Strong Nash equilibrium is in smoothed $\mathcal{P}$

Nicola Gatti and **Marco Rocco**
Politecnico di Milano
Piazza Leonardo da Vinci 32
Milano, Italy
{ngatti, mrocco}@elet.polimi.it

Tuomas Sandholm
Carnegie Mellon University
Computer Science Department
5000 Forbes Avenue
Pittsburgh, PA 15213, USA
sandholm@cs.cmu.edu

Abstract

The computational characterization of game–theoretic solution concepts is a prominent topic in computer science. The central solution concept is *Nash equilibrium* (NE). However, it fails to capture the possibility that agents can form coalitions. *Strong Nash equilibrium* (SNE) refines NE to this setting. It is known that finding an SNE is $\mathcal{NP}$–complete when the number of agents is constant. This hardness is solely due to the existence of mixed–strategy SNEs, given that the problem of enumerating all pure–strategy SNEs is trivially in $\mathcal{P}$. Our central result is that, in order for an n–agent game to have at least one non–pure–strategy SNE, the agents' payoffs restricted to the agents' supports must lie on an $(n-1)$–dimensional space. Small perturbations make the payoffs fall outside such a space and thus, unlike NE, finding an SNE is in smoothed polynomial time.

Introduction

The central solution concept provided by game theory is *Nash equilibrium* (NE). Finding an NE of a strategic–form (aka normal–form) game is $\mathcal{PPAD}$–complete (Daskalakis *et al.* 2006) even with just two agents (Chen *et al.* 2009). Although $\mathcal{PPAD} \subseteq \mathcal{NP}$, it is generally believed that $\mathcal{PPAD} \neq \mathcal{P}$ and therefore that there does not exist any polynomial–time algorithm to find an NE unless $\mathcal{P} = \mathcal{NP}$. Furthermore, 2–agent games do not have a fully polynomial–time approximation scheme unless $\mathcal{PPAD} \subseteq \mathcal{P}$ (Chen *et al.* 2009) and finding an NE is not in smoothed $\mathcal{P}$ unless $\mathcal{PPAD} \subseteq \mathcal{RP}$ (Chen *et al.* 2006) and, therefore, by definition of smoothed complexity, game instances remain hard even if subjected to small perturbations.

NE captures the situation in which no agent can gain more by unilaterally changing her strategy. When agents can form coalitions and change their strategies multilaterally in a coordinated way, the most natural solution concept is *strong Nash equilibrium* (SNE) (Aumann 1960). An SNE is a strategy profile from which no coalition can deviate in a way that benefits each of the deviators. Thus, a strategy profile is an SNE when it is weakly Pareto efficient over the space of all the strategy profiles for each possible coalition. An SNE is not assured to exist. Finding an SNE (determining whether one exists) is $\mathcal{NP}$–complete when the number of agents is constant; $\mathcal{NP}$–hardness was proven in (Conitzer and Sandholm 2008) and membership in $\mathcal{NP}$ in (Gatti *et al.* 2013b). Unlike for NE, the literature has very few algorithms for SNE and almost all of them focus only on pure–strategy SNEs for specific classes of games, e.g., congestion games (Holzman and Law-Yone 1997; Hayrapetyan *et al.* 2006; Rozenfeld and Tennenholtz 2006; Hoefer and Skopalik 2010), connection games (Epstein *et al.* 2007), maxcut games (Gourvès and Monnot 2009), and continuous games (Nessah and Tian 2012). The only algorithms for finding mixed–strategy SNEs with general games are presented in (Gatti *et al.* 2013a; 2013b). SNE hardness is only due to the existence of mixed–strategy equilibria.

In this paper, we show that if there is a mixed–strategy SNE, then the payoffs restricted to the actions in the support must satisfy strict geometric conditions. For example, in 2–agent games, they must lie on the same line in agents' utilities space. Leveraging this result, we show that finding an SNE is in smoothed $\mathcal{P}$ since, in the generic case (i.e., in all except knife–edge cases), all SNEs are pure.

Preliminaries

A strategic–form game is a tuple (N, A, U) where (Shoham and Leyton-Brown 2008): $N = \{1, \ldots, n\}$ is the set of agents (we denote by i a generic agent); $A = A_1 \times \ldots \times A_n$ is the set of agents' joint actions and A_i is the set of agent i's actions (we denote a generic action by a, and by m_i the number of actions in A_i); $U = \{U_1 \ldots, U_n\}$ is the set of agents' utility arrays where $U_i(a_1, \ldots, a_n)$ is agent i's utility when the agents play actions $a_1, \ldots, a_n$. We denote by $x_i(a_i)$ the probability with which agent i plays action $a_i \in A_i$ and by $\mathbf{x}_i$ the vector of probabilities $x_i(a_i)$ of agent i. We denote by Δ_i the space of well–defined probability vectors over A_i. We denote by S_i the support of agent i, that is, the set of actions played with positive probability.

Games and mixed–strategy SNEs

We study the properties of mixed–strategy SNEs. We first discuss the 2–agent case and then the n–agent case. We denote by P_{mix} and by P_{cor} the sets of points in the agents' utility spaces $\mathbb{E}[U_1] \times \mathbb{E}[U_2]$ that are on the Pareto frontier when the agents play *mixed* and *correlated* strategies, respectively. Obviously, points in P_{cor} non–strictly Pareto dominate points in P_{mix}, given that mixed strategies constitute a subset of correlated strategies. We denote by $P_{mix}(S_1, S_2)$ and $P_{cor}(S_1, S_2)$ the Pareto frontiers in mixed and correlated strategies, respectively, when all the actions outside supports S_1 and S_2 are removed.

Theorem 1 *Consider a non–degenerate 2–agent game with two actions per agent. If there is a mixed–strategy SNE, then* $P_{mix} = P_{cor}$.

Proof. We can write down the game as follows:

agent 1 \ agent 2	a_3	a_4
a_1	p_1, q_1	p_2, q_2
a_2	p_3, q_3	p_4, q_4

Since there is a mixed–strategy NE:

$$x_2(\mathsf{a}_3) \cdot p_1 + x_2(\mathsf{a}_4) \cdot p_2 = x_2(\mathsf{a}_3) \cdot p_3 + x_2(\mathsf{a}_4) \cdot p_4$$
$$x_1(\mathsf{a}_1) \cdot q_1 + x_1(\mathsf{a}_2) \cdot q_3 = x_1(\mathsf{a}_1) \cdot q_2 + x_1(\mathsf{a}_3) \cdot q_4$$

Moreover, being an SNE, the mixed–strategy profile has to satisfy the Karush–Kuhn–Tucker conditions necessary conditions for local weak Pareto efficiency (Miettinen 1999):

$$-\lambda_1 \cdot (x_2(\mathsf{a}_3) \cdot p_1 + x_2(\mathsf{a}_4) \cdot p_2) - \lambda_2 \cdot (x_2(\mathsf{a}_3) \cdot q_1 + x_2(\mathsf{a}_4) \cdot q_2) = \nu_1$$
$$-\lambda_1 \cdot (x_2(\mathsf{a}_3) \cdot p_3 + x_2(\mathsf{a}_4) \cdot p_4) - \lambda_2 \cdot (x_2(\mathsf{a}_3) \cdot q_3 + x_2(\mathsf{a}_4) \cdot q_4) = \nu_1$$
$$-\lambda_1 \cdot (x_1(\mathsf{a}_1) \cdot p_1 + x_1(\mathsf{a}_3) \cdot p_3) - \lambda_2 \cdot (x_1(\mathsf{a}_1) \cdot q_1 + x_1(\mathsf{a}_3) \cdot q_3) = \nu_2$$
$$-\lambda_1 \cdot (x_1(\mathsf{a}_1) \cdot p_2 + x_1(\mathsf{a}_2) \cdot p_4) - \lambda_2 \cdot (x_1(\mathsf{a}_1) \cdot q_2 + x_1(\mathsf{a}_3) \cdot q_4) = \nu_2$$

By combining the above conditions and excluding degenerate cases, we obtain:

$$\frac{p_1 - p_2}{p_4 - p_3} = \frac{q_1 - q_2}{q_4 - q_3} \qquad \frac{p_1 - p_3}{p_4 - p_2} = \frac{q_1 - q_3}{q_4 - q_2}$$

We can give a simple geometric interpretation of the above conditions. Call $R_i = (p_i, q_i)$. Each R_i is a point in the space $\mathbb{E}[U_1] \times \mathbb{E}[U_2]$. The above conditions state that:

- $\overline{R_1 R_2}$ is parallel to $\overline{R_3 R_4}$,
- $\overline{R_1 R_3}$ is parallel to $\overline{R_2 R_4}$,

and therefore R_1, R_2, R_3, R_4 are the vertices of a parallelogram, see Fig. 1(a). Given that

- a mixed–strategy NE is strictly inside the parallelogram (it being the convex (non–degenerate) combination of the vertices), see Fig. 1(a), and that
- it must be on a Pareto efficient edge (since, if it is strictly inside the parallelogram—as in Fig. 1(a)—then it is Pareto dominated by some point on some edge),

we have that R_1, R_2, R_3, R_4 must be aligned according to a line of the form $\mathbb{E}[U_1] + \phi \cdot \mathbb{E}[U_2] = const$ with $\phi \in (-1, 0)$, see, e.g., Fig. 1(b). Thus, the combination of R_1, R_2, R_3, R_4 through every mixed–strategy profile lies on the line connecting the two extreme vertices; e.g., in Fig. 1(b) the extreme vertices are R_2 and R_1. Thus, $P_{mix} = P_{cor}$. □

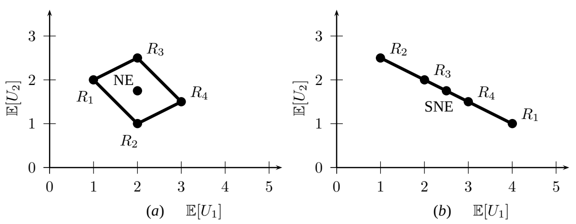

Figure 1: Examples used in the proof of Theorem 1.

We now extend the previous result to the setting in which each agent has m actions and $|S_1| = |S_2| = 2$.

Corollary 2 *Consider a non–degenerate 2–agent game with m actions per agent. If there is a mixed–strategy SNE with support sizes $|S_1| = |S_2| = 2$, then $P_{mix}(S_1, S_2) = P_{cor}(S_1, S_2)$.*

Proof. We can split the NE constraints and KKT conditions into two groups: those generated considering deviations towards pure or mixed strategies over the supports S_1 and S_2 and those generated considering deviations towards pure or mixed strategies over actions off the supports S_1 and S_2. The constraints belonging to the first group are the same as in the case with two actions per agent considered in the proof of Theorem 1. The second group overconstrains the problem and it is not necessary for the proof. Thus, restricting the game to the actions in S_1 and S_2, Theorem 1 holds and therefore $P_{mix}(S_1, S_2) = P_{cor}(S_1, S_2)$. □

The extension to the general case follows.

Corollary 3 *Consider a 2–agent game,if there is a mixed–strategy SNE in which agents' supports are S_1, S_2, then $P_{mix}(S_1, S_2) = P_{cor}(S_1, S_2)$.*

We will now discuss how the above results extend to more than two agents. For example, in the 3–agent setting, the vector of payoffs for each action profile is $R_{i,j,k} = (U_1(i, j, k), U_2(i, j, k), U_3(i, j, k))$. The crucial result is that necessary conditions, generated for only the actions in the supports, for mixed–strategy SNEs forced by NE constraints with KKT conditions for all the coalitions (i.e., $\{1, 2\}$, $\{1, 3\}$, $\{2, 3\}$, $\{1, 2, 3\}$) require that all the $R_{i,j,k}$ lie on a plane (with n–agent games, all the payoff vectors on the support must lie on an $(n-1)$–dimensional hyperplane). Thus:

Theorem 4 *Consider an n–agent game. If there is a mixed–strategy SNE with in which agents' supports are S_1, S_2, then $P_{mix}(S_1, S_2) = P_{cor}(S_1, S_2)$.*

Leveraging the above results we can state the following.

Theorem 5 *Given n agents, searching for an SNE is in smoothed $\mathcal{P}$.*

Proof sketch. With non–degenerate games, we can develop a support–enumeration algorithm scanning the pure strategies first and then, if no pure SNE exists, it checks whether there are payoffs on supports $|S_1| = |S_2| = 2$ that are aligned. If there are no such payoffs, the algorithm terminates, otherwise it enumerates all the possible supports. This algorithm goes into the exponential support enumeration with zero probability and therefore its expected compute time is polynomial. The case with degenerate games is similar. □

Future research

In future research we plan to study the computational complexity of approximating SNE and to design algorithms to do so. We also plan to study computational issues related to *strong correlated equilibrium*. This concept should present different properties than SNE, e.g., the convexity of the Pareto frontier with this solution concept could make the computation of an equilibrium easier and could make equilibria not sensitive to small perturbations.

Acknowledgments

Nicola Gatti was supported by MIUR under grant 2009BZM837. Tuomas Sandholm was supported by NSF under grants IIS–0964579 and CCF–1101668.

References

R. Aumann. Acceptable points in games of perfect information. *PAC J MATH*, 10:381–417, 1960.

X. Chen, X. Deng, and S.H. Teng. Computing Nash equilibria: approximation and smoothed complexity. In *FOCS*, pages 603–612, 2006.

X. Chen, X. Deng, and S.-H. Teng. Settling the complexity of computing two–player Nash equilibria. *J ACM*, 56(3):14:1–14:57, 2009.

V. Conitzer and T. Sandholm. New complexity results about Nash equilibria. *GAME ECON BEHAV*, 63(2):621–641, 2008.

C. Daskalakis, P. Goldberg, and C. Papadimitriou. The complexity of computing a Nash equilibrium. In *STOC*, pages 71–78, 2006.

A. Epstein, M. Feldman, and Y. Mansour. Strong equilibrium in cost sharing connection games. In *ACM EC*, pages 84–92, 2007.

N. Gatti, M. Rocco, and T. Sandholm. Algorithms for strong Nash equilibrium with more than two agents. In *AAAI*, 2013.

N. Gatti, M. Rocco, and T. Sandholm. On the verification and computation of strong Nash equilibrium. In *AAMAS*, 2013.

L. Gourvès and J. Monnot. On strong equilibria in the max cut game. In *WINE*, pages 608–615, 2009.

A. Hayrapetyan, E. Tardos, and T. Wexler. The effect of collusion in congestion games. In *STOC*, pages 89–98, 2006.

M. Hoefer and A. Skopalik. On the complexity of Pareto–optimal Nash and strong equilibria. In *SAGT*, pages 312–322, 2010.

R. Holzman and N. Law-Yone. Strong equilibrium in congestion games. *GAME ECON BEHAV*, 21:85–101, 1997.

K. Miettinen. *Multiobjective Optimization*. Kluwer, 1999.

R. Nessah and G. Tian. On the existence of strong Nash equilibria. *IESEG School of Management, Working Paper*, 2012.

O. Rozenfeld and M. Tennenholtz. Strong and correlated strong equilibria in monotone congestion games. In *WINE*, pages 74–86, 2006.

Y. Shoham and K. Leyton-Brown. *Multiagent Systems: Algorithmic, Game Theoretic and Logical Foundations*. Cambridge University Press, 2008.