Games and Economic Behavior 86 (2014) 438–457

Contents lists available at SciVerse ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Fault tolerance in large games

Ronen Gradwohl [a,*], Omer Reingold [b,c]

[a] Kellogg School of Management, Northwestern University, 2001 Sheridan Road, Evanston, IL 60208, United States
[b] Microsoft Research Silicon Valley, 1065 La Avenida, Mountain View, CA, 94043, United States
[c] The Weizmann Institute of Science, Rehovot 76100, Israel

ARTICLE INFO

Article history:
Received 27 June 2011
Available online 27 June 2013

JEL classification:
C72

Keywords:
Nash equilibrium
Large games
Fault tolerance

ABSTRACT

A Nash equilibrium is an optimal strategy for each player under the assumption that others play according to their respective Nash strategies, but it provides no guarantees in the presence of irrational players or coalitions of colluding players. In fact, no such guarantees exist in general. However, in this paper we show that large games are innately fault tolerant. We quantify the ways in which two subclasses of large games – λ-continuous games and anonymous games – are resilient against Byzantine faults (i.e. irrational behavior), coalitions, and asynchronous play. We also show that general large games have some non-trivial resilience against faults.

1. Introduction

Game theory provides a means of studying an interaction between many agents by modeling it as a game. One of the simplest models is a normal-form game, in which all agents play simultaneously and only once. Given such a game, a solution concept – most commonly a Nash equilibrium – is hypothesized to explain or predict the behavior of the agents.

Nowadays, with the advent of the Internet, many of the interactions we wish to model are taking place online, in a distributed setting. Such interactions are characterized by unpredictably asynchronous communication, the occasional presence of faults, and perhaps some exogenous communication between the players. The simple model of a normal-form game is far from accurate in describing this setting. Furthermore, even if it were somehow sufficient, the concept of a Nash equilibrium would provide no guarantees about the optimality of strategies in this fault-ridden environment, and hence would not suffice.

One approach to correcting this deficiency is to incorporate the characteristics of the distributed setting into the model. This, however, requires precise knowledge of the setting to be modeled, and may make the model much more complicated. A second approach that is taken in some recent literature is to quantify the damage caused by the presence of faulty players (see the section Related Work below).

In this paper we take a different approach, and show that sometimes the model of a normal-form game and the concept of a Nash equilibrium do suffice. More specifically, we show that sometimes the strategies of a Nash equilibrium in a normal-form game remain close to best responses for each player even if the game is modified to permit asynchronous communication, faulty behavior, and the possibility of collusion.

We study large games – games that involve many agents – because in such games it is particularly relevant to examine the robustness of equilibria against faults. Additionally, in large games it seems plausible that the sheer size of the game renders it robust, whereas in small games such an approach is easily seen to fail. Of course, a large game could be comprised

* Corresponding author.
E-mail addresses: r-gradwohl@kellogg.northwestern.edu (R. Gradwohl), omer.reingold@microsoft.com (O. Reingold).

0899-8256/$ – see front matter
http://dx.doi.org/10.1016/j.geb.2013.06.007

of a small game embedded in a setting with many players, in which case our approach would not work. But then again, such a game is not really a "large" game. Thus, we consider various properties of the game that are reasonable characterizations of large games, and show how these properties relate to the fault tolerance of the game.

1.1. Faulty behavior

The guarantee provided by a Nash equilibrium is that each player's strategy is optimal, assuming all others play their designated strategies. One type of faulty behavior consists of Byzantine faults,[1] wherein some of the players do not play according to their equilibrium strategies. Perhaps the actions of these faulty players are altered due to an error in communication, perhaps the players are irrational, or perhaps they have some unknown utility. In any case, their actions can be arbitrary or even adversarial.

Two types of resilience we wish to have against Byzantine faults are immunity and tolerance. Immunity means that even if some players fault, the utility of the non-faulty players is not affected by much. Such a property is useful in that it provides the (non-faulty) players a guaranteed utility under the equilibrium strategies, even if some others are faulty. More formally, we say that an equilibrium is (ε, t)-immune if players' expected utilities do not decrease by more than ε when any t other players deviate arbitrarily. Tolerance means that even if some players fault, the original strategies of non-faulty players are still optimal (even conditioned on the actions of the faulty players), although their payoffs may be different from the case in which no players fault. More formally, an equilibrium is (ε, t)-tolerant if players' equilibrium strategies remain best responses (within ε) even if t other players deviate arbitrarily.

Another type of faulty behavior is the potential strategizing enabled by asynchrony. Nash equilibrium strategies are optimal if all players play simultaneously, but suppose some player was delayed, and was then exposed to the chosen actions of some of the other players. This player may then use the information he obtained from the exposed actions to get a better payoff for himself. In this case, his original Nash strategy may no longer be optimal. Resilience to asynchrony means that this does not occur – that a player cannot improve his payoff by too much, even after he is exposed to the actions of some other players. More formally, a strategy profile is an (ε, p, t)-ex post Nash equilibrium if with probability $1 - p$ any player's original strategy is still nearly optimal (up to an additive ε), even after he sees the chosen actions of a set of t players.

A final type of faulty behavior consists of collusion among players. Nash equilibria only guarantee that no player can benefit by a unilateral deviation from his strategy, but imply nothing about deviations by a colluding set of players. Resilience against collusion means that coalitions of players cannot improve their payoffs even by a coordinated deviation. More formally, an equilibrium is (ε, t)-coalitional if no player can gain more than ε even with a joint deviation by himself and $t - 1$ other players.

1.2. Our results

The first subclass of games we examine is one that has received much attention in the economics literature (see, for example, Kalai, 2004; Azrieli and Shmaya, 2010, and the references therein). This subclass consists of λ-continuous games, and it turns out that it is particularly well-suited to our demands of fault tolerance. Roughly, in such games no player is influenced too much by the other players. More precisely, if a coalition consisting of a δ-fraction of the players changes their actions, then this can change the utility of a non-member by at most $\lambda \cdot \delta$. For games that are λ-continuous, we observe that all Nash equilibria (ε, t)-tolerant, (ε, t)-immune, and (ε, t)-coalitional for $t = O(n)$ (assuming λ and ε are arbitrary constants). They are also resilient to asynchronous exposure to a large number of players – always to at least $O(n)$ players, and sometimes to all $n - 1$ other players.

λ-continuous games have very strong fault tolerance, but perhaps the restriction on the game is too strong. More precisely, in these games a player's utility given **any** profile of the others' actions changes little if a small fraction of players changes their action. For fault tolerance, however, we are really only interested in what happens when the actions of (non-faulty) players come from specific distributions – namely, the strategies of a Nash equilibrium.

The next natural subclass of games we study is that of anonymous games, in which the utility of each player is a function of his own action and the empirical distribution of other players' actions. Thus, in anonymous games a player does not care which player performed each action, but only how many performed each action. Anonymous games are widely studied in economics. Note that anonymous games can be very far from being λ-continuous: For example, if a player's utility is determined by the majority of the other players' binary actions, then a change of even 1 player's action can flip the utility from 1 to 0.

We show that every anonymous game has an equilibrium that is (ε, t)-tolerant, (ε, t)-immune, and (ε, t)-coalitional, for $t = O(\sqrt{n})$. Additionally, the equilibrium is an (ε, p, t')-ex post Nash equilibrium for $t' = O(n)$. In order to prove this theorem we actually show that all equilibria that are "mixed enough" – in which players play every action with at least some minimal probability – have these fault tolerance guarantees. Section 5.2 contains some discussion about the applicability and naturalness of this assumption, as well examples demonstrating its necessity.

[1] Byzantine faults are named after the Byzantine generals' problem of Pease et al. (1980).