Games and Economic Behavior 85 (2014) 19–31

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

On the communication complexity of approximate Nash equilibria ☆

Paul W. Goldberg[a,*], Arnoud Pastink[b]

[a] University of Oxford, Dept. of Computer Science, Wolfson Building, Parks Road, Oxford OX1 3QD, UK
[b] Utrecht University, Department of Information and Computing Science, PO Box 80089, 3508TB Utrecht, The Netherlands

ARTICLE INFO

Article history:
Received 25 February 2013
Available online 3 February 2014

JEL classification:
C7
D83

Keywords:
Two player
Mixed strategy
Approximate equilibria
Efficient algorithms

ABSTRACT

We study the problem of computing approximate Nash equilibria of bimatrix games, in a setting where players initially know their own payoffs but not the other player's. In order to find a solution of reasonable quality, some amount of communication is required. We study algorithms where the communication is substantially less than the size of the game. When the communication is polylogarithmic in the number of strategies, we show how to obtain ϵ-approximate Nash equilibrium for $\epsilon \approx 0.438$, and for well-supported approximate equilibria we obtain $\epsilon \approx 0.732$. For one-way communication we show that $\epsilon = \frac{1}{2}$ is the best approximation quality achievable, while for well-supported equilibria, no value of $\epsilon < 1$ is achievable. When the players do not communicate at all, ϵ-Nash equilibria can be obtained for $\epsilon = \frac{3}{4}$; we also provide a corresponding lower bound of slightly more than $\frac{1}{2}$ on the smallest constant ϵ achievable.

1. Introduction

Algorithmic game theory is concerned not just with properties of a solution concept, but also how that solution can be obtained. It is considered desirable that the outcome of a game should be "easy to compute", which is typically formalized as polynomial-time computability, in the algorithms community. In that respect the PPAD-completeness results of Daskalakis et al. (2009a) and Chen and Deng (2006) are interpreted as a "complexity-theoretic critique" of Nash equilibrium. Following those results, a line of work addressed the problem of computing ϵ-Nash equilibrium, where $\epsilon > 0$ is a parameter that bounds a player's incentive to deviate, in a solution. Thus, ϵ-Nash equilibrium imposes a weaker constraint on how players are assumed to behave, and an exact Nash equilibrium is obtained for $\epsilon = 0$. The main open problem is to find out what values of ϵ admit a polynomial-time algorithm. Below we summarize some of the progress in this direction.

Beyond the existence of a fast algorithm, it is also desirable that a solution should be obtained by a process that is simple and decentralized, since that is likely to be a better model for how players in a game may eventually reach a solution. In that respect, most of the known efficient algorithms for computing ϵ-Nash equilibria are not entirely satisfying. They take as input the payoff matrices and output the approximate Nash equilibrium. If we try to translate such an algorithm into real life, it would correspond to a process where the players pass their payoffs to a central authority, which returns to them some mixed strategies that have the "low incentive to deviate" guarantee. In this paper we aim to model a setting where

☆ A preliminary version of this paper appeared in the Proceedings of the 5th SAGT. The first author was supported by EPSRC Grant EP/G069239/1 "Efficient Decentralised Approaches in Algorithmic Game Theory".

* Corresponding author.
E-mail addresses: Paul.Goldberg@cs.ox.ac.uk (P.W. Goldberg), A.J.Pastink@uu.nl (A. Pastink).

http://dx.doi.org/10.1016/j.geb.2014.01.009
0899-8256/© 2014 Elsevier Inc. All rights reserved.

players perform individual computations and exchange some limited information. We revisit the question of what values of ϵ are achievable, subject to this restriction to more "realistic" algorithms.

There are various ways in which one can try to model the notion of a decentralized algorithm; here we consider a general approach that has previously been studied in Conitzer and Sandholm (2004) and Hart and Mansour (2010) in the context of computing exact Nash equilibria. The players begin with knowledge of their own payoffs but not the payoffs of the other players; this is often called an *uncoupled* setting (see Section 1.2.4 for an overview). An algorithm involves communication in addition to computation; to find a game-theoretic solution, a player usually has to know something about the other players' matrices, but hopefully not all of that information. We study the computation of ϵ-Nash equilibria in this setting, and the general topic is the trade-off between the amount of communication that takes place, and the value of ϵ that can be obtained. In uncoupled settings, there are natural dynamic processes that converge to correlated equilibria, but the results are less positive for exact and approximate Nash equilibria. This paper aims to contribute to the general goal of evaluating the merits of approximate Nash equilibrium as a solution concept, as opposed to (for example) exact or approximate correlated equilibrium.

1.1. Definitions

We consider 2-player games, with a *row player* and a *column player*, who both have n *pure strategies*. The game (R, C) is defined by two $n \times n$ *payoff matrices*, R for the row player, and C for the column player. The pure strategies for the row player are his rows and the pure strategies of the column player are her columns. If the row player plays row i and the column player plays column j, the *payoff* for the row player is R_{ij}, and C_{ij} for the column player. For the row player a *mixed strategy* is a probability distribution $\mathbf{x}$ over the rows, and a mixed strategy for the column player is a probability distribution $\mathbf{y}$ over the columns, where $\mathbf{x}$ and $\mathbf{y}$ are column vectors and $(\mathbf{x}, \mathbf{y})$ is a *mixed strategy profile*. The payoffs resulting from these mixed strategies $\mathbf{x}$ and $\mathbf{y}$ are $\mathbf{x}^{\mathrm{T}} R\mathbf{y}$ for the row player and $\mathbf{x}^{\mathrm{T}} C\mathbf{y}$ for the column player.

A *Nash equilibrium* is a pair of mixed strategies $(\mathbf{x}^*, \mathbf{y}^*)$ where neither player can get a higher payoff by playing another strategy assuming the other player does not change his strategy. Because of the linearity of a mixed strategy, the largest gain can be achieved by defecting to a pure strategy. Let $\mathbf{e}_i$ be the vector with a 1 at the i-th position and a 0 at every other position. Thus a Nash equilibrium $(\mathbf{x}^*, \mathbf{y}^*)$ satisfies

$$\forall i = 1 \ldots n \quad \mathbf{e}_i^{\mathrm{T}} R\mathbf{y}^* \leqslant \left(\mathbf{x}^*\right)^{\mathrm{T}} R\mathbf{y}^* \quad \text{and} \quad \left(\mathbf{x}^*\right)^{\mathrm{T}} C\mathbf{e}_i \leqslant \left(\mathbf{x}^*\right)^{\mathrm{T}} C\mathbf{y}^*.$$

We assume that the payoffs of R and C are between 0 and 1, which can be achieved by affine transformations. An *ϵ-approximate Nash equilibrium* (or, ϵ-Nash equilibrium) is a strategy pair $(\mathbf{x}^*, \mathbf{y}^*)$ such that each player can gain at most ϵ by unilaterally deviating to a different strategy. Thus, it is $(\mathbf{x}^*, \mathbf{y}^*)$ satisfying

$$\forall i = 1 \ldots n \quad \mathbf{e}_i^{\mathrm{T}} R\mathbf{y}^* \leqslant \left(\mathbf{x}^*\right)^{\mathrm{T}} R\mathbf{y}^* + \epsilon \quad \text{and} \quad \left(\mathbf{x}^*\right)^{\mathrm{T}} C\mathbf{e}_i \leqslant \left(\mathbf{x}^*\right)^{\mathrm{T}} C\mathbf{y}^* + \epsilon.$$

We say that the *regret* of a player is the difference between his payoff and the payoff of his best response.

The *support* of a mixed strategy $\mathbf{x}$, denoted $\mathrm{Supp}(\mathbf{x})$, is the set of pure strategies that are played with non-zero probability by $\mathbf{x}$. An *approximate well-supported Nash equilibrium* strengthens the requirements of an approximate Nash equilibrium. For a mixed strategy $\mathbf{y}$ of the column player, a pure strategy $i \in [n]$ is an *ϵ-best response* for the row player if, for all pure strategies $i' \in [n]$ we have: $\mathbf{e}_i^{\mathrm{T}} R\mathbf{y} \geqslant \mathbf{e}_{i'}^{\mathrm{T}} R\mathbf{y} - \epsilon$. We define ϵ-best responses for the column player analogously. A mixed strategy profile $(\mathbf{x}, \mathbf{y})$ is an *ϵ-well-supported Nash equilibrium* (ϵ-WSNE) if every pure strategy in $\mathrm{Supp}(\mathbf{x})$ is an ϵ-best response against $\mathbf{y}$, and every pure strategy in $\mathrm{Supp}(\mathbf{y})$ is an ϵ-best response against $\mathbf{x}$.

The communication model: Each player $q \in \{r, c\}$ has an algorithm $\mathcal{A}_q$ whose initial input data is q's $n \times n$ payoff matrix. Communication proceeds in a number of rounds, where in each round, each player may send a single bit of information to the other player. During each round, each player may also carry out a polynomial (in n) amount of computation. (A natural variant of the model would omit the restriction to polynomial computation. Indeed, our lower bounds on communication requirement do not depend on computational limits.) At the end, each player q outputs a mixed strategy $\mathbf{x}_q$. We aim to design (pairs of) algorithms $(\mathcal{A}_r, \mathcal{A}_c)$ that output ϵ-Nash strategy profiles $(\mathbf{x}_r, \mathbf{x}_c)$, and are economical with the number of rounds of communication. This is similar to the *mixed Nash equilibrium procedure* of Hart and Mansour (2010), here applied to approximate rather than exact equilibria.

Notice that given $\Theta(n^2)$ rounds of communication, we can apply any centralized algorithm $\mathcal{A}$ by getting (say) the row player to pass additive approximations of all his payoffs to the column player, who applies $\mathcal{A}$ and passes to the row player the mixed strategy obtained by $\mathcal{A}$ for the row player. (The quality of the ϵ-Nash equilibrium is proportional to the quality of the additive approximations used.) For this reason we focus on algorithms with many fewer rounds, and we obtain results for logarithmic or polylogarithmic (in n) rounds.

We also consider a restriction to *one-way communication*, where one player may send but not receive information.