making, to have perfect foresight of the game and to be unbounded in her computational abilities. Intuitively, players such as network users (which are not necessarily perfectly rational) or automated agents (which can be faulty, due to software bugs or misconfiguration, or have limited computational resources) will likely deviate from these rigid assumptions.

As an illustration of how the assumptions required in a Nash equilibrium may need to be relaxed in practice, consider, an experienced player whose strategy choice is almost perfectly correlated

with a Nash prediction of a game but always contains a small error. She is playing in an auction with an asymmetry between the expected cost of overshooting and undershooting the Nash solution. If overshooting is less costly, the player's strategy will most likely contain a small upward bias. If a substantial part of the other players shares this marginal bias the outcome of the auction can be surprisingly far away from a Nash prediction [15]. Similarly, in a sealed-bid auction the Nash equilibrium outcome predicts that a player with a lower valuation will only sometimes win the auctioned good. However, this outcome is more likely if players share small imperfections in the execution of their Nash strategies [21].

Such systematic and non-systematic deviations and their outcomes have been motivation to formulate more generalized models of strategic behavior that include the notion of the Nash equilibrium as a special case. In particular, to relax the assumption of perfect rationality required by the concept of Nash equilibrium, some have introduced the concept of *bounded rationality*. Players that are bounded rational are not necessarily picking the best strategy available across the entire decision space, but instead are allowed to make small errors on a number of levels, such as the evaluation of the payoffs associated with a strategy, the assessment of the best available strategy, or the execution of a specific strategy.

Some techniques to model bounded rationality introduce (possibly small) amounts of noise into the decision-making process. As an example, the noisy introspection model [16] relies on (possibly many) layers of speculation on the beliefs about other players' beliefs. When all beliefs are perfectly correct, adding layers of speculation can only converge to a Nash equilibrium. However, the authors of [16] show that, by including some small (noisy) uncertainty in the conjectures about other players beliefs, the game can produce outcomes significantly different from a Nash equilibrium.

Another model of equilibrium with bounded rationality, called Quantal Response Equilibrium, or QRE [23], has been used to characterize equilibria in games where users make errors on the computation of the payoffs associated with a given strategy. Given two strategies and their associated payoffs (u_1, u_2) , with $u_1 > u_2$, an a priori rational player may choose the strategy yielding u_2 if she makes errors $(\varepsilon_1, \varepsilon_2)$ in the computation of the payoffs (u_1, u_2) such that $u_1 + \varepsilon_1 < u_2 + \varepsilon_2$. McFadden showed that, in such a context, one could express the probability of choosing the strategy yielding u_1 as a *power function* $\Pr(\text{choose } 1) = e^{\lambda u_1} / (e^{\lambda u_1} + e^{\lambda u_2})$, where λ is a factor that characterizes the probability of the player making a mistake in the choice of the "right" strategy [22]. Namely, $\lambda = 0$ indicates that the player chooses at random regardless of the payoffs, while $\lambda \rightarrow \infty$ converges to the Nash behavior of always selecting the strategy with the higher payoff. At the Quantal Response Equilibrium, the vector of mixed strategies σ^* used by the players satisfies $\sigma^* = T(u(\sigma^*))$, where $u(\cdot)$ is the function that maps a set of mixed strategies to a set of payoffs, and $T(\cdot)$ is the function that maps a set of payoffs to a set of mixed strategies.¹ Thus, a QRE results in a set of conditions on the power functions.

Models such as QRE or noisy introspection are very useful as an empirical structure for uncovering features of payoffs from field data, or to obtain relationships between observables and primitives of interest [17]. In other words, models and equilibrium concepts for bounded rationality can help reconcile data observed experimentally with a game theoretic analysis.

For the preliminary analysis we present in this paper, we will focus on a particular type of bounded rational behavior, called *near rationality* [3, 27]. In near rational equilibria a player who is not

¹The Brouwer fixed point theorem indicates that such an equilibrium exists as soon as the function $T(e(\cdot))$ is continuous.

perfectly maximizing her utility cannot improve her payoff by a substantial amount by playing her Nash strategy more accurately. While the personal losses for a player are potentially very small, the equilibria derived often represent substantial departures from a prediction based on perfect Nash optimizing behavior. As we will discuss in the remainder of this paper, models of near rationality are appropriate for the description of empirical phenomena but can also contribute explanations and predictions of strategic behavior. In addition, we will also illustrate how models of near rationality can be used to assess the accuracy and robustness of a given game-theoretic model, which can possibly lead to refinements of the model.

3. CASE STUDIES

In this section, we present three case studies on network formation, security, and TCP congestion control. For each of the case studies, we describe the interaction between the different participants in terms of a game, and note the discrepancies between the game outcome as predicted by a Nash equilibrium and the behavior observed empirically. Allowing for near rationality allows us to determine if the outcome is highly sensitive to small variations in player behavior, in which case the model might need to be refined. In cases where the model seems sufficiently robust, we then discuss whether considering near rational players can lead to more accurate predictions.

3.1 Preliminaries

Before we delve into the details of each case study, we discuss in more details the equilibrium concept that we use for the analysis in the remainder of this paper. As we mentioned before, we are interested in assessing if near rationality can help us build better game theoretic models for networked systems. Here, we will focus on a simple, but powerful model of near rationality, called the ε -equilibrium [27].

The ε -equilibrium concept [27] is relaxing the conception of a fully rational player to a model where each player is satisfied to get close to (but does not necessarily achieve) her best response to the other player's strategies. No player can increase her utility by more than ε by choosing another strategy. Therefore, we locate an ε -equilibrium by identifying a strategy for each player so that her payoff is within ε of the maximum possible payoff given the other players' strategies.

Formally, an ε -equilibrium can be defined as follows:

DEFINITION 2. A vector of mixed strategies $\sigma^\varepsilon = (\sigma_1^\varepsilon, \dots, \sigma_n^\varepsilon) \in \Sigma$ comprises a mixed-strategy ε -equilibrium of a game G if, for all $i \in N$, for all $\sigma'_i \in \Sigma_i$, and a fixed $\varepsilon > 0$, $u_i(\sigma'_i, \sigma_{-i}^\varepsilon) - u_i(\sigma_i^\varepsilon, \sigma_{-i}^\varepsilon) \leq \varepsilon$.

A pure-strategy ε -equilibrium is a vector of pure strategies, $s^\varepsilon \in S$, that satisfies the equivalent condition. If we allow $\varepsilon = 0$ this condition reduces to the special case of a Nash equilibrium. Thus, one can consider ε -equilibria as a more generalized solution concept for competitive equilibria.

We emphasize again that other equilibrium concepts, including bounded rationality models, are probably equally, if not more, useful in modeling and analyzing networked systems. However, since the main objective of this paper is to show how equilibrium concepts that are conceptually quite close to Nash equilibria can help improve the game theoretic analysis of networked systems, we defer the discussion of the applicability of more elaborate models of near (or bounded) rationality to future work, and will use ε -equilibria for the analysis in the remainder of this section.

3.2 Network Formation

For our first case study, we briefly discuss network formation by self-interested parties. Following seminal work in economics [20], network formation has lately received relatively significant attention in the networking research community. We refer the interested reader to recent studies, such as [7, 11], for an in-depth discussion of the problem. Here, our only focus is to show how considering near rational behavior, as opposed to Nash behavior, can help us validate or invalidate a game theoretic model.

We define a network as a set of n nodes connected by a set of k directed links (where $k \leq 2n(n-1)$). Each node is used to store items that are of interest to other nodes. We follow the generic network model described in [6] where each node can request items, serve items, or forward requests between other nodes. As in [6], we assume shortest-path routing. Using a few simplifying assumptions (e.g., all nodes are considered to have the same capabilities, all links have the same establishment cost, and requests for items are uniformly distributed over the entire network), the authors of [6] express the cost associated to each node i as

$$C_i = \frac{s}{n} + lEd_{i,j} + rEb_{j,k}(i) + m \deg(i),$$

where $Ed_{i,j}$ is the expected value of the topological distance (hop-count) between node i and another node j , $Eb_{j,k}(i)$ is the expected value of the probability that node i is on the path between two arbitrary nodes j and k , and $\deg(i)$ is the out-degree of node i , that is, the number of nodes node i links to. The constants s , l , r and m represent the nominal costs associated with storing an item, retrieving an item one hop away, routing a request between two other nodes, and maintaining a connection to another node, respectively. From this cost model, we can immediately define the utility of node i , u_i , as

$$u_i = -C_i. \quad (1)$$

Assume that nodes can choose which links they maintain, but do not have any control over the items they hold, and honor all routing requests. In other words, nodes are selfish when it comes to link establishment, but are obedient once links are established.

PROPOSITION 1. *With the utility function given in Eq. (1), if $m < l/n$, the fully connected network where each node links to every other node is a unique pure Nash equilibrium.*

PROPOSITION 2. *If $m > l/n$, the star-shaped network, where all links connect to or from a central node, is a pure Nash equilibrium.*

Propositions 1 and 2, whose proofs are in Appendix A, tell us that, if maintaining links is cheap, or if the network is small, the only Nash equilibrium is the fully connected network. If maintaining links is more expensive, or if the network is large, a star-shaped network is a possible Nash equilibrium.² While the star may not be a unique Nash equilibrium, the high aggregate utility of the star [6] suggests it may dominate other potential Nash equilibria. We note that the authors of [20] obtain comparable results using a slightly different cost model.

Thus, if the model is accurate, and if nodes behave fully rationally, we should expect predominance of fully-connected or star-shaped networks in practice. However, if instead of considering Nash equilibrium, we now consider an ε -equilibrium, then, for any $m \in [l/n - \varepsilon, l/n + \varepsilon]$, any network topology constitutes an ε -equilibrium. (This can be proven by simply including ε in all the

²In the limit case where m is exactly equal to l/n , any network constitutes a Nash equilibrium.

derivations of Appendix A.) Additionally, if, to account for failures in link establishment due for instance to lossy channels, we allow nodes to use mixed strategies instead of being restricted to pure strategies, we conjecture that the range of possible values for m such that any network is an ε -equilibrium is much larger than 2ε .

In other words, even a detailed cost model as proposed in [6] appears to be highly sensitive to small changes in player behavior. This high sensitivity hints that the game model needs to be refined, an intuition which is corroborated by empirical measurements. While star-shaped topologies or full meshes can indeed be found in existing networks (e.g., many small local area networks use star topologies), measurement studies of Internet topologies exhibit much more varied results [12]. Among the reasons why Internet topologies do not solely consist of an interconnection of star-shaped and fully connected networks, one can cite capacity constraints [7] or monetary incentives, which are not included in the cost model described above.

While proposing a game-theoretic model that captures these additional factors is outside of the scope of this paper, we summarize the outcome of this first case study as follows: considering near rationality instead of perfect rationality can help us evaluate the robustness of a model to small perturbations in the players actions. If the model seems to lack robustness, its chances of being an accurate model of reality decrease.

3.3 Protection Against Security Threats

With our second case study, we hope to show that slightly modifying the description of a game to account for near rationality allows to significantly improve the predictive quality of the game when compared to empirically observed behavior. To that effect, we look at the level of security users choose in a network subject to a security threat. Specifically, we focus on protection against potential distributed denial of service (DDoS) attacks. In the first stage of a DDoS attack, an attacker looks for a (set of) machine(s) whose control she can easily seize, to use as a platform to launch an attack of larger magnitude. For instance, by obtaining total control of a machine on a network, an attacker may be able to retrieve passwords and gain access to more secure machines on the same network.

We model here a network of n users, who are all potential targets in the initial stage of a DDoS attack. If we characterize the level of computer security that each user i adopts by a variable s_i , the user(s) with the lowest s_i (i.e., $s_i = s_{\min} = \min_i\{s_i\}$) will be compromised. We assume that each user can infer the security level s_i used by every other user (e.g., by probing), and no finite security level s_i can be selected to guarantee a protection against all attacks. We further assume that the cost of implementing a security policy s_i is a monotonic increasing function of s_i . Specifically, to simplify the notations, we consider here that each user i that is not compromised pays s_i to implement their security policy. The compromised user(s), say user j , pays a fixed penalty $P \geq s_i$ (for any s_i), independent of the security level s_{\min} she has chosen.

While very simplified, we conjecture this game is a relatively accurate model of the first stage of DDoS attacks that have been carried out in practice [10].³ We defer the study of the deployment of the attack beyond the first stage to future work. Also, we stress

³While this type of attack shares some similarities with worm propagation, notably searching for insecure machines [24], a worm typically propagates by infecting all machines on a network that are below a certain, *fixed*, security level, which is different from our hypothesis that only the machines with the lowest level of security are compromised.

that, while the size n of the network does not play a central role in our model, we expect the scenario we describe to be more realistic in the case of a small corporate or university network, that is when n is reasonably small.

PROPOSITION 3. *The game described above has a unique pure Nash equilibrium, where all users choose an identical security level $s_i = P$.*

Proposition 3, whose proof we derive in Appendix B, tells us that, for a Nash equilibrium to hold, all users have to choose the highest level of security available. As in the first case study, one could consider pure ε -equilibria to check the sensitivity of the model to small changes in player behavior. It can be shown that, for this specific game, pure ε -equilibria produce results very close to Proposition 3, and thus the model seems reasonably robust to small perturbations. In fact, due to the symmetry of the game, one can conjecture that, regardless of the value of the penalty P , all users would implement the same security level in a pure equilibrium.

However, available data from large networks, e.g., [8], documents that different systems present highly heterogeneous security vulnerabilities, which in turn indicates that implemented security levels are highly disparate across machines. Hence, in the context of the security game we just described, a Nash equilibrium does not seem to accurately describe observed behavior.

Some of the possible explanations for the heterogeneity of the implemented security levels can be captured by more elaborate equilibrium models. In particular, (1) users have incomplete information on the levels of security deployed by other users, (2) the *perceived* benefit of installing security patches may be smaller than the overhead patching incurs, and (3) some users may be gambling (knowingly or not) on the seriousness of the security threats they face. These three arguments all make the case for considering ε -equilibria (to account for (1) and (2)) with *mixed* strategies (to account for (2) and (3)), rather than a pure Nash equilibrium.

PROPOSITION 4. *There exist mixed-strategy ε -equilibria with $\varepsilon \leq P/4$ where all chosen security levels are distributed over the interval $[0, P]$.*

Proposition 4, which we prove in Appendix B, indicates that considering ε -equilibria with mixed strategies allows us to predict large dispersion of the chosen security levels, even for relatively low values of ε . This result seems to be more in line with the available measurement data. We further note that analogous results have been recently derived to quantitatively model price dispersion phenomena [5], where assuming a Nash equilibrium likewise fails to corroborate empirical measurements.

We emphasize that Proposition 4 only states that near rationality *can* explain dispersion. In particular, the proof of Proposition 4 we present in Appendix B uses a specific distribution function to show that there exist mixed ε -equilibria, with ε relatively small, that result in dispersion of the security levels over the entire spectrum $(0, P)$. However, we do not claim that the specific distribution used in the proof is a realistic characterization of user behavior; in fact, we believe that further work is required to provide a distribution function that accurately describes the mixed strategies implemented by each user.

Additionally, one can direct two critiques at the discussion on the security game we just presented. First, the discrepancies between the behavior predicted by a Nash equilibrium and that observed in practice may be due to an inaccurate game model, rather than from assuming a given type of equilibrium. Second, one can argue that while the assumption of perfect rationality, as required in a pure

Nash equilibrium, is very debatable when strategies are selected by humans (such as in the security game), perfect rationality is a much more reasonable assumption in the case of automated agents. However, neither of these arguments invalidates our case for considering near rationality as an additional tool to improve game theoretic modeling of networked systems.

3.4 TCP Congestion Control

The last case study we propose has for object to further our case that considering near rationality is helpful in refining game-theoretic models for networked systems. This third case study relies on a game-theoretic analysis of the TCP transport protocol [2]. Each TCP sender relies on an additive-increase-multiplicative-decrease (AIMD) algorithm to adjust its sending rate in function of the congestion experienced on the path from sender to receiver.

In [2], Akella et al. present a game-theoretic analysis to model competition between different TCP senders for three of the most popular variants of TCP, namely, TCP Tahoe, TCP Reno and TCP SACK. In the *TCP Game* they describe, players are the TCP sources ($i \in \{1, \dots, n\}$), which are allowed to adjust their individual additive increase (α_i) and multiplicative decrease (β_i) parameters. In the TCP Game, the utility of each player is equal to her goodput, which is defined as the total amount of data transferred over a time interval, minus the amount of data that had to be retransmitted (presumably because of losses in the network) over the same time interval.

One of the insights presented in [2] is that, for TCP SACK, a pure Nash equilibrium results in $\alpha_i \rightarrow \infty$ (infinite additive increase) if β_i is held fixed, while $\beta_i \rightarrow 1$ (no multiplicative decrease) if α_i is held fixed. Simply stated, if all players in a TCP SACK network were behaving according to a Nash equilibrium, they would simply turn off congestion control, which would likely result in the network suffering from complete congestion collapse. However, TCP SACK is increasingly deployed on the Internet [4], and yet, we do not observe congestion collapse phenomena due to misbehaving TCP sources.⁴

One of the possible reasons proposed by the authors of [2] for the continued stable operation of the Internet is that a given user may face technical difficulties to modify the behavior of her machine to behave greedily. We submit this potential explanation can be partially captured by considering an ε -equilibrium instead of a Nash equilibrium. The cost of modifying the behavior of a given machine can indeed be viewed as a switching cost, to be included in the factor ε .

For simplicity, we assume here that players can only modify their additive increase parameter α_i . (An analogous study can be carried out if we allow changes to β_i .) The authors of [2] show that, with TCP SACK, player i 's utility (goodput) is given by

$$u_i(\alpha_i, \alpha_{-i}) = c \frac{\alpha_i}{A + \alpha_i},$$

where c denotes the total capacity (bandwidth-delay product divided by the round-trip-time) of the bottleneck link, and

$$A = \sum_{j \neq i} \alpha_j.$$

⁴In fact, the authors of [2] point out that the Nash equilibria for TCP NewReno and TCP SACK are similar. TCP NewReno and TCP SACK combined currently account for an overwhelming majority of all traffic on the Internet, which hints that the observed stable operation of the Internet probably does not result from having a mix of different TCP variants in the network.

Therefore, having an ε -equilibrium implies that, for any α'_i ,

$$u_i(\alpha'_i, \alpha_{-i}) - u_i(\alpha_i, \alpha_{-i}) \leq \varepsilon,$$

so that

$$c \frac{A(\alpha'_i - \alpha_i)}{(A + \alpha'_i)(A + \alpha_i)} \leq \varepsilon. \quad (2)$$

If we allow $\alpha_i = 0$ and $\alpha'_i \rightarrow \infty$, an ε -equilibrium can only occur for $\varepsilon \geq c$, that is, when ε is larger than the maximum utility achievable. In such a scenario, ε is so large that all players select a value for their parameter α_i at random.

Adding the assumption that variations of α_i are bounded leads to much more interesting results.⁵ Specifically, let us impose $\alpha'_i - \alpha_i \leq K$ for $K \in \mathbb{N}$. For simplicity, let us set the initial values for α_i to the default value in TCP implementations, that is, $\alpha_i = 1$ for all i . Then, we have $A = n - 1$ and $0 \leq \alpha'_i \leq K + 1$. Substituting in Eq. (2), we have a ε -equilibrium as soon as

$$\varepsilon \geq c \frac{K}{n}.$$

Hence, in a network with a large number of TCP senders, the default TCP implementation can be an ε -equilibrium for small values of ε . This is one of the possible explanations why the predicted Nash behavior that all users would turn off TCP congestion control primitives is not observed.

4. DISCUSSION

We have shown through three case studies that considering competitive equilibria of a more general form than pure Nash equilibria can be beneficial in systems research. In particular, we discussed how allowing players to slightly deviate from their optimal utility can help reconcile game-theoretic models and observed player behavior.

The first case study on network formation showed how slightly relaxing the assumption that the players are perfectly rational was useful in assessing the sensitivity of a game theoretic model to small perturbations, thereby helping to evaluate its likelihood to be a realistic model. The second case study on the implementation of security policies showed how the interaction of uncertainties in the decision making process and the payoff evaluation could be used to explain empirical behavior. The third case study on TCP congestion control showed how combining refinements to a model description with a relaxation of the assumptions in force by employing a more generalized equilibrium concept can improve the match between the (analytically) predicted outcome and the (empirically) observed behavior.

From a system design perspective, we note that, even in games in which the actions of each player resulting in a pure Nash equilibrium are undesirable from the system designer's perspective, near rational players may actually settle for a desirable outcome. This is a possible explanation why the Internet does not suffer from congestion collapse, despite the inefficiency of the Nash equilibrium in the TCP SACK game. Conversely, potentially desirable outcomes associated with a Nash equilibrium may prove difficult to reach unless all players are perfectly rational. The security game

⁵Note that there are several possible justifications for bounding the variations on α_i . For instance, because obtaining perfect knowledge of the state of the entire network is difficult (or impossible) for a given user, each user may instead incrementally probe the network to discover her optimal setting for α_i . Such a probing behavior can be captured as a repeated game where, for each repetition, $\alpha'_i - \alpha_i \leq K$. This type of model thus presents some similarities with the learning behavior that is thoroughly studied in [14].

we described presents an instance of such a phenomenon. Thus, it appears that taking into account uncertainty factors can be useful in both game specification and mechanism design.

An alternative to modeling near rationality is to refine the game specification, to capture all factors with any conceivable influence on the game outcome. However, we argue that the two approaches are not exclusive. In fact, as shown in the first case study, refinements to the game description are probably of interest when the near rationality assumption yields substantial deviations from the outcome predicted by a Nash equilibrium. Research on bounded-reasoning and bounded-optimality models [28] provides a solid framework for such refinements.

As a follow-up on our case studies, we are interested in gathering experimental data, through user surveys, on how security levels are chosen in practice, and in investigating how well this data can be described using game-theoretic models. We are also planning on conducting simulation studies to assess the actual impact of uncertainties and of mixed strategies on network formation.

Last, we believe that this research has uncovered a few open problems that may warrant future investigation. First, our case studies seem to show that considering other types of equilibria besides Nash equilibria can help expand the applicability of game-theoretic models to networked systems. While the ε -equilibrium used in this paper is an interesting tool, many other equilibrium models have been investigated in the literature, e.g., [3, 16, 23]. We conjecture that different types of equilibrium may be appropriate for different networking problems, and believe that providing a classification of networking problems according to the specific types of equilibrium that best characterize them would be valuable.

More generally, one can also ask how a game-theoretic model can capture that the rationality of each participant may vary across users: some users may be obedient, some others may be fully rational, some may be faulty [13]. Finding if and how game-theoretical models can accommodate heterogeneous populations of players may help us design better systems, and certainly poses a number of interesting research questions.

5. ACKNOWLEDGMENTS

This work is supported in part by the National Science Foundation through grants ANI-0085879 and ANI-0331659. We thank the anonymous reviewers for their insightful comments, which greatly helped in improving this paper.

6. REFERENCES

- [1] E. Adar and B. Huberman. Free riding on Gnutella. *First Monday*, 5(10), Oct. 2000.
- [2] A. Akella, S. Seshan, R. Karp, S. Shenker, and C. Papadimitriou. Selfish behavior and stability of the Internet: A game-theoretic analysis of TCP. In *Proc. ACM SIGCOMM'02*, pages 117–130, Pittsburgh, PA, Aug. 2002.
- [3] G. Akerlof and J. Yellen. Can small deviations from rationality make significant differences to economic equilibria? *Amer. Econ. Rev.*, 75(4):708–720, Sept. 1985.
- [4] M. Allman. A web server's view of the transport layer. *ACM Comp. Comm. Rev.*, 30(5):10–20, Oct. 2000.
- [5] M. Baye and J. Morgan. Price dispersion in the lab and on the Internet: Theory and evidence. *RAND J. Econ.*, 35(3), Autumn 2004. To appear.
- [6] N. Christin and J. Chuang. On the cost of participating in a peer-to-peer network. In *Proc. IPTPS'04*, San Diego, CA, Feb. 2004.
- [7] B.-G. Chun, R. Fonseca, I. Stoica, and J. Kubiatowicz. Characterizing selfishly constructed overlay networks. In *Proc. IEEE INFOCOM'04*, Hong Kong, Mar. 2004.
- [8] Cisco Secure Consulting. Vulnerability statistics report. http://www.cisco.com/warp/public/778/security/vuln_stats_02-03-00.html.
- [9] B. Cohen. Incentives build robustness in BitTorrent. In *Proc. 1st Workshop on Econ. of Peer-to-Peer Syst.*, Berkeley, CA, June 2003.
- [10] D. Dittrich. The DoS project's "trinoo" distributed denial of service attack tool, Oct. 1999. Available from <http://staff.washington.edu/dittrich/misc/trinoo.analysis>.
- [11] A. Fabrikant, A. Luthra, E. Maneva, C. Papadimitriou, and S. Shenker. On a network creation game. In *Proc. ACM PODC'03*, pages 347–351, Boston, MA, July 2003.
- [12] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the Internet topology. In *Proc. ACM SIGCOMM'99*, pages 251–262, Boston, MA., Aug. 1999.
- [13] J. Feigenbaum and S. Shenker. Distributed algorithmic mechanism design: Recent results and future directions. In *Proc. DIAL-M'02*, pages 1–13, Atlanta, GA, Sept. 2002.
- [14] E. Friedman and S. Shenker. Learning and implementation on the Internet. Working paper. Available from <http://ideas.repec.org/p/rut/rutres/199821.html>, Aug. 1998.
- [15] J. Goeree and C. Holt. Ten little treasures of game theory and ten intuitive contradictions. *Amer. Econ. Rev.*, 91(5):1402–1422, Dec. 2001.
- [16] J. Goeree and C. Holt. A model of noisy introspection. *Games Econ. Behav.*, 46(2):365–382, Feb. 2004.
- [17] P. Haile, A. Hortaçsu, and G. Kosenok. On the empirical content of quantal response equilibrium. *Cowles Foundation Discussion Paper Series*, (1432), Aug. 2003.
- [18] G. Hardin. The tragedy of the commons. *Science*, 162(3859):1243–1248, Dec. 1968.
- [19] H.-Y. Hsieh and R. Sivakumar. Performance comparison of cellular and multi-hop wireless networks: A quantitative study. In *Proc. ACM SIGMETRICS'01*, pages 113–122, Cambridge, MA, June 2001.
- [20] M. Jackson and A. Wolinsky. A strategic model for social and economic networks. *J. Econ. Theory*, 71(1):44–74, Oct. 1996.
- [21] P. Klemperer. Using and abusing economic theory. *J. Europ. Econ. Assoc.*, 1(2/3):272–300, April-May 2003.
- [22] D. McFadden. Econometric analysis of qualitative response models. In Z. Grilliches and M. Intriligator, editors, *Handbook of Econometrics*, volume II, pages 1396–1456. Elsevier, Amsterdam, Netherlands, 1984.
- [23] R. McKelvey and T. Palfrey. Quantal response equilibria for normal form games. *Games Econ. Behav.*, 10(1):6–38, July 1995.
- [24] D. Moore, V. Paxson, S. Savage, C. Shannon, S. Staniford, and N. Weaver. Inside the Slammer worm. *IEEE Security and Privacy*, 1(4):33–39, July 2003.
- [25] J. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, Sept. 1951.
- [26] C. Papadimitriou. Algorithms, games and the Internet. In *Proc. ACM STOC'01*, pages 749–753, Heraklion, Crete, Greece, July 2001.
- [27] R. Radner. Collusive behavior in noncooperative

epsilon-equilibria of oligopolies with long but finite lives. *J. Econ. Theory*, 22:136–154, 1980.

- [28] S. Russell and D. Subramanian. Provably bounded-optimal agents. *J. Artif. Intel. Res.*, 2:575–609, May 1995.
- [29] S. Shenker. Making greed work in networks: A game-theoretic analysis of switch service disciplines. *IEEE/ACM Trans. Networking*, 3(6):819–831, Dec. 1995.
- [30] J. Shneidman and D. Parkes. Rationality and self-interest in peer-to-peer networks. In *Proc. IPTPS'03*, pages 139–148, Berkeley, CA, Feb. 2003.

APPENDIX

A. PROOFS OF PROPOSITIONS 1 AND 2

Here, we first show that the fully connected network is the only Nash equilibrium if and only if $m < l/n$, before showing that, if $m > l/n$, the star-shaped network characterizes a Nash equilibrium.

PROOF OF PROPOSITION 1. In a fully connected network, no node can create additional links. If a given node i removes one of its links, $\deg(i)$ decreases from $(n-1)$ to $(n-2)$, but, at the same time, one of the nodes $i' \neq i$ is now at a distance of 2 from i . Thus, $Ed_{i,j}$ increases from 1 to

$$Ed_{i,j} = \frac{n-1}{n} + \frac{2}{n} = 1 + \frac{1}{n},$$

and the difference in utility for node i , between the strategy of removing one link and the strategy consisting in maintaining all links, is $m - l/n$. To have a pure Nash equilibrium, we therefore need to have $m - l/n \leq 0$, which is true if and only if $m \leq l/n$.

Suppose now that we have $m < l/n$, and a network that is not fully connected. In particular, consider that a node i can decide whether to create a link to another node $i' \neq i$. Before addition of the link $i \rightarrow i'$, i' is at a distance $2 \leq d_{i,i'} \leq n-1$ of i . After creation of the link $i \rightarrow i'$, i' is at a distance 1 of i . Thus, by creating the link $i \rightarrow i'$, $Ed_{i,j}$ at least decreases by $(2-1)/n = 1/n$. Adding the link $i \rightarrow i'$ also results in $\deg(i)$ increasing by one, so that the addition of the link $i \rightarrow i'$ eventually results in a change in the node i 's utility equal to $-m + l/n$, which, by hypothesis, is strictly positive. Hence, node i always has an incentive to add links to nodes it is not connected to. Using the same reasoning for all nodes, we conclude that the fully connected network is the unique Nash equilibrium if $m < l/n$. \square

Consider now a star-shaped network, where all links connect to or from a central node, say node 0, and assume that $m > l/n$.

PROOF OF PROPOSITION 2. Node 0 is fully connected to the rest of the network, and therefore cannot create additional links. If node 0 removes one of its links, one of the $n-1$ other nodes becomes unreachable, which implies $Ed_{0,j} \rightarrow \infty$, and $u_0 \rightarrow -\infty$. Thus, node 0 has no incentive in modifying its set of links. Likewise, peripheral nodes do not remove their (only) link to the central node, to avoid having their utility $u_i \rightarrow -\infty$.

Suppose now that a peripheral node i creates an additional link to another peripheral node $i' \neq i$. An argument identical to that used in the proof of Proposition 1 shows that the addition of the link $i \rightarrow i'$ results in a change in the node i 's utility equal to $-m + l/n$. Here, however, $m > l/n$, so that $-m + l/n < 0$, and node i has no incentive in adding the link $i \rightarrow i'$. Thus, the star-shaped network is a pure Nash equilibrium, which may not be unique. \square

B. PROOFS OF PROPOSITIONS 3 AND 4

We first consider that users are only allowed pure strategies, and prove Proposition 3.

PROOF OF PROPOSITION 3. Without loss of generality, we assume that users $\{1, \dots, k\}$, with $1 \leq k \leq n$, choose a security level $s_{\min} < s_i$ for all $i \in \{k+1, \dots, n\}$. Thus, each user i for $i \in \{1, \dots, k\}$ is compromised, and has a utility $u_i = -P$. Users in $i \in \{k+1, \dots, n\}$ cannot be compromised because $s_i > s_{\min}$ and therefore have a utility $u_i = -s_i$.

Suppose a user i in $\{1, \dots, k\}$ were to increase her security level to $s_i = s_{\min} + h$ for $h > 0$. User i 's utility would become $-s_{\min} - h$. However, because the original constellation of security levels forms a Nash equilibrium, we know that such a change of strategy results in a decrease of user i 's utility for any $h > 0$. That is, for any $h > 0$,

$$-s_{\min} - h + P \leq 0,$$

which reduces to $s_{\min} \geq P - h$ for any $h > 0$, so that $s_{\min} \geq P$ by continuity. By hypothesis, $s_{\min} \leq P$, which implies that $s_{\min} = P$. Since for any i , $s_{\min} \leq s_i \leq P$, we obtain $k = n$, and, for any i , $s_i = P$ is the only possible Nash equilibrium. The utility of each user is $u_i = -P$, and cannot be increased by picking a different security level, which confirms that $s_i = P$ for all i constitutes a Nash equilibrium. \square

Suppose now that users choose their security level probabilistically. More precisely, the probability that user i picks a security level s_i below a value s is characterized by the cumulative distribution function (c.d.f.) $F_{s_i}(s) = Pr[s_i \leq s]$.

PROOF OF PROPOSITION 4. Consider the following continuous c.d.f. $F_{s_i}(s)$:

$$F_{s_i}(s) = \begin{cases} 0 & \text{if } s \leq 0, \\ 1 - (1 - \frac{s}{P})^{\frac{1}{n-1}} & \text{if } 0 \leq s < P, \\ 1 & \text{if } s \geq P. \end{cases} \quad (3)$$

We use $Eu_i(s)$ to denote the expected value of the utility $u_i(s)$ in function of a security level s . Because $u_i(s) = -P$ if all users $j \neq i$ choose security levels higher than s , and $u_i(s) = -s$ otherwise, we have

$$Eu_i(s) = -P(Pr[s_j > s])^{n-1} - s(1 - (Pr[s_j > s])^{n-1}),$$

which can be expressed in terms of $F_{s_i}(s)$ as

$$Eu_i(s) = -P(1 - F_{s_i}(s))^{n-1} - s(1 - (1 - F_{s_i}(s))^{n-1}). \quad (4)$$

Substituting $F_{s_i}(s)$ by its expression given in Eq. (3), Eq. (4) reduces to

$$Eu_i(s) = -P + s \left(1 - \frac{s}{P}\right).$$

A study of the variations of $Eu_i(s)$ in function of $s \in [0, P]$ indicates that $Eu_i(s) \geq Eu_i(0) = -P$ and that $Eu_i(s) \leq Eu_i(P/2)$, with $Eu_i(P/2) = -3P/4$. Thus, if we have $\varepsilon = P/4$, any variation of the expected utility is smaller ε , which characterizes an ε -equilibrium. In other words, we have shown, by providing a specific c.d.f. $F_{s_i}(s)$, that there exist ε -equilibria with $\varepsilon \leq P/4$ where the security levels s_i can be spread out over the entire interval $[0, P]$. Note that we only present an existence proof here. It is unclear whether the chosen c.d.f. $F_{s_i}(s)$ is an accurate depiction of how security levels are chosen in reality, and it is likewise entirely possible that there exist other distributions of the security levels over $[0, P]$ that result in ε -equilibria for $\varepsilon \ll P/4$. \square