Malicious Bayesian Congestion Games*

Martin Gairing

International Computer Science Institute, Berkeley, CA, USA
gairing@icsi.berkeley.edu

Abstract. In this paper, we introduce *malicious Bayesian congestion games* as an extension to congestion games where players might act in a malicious way. In such a game each player has two *types*. Either the player is a rational player seeking to minimize her own delay, or – with a certain probability – the player is *malicious* in which case her only goal is to disturb the other players as much as possible.

We show that such games do in general not possess a Bayesian Nash equilibrium in pure strategies (i.e. a *pure Bayesian Nash equilibrium*). Moreover, given a game, we show that it is NP-complete to decide whether it admits a pure Bayesian Nash equilibrium. This result even holds when resource latency functions are linear, each player is malicious with the same probability, and all strategy sets consist of singleton sets of resources. For a slightly more restricted class of malicious Bayesian congestion games, we provide easy checkable properties that are necessary and sufficient for the existence of a pure Bayesian Nash equilibrium.

In the second part of the paper we study the impact of the malicious types on the overall performance of the system (i.e. the *social cost*). To measure this impact, we use the *Price of Malice*. We provide (tight) bounds on the Price of Malice for an interesting class of malicious Bayesian congestion games. Moreover, we show that for certain congestion games the advent of malicious types can also be beneficial to the system in the sense that the social cost of the worst case equilibrium decreases. We provide a tight bound on the maximum factor by which this happens.

1 Introduction

Motivation and Framework. Over the last decade, the study of strategic behavior in distributed systems has improved our understanding of modern computer artifacts such as the Internet. Normally, the users of such distributed systems are modeled as rational, utility optimizing players. However, in many real world scenarios, users do not necessarily act rationally, but rather *irrationally.* In this paper, we address one form of irrationality, namely, we allow that players act in a *malicious* way. In this case, the only goal of a malicious player is to disturb the (non-malicious) players as much as possible. The presence of *Denial of Service attacks* in the Internet is an example showing that such systems

* This work was supported by a fellowship within the Postdoc-Programme of the German Academic Exchange Service (DAAD).

E. Bampis and M. Skutella (Eds.): WAOA 2008, LNCS 5426, pp. 119–132, 2009.

are quite realistic. In many such systems with malicious players, the players have only *incomplete information* about the set of malicious players. A standard approach for modeling games with incomplete information uses the Harsanyi transformation [14], which converts a game with incomplete information to a game where players have different *types*. The type of a player represents its private information that is not common knowledge to all players. In the resulting *Bayesian game*, each player's uncertainty about each other's type is described by a probability distribution.

One aspect of Game Theory that was studied extensively in recent years is the *Price of Anarchy* as introduced by Koutsoupias and Papadimitriou [16]. The Price of Anarchy is the worst case ratio between the value of the *social cost* in an equilibrium state of the system and that of some social optimum. Usually, the equilibrium state is defined as *Nash equilibrium* – a state in which no player can unilaterally improve her private objective function, also coined as *private cost*. A Nash equilibrium is *pure* if all players choose a pure strategy and *mixed* if players choose probability distributions over pure strategies.

While the celebrated result of Nash [20] guarantees the existence of a mixed Nash equilibrium for every finite game, pure Nash equilibria are not guaranteed to exist (see e.g. [9,12,17,18]). A natural question to ask, is whether a given game possesses a pure Nash equilibrium or not. We address this question by asking about the complexity of this decision problem.

A class of games that always possess pure Nash equilibria is the class of congestion games as introduced by Rosenthal [21]. Here, the strategy set of each player is a subset of the power set of given resources, the latency on each resource is described by a latency function in the number of players sharing this resource, and the private cost of each player is the sum of the latencies of its chosen resources. Milchtaich [18] considered weighted congestion games as an extension to congestion games in which the players have weights and thus different influence on the latency of the resources.

To measure the influence of malicious behavior, Moscibroda et al. [19] introduced the *Price of Byzantine Anarchy* as the worst case ratio between the social cost in an equilibrium state of the system under some assumption on the malicious players and the social cost of some social optimum without malicious players. They further define the *Price of Malice* as the ratio between the Price of Byzantine Anarchy and the Price of Anarchy. We will use a similar definition and define the equilibrium state as a Bayesian Nash equilibrium.

Contribution. In this paper, we introduce *malicious Bayesian congestion games* as an extension to congestion games where players might act in a malicious way. Following Harsanyi's transformation [14], we allow each player to be of two *types*. Either the player is a rational player seeking to minimize her own delay, or – with a certain probability – the player is *malicious* in which case her only goal is to disturb the other players as much as possible. For such games we study the complexity of deciding whether a given game has a pure Bayesian Nash equilibrium. Moreover, we study the impact of the malicious types on the