Zero-Sum Game Techniques for Approximate Nash Equilibria*

(Extended Abstract)

Artur Czumaj[1] Michail Fasoulakis[1,2] Marcin Jurdziński[1]
Department of Computer Science and Centre for Discrete Mathematics and its Applications (DIMAP)
University of Warwick[1]
Institute of Computer Science, Foundation for Research and Technology-Hellas (ICS-FORTH)[2]
{A.Czumaj, M.Fasoulakis, M.Jurdzinski}@warwick.ac.uk

ABSTRACT

We apply existing, and develop new, zero-sum game techniques for designing polynomial-time algorithms to compute additive approximate Nash equilibria in bimatrix games. In particular, we give a polynomial-time algorithm that given an arbitrary bimatrix game as an input, outputs either an additive $\frac{1}{3}$-Nash equilibrium or an additive $\frac{1}{2}$-well-supported Nash equilibrium; and we give a polynomial-time algorithm that given a bimatrix game in which both payoff matrices are symmetric as an input, computes an additive $\frac{1}{2}$-well-supported Nash equilibrium. The former result is unusual: the obvious weakness is that the algorithm does not guarantee which of the two kinds of approximate equilibria it will output, but on the other hand each of the two approximation guarantees it gives are better than the best unconditional bounds known to be computable in polynomial time: 0.3393 for Nash equilibria and 0.6528 for well-supported Nash equilibria. In the latter case, we motivate the interest in computing additive approximate Nash equilibria efficiently for bimatrix games with symmetric payoff matrices by proving that computing Nash equilibria in bimatrix games is PPAD-complete even if both of the payoff matrices are symmetric.

1. INTRODUCTION

One of the most fundamental problems in algorithmic game theory is the *computation of a Nash equilibrium*. After the computational complexity of this problem has been established to be complete in the class PPAD, even for bimatrix games [5, 8], further research has been focusing on the *computation of approximate Nash equilibria*. One very natural kind of approximation of the Nash equilibria is the additive approximation and this has two different notions: the ε-*Nash equilibria* and the ε-*well-supported Nash equilibria*.

Let (R, C) be a bimatrix game: $R \in [0,1]^{n\times n}$ is the payoff matrix of the *row player* and $C \in [0,1]^{n\times n}$ is the payoff matrix of the *column player*. If the row player uses the *pure strategy* $i \in \{1,\ldots,n\}$ and the column player uses the pure strategy $j \in \{1,\ldots,n\}$, then the row player receives payoff R_{ij} and the column player receives payoff C_{ij}.

A *(mixed) strategy* $x \in [0,1]^n$ is a probability vector: for every pure strategy $i \in \{1,\ldots,n\}$, the probability with which the mixed strategy x uses the pure strategy i is x_i; naturally, we require that $\sum_{i=1}^{n} x_i = 1$. For a pure strategy $i \in \{1,\ldots,n\}$, we write e_i for the strategy that uses the pure strategy i with probability 1 (and all other pure strategies with probability 0). The *support* of a strategy x, denoted by $\mathrm{supp}(x)$, is the set of pure strategies that are *used* (that is, have non-zero probability) in x. A *strategy profile* is a pair (x, y) of strategies: the strategy x is used by the row player and the strategy y is used by the column player. The *expected payoffs* in a strategy profile that the row player and the column player receive are $x^T Ry$ and $x^T Cy$, respectively.

A *best response* to a strategy x of the row player is a strategy y^* of the column player, such that for every strategy y, we have $x^T Cy^* \geq x^T Cy$. Note that this is equivalent to stating that all pure strategies used (with non-zero probability) in y^* are best responses to x, that is for every pure strategy $j^* \in \mathrm{supp}(y^*)$ and for every strategy $j \in \{1,\ldots,n\}$, we have $x^T Ce_{j^*} \geq x^T Ce_j$. Analogously, a best response to a strategy y of the column player is a strategy x^* of the row player, such that for every strategy x, we have $(x^*)^T Ry \geq x^T Ry$, which is equivalent to stating that all pure strategies used in x^* are best responses to y. A *Nash equilibrium* (NE) is a strategy profile (x^*, y^*), such that y^* is a best response to x^*, and x^* is a best response to y^*.

For an $\varepsilon \geq 0$, an ε-*best response* to a strategy x of the row player is a strategy y^* of the column player such that for every strategy y, we have $x^T Cy^* \geq x^T Cy - \varepsilon$. Analogously an ε-best response to a strategy y of the column player is a strategy x^* of the row player such that for every strategy x, we have $(x^*)^T Ry \geq x^T Ry - \varepsilon$. An additive ε-*Nash equilibrium* (ε-NE) is a strategy profile (x^*, y^*) such that y^* is an ε-best response to x^*, and x^* is an ε-best response to y^*.

Note that—unlike for best responses—if y^* is an ε-best response to x then it is not necessarily the case that every pure strategy used in y^* is an ε-best response to x. By demanding the latter—and stronger—condition to hold instead, we obtain a refinement of the concept of an ε-NE: an additive ε-*well-supported Nash equilibrium* (ε-WSNE) is a strategy profile (x^*, y^*) such that every pure strategy used in y^* is an ε-best response to x^*, and every pure strategy used in x^* is an ε-best response to y^*.

We are interested in the problem of computing additive approximate Nash equilibria in polynomial time. The research programme here is to find polynomial-time algorithms to compute additive ε-NE or additive ε-WSNE for as small values of $\varepsilon > 0$ as possible.

*Research partially supported by the Centre for Discrete Mathematics and its Applications (DIMAP) and by the EPSRC awards EP/D063191/1 and EP/G069034/1.

Appears in: *Proc. of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017), S. Das, E. Durfee, K. Larson, M. Winikoff (eds.), May 8–12, 2017, São Paulo, Brazil.*

The problem of computing additive ε-NE in polynomial time has been extensively studied [15, 10, 9, 4], and the currently best approximation guarantee is due to Tsaknakis and Spirakis [19], whose algorithm achieves it for $\varepsilon = 0.3393$. On the other hand, for every $\varepsilon > 0$, there is a quasi-polynomial time algorithm for computing additive ε-NE [17, 14, 3]. Some lower bounds are also known: poly-logarithmic support size is necessary for constructing additive ε-NE for every $\varepsilon < \frac{1}{2}$ [12], and complexity-theoretic evidence suggests that there is a constant $\varepsilon > 0$, such that additive ε-NE cannot be computed significantly faster than in quasi-polynomial time [18].

Every additive ε-WSNE is an additive ε-NE, but not vice versa. As for additive ε-NE, for every $\varepsilon > 0$, there is a quasi-polynomial time algorithm for computing additive ε-WSNE [16]. The problem of computing additive ε-WSNE in polynomial time has seen less progress [16, 11] than for additive ε-NE, and the currently best approximation guarantee is due to Czumaj et al. [6], whose algorithm achieves it for $\varepsilon = 0.6528$. Better approximation guarantees are known to be achievable in polynomial time for two classes of bimatrix games: $\varepsilon = \frac{1}{2}$ for win-lose games [16]; and $\varepsilon = \frac{1}{2} + \delta$, for every $\delta > 0$, for symmetric games [7]. Finally, poly-logarithmic support size is necessary for constructing additive ε-WSNE for all $\varepsilon < 1$ [2, 1].

Because of space limitations we defer most proofs and detailed discussion to the full version of this paper.

2. HARDNESS RESULTS

We motivate the study of polynomial-time algorithms for computing additive *approximate* Nash equilibria in bimatrix games with symmetric payoff matrices (to be discussed in Section 4.1) by establishing the following hardness results.

THEOREM 1. *The problems of computing a Nash equilibrium in bimatrix games in which either one or both payoff matrices are symmetric are* PPAD*-complete.*

3. ZERO-SUM GAME TECHNIQUES

In this section we recall an approximate equilibrium construction technique for a bimatrix game (R, C) due to Czumaj et al. [6], based on solving the zero-sum games $(R, -R)$ and $(-C, C)$ considered earlier by Goldberg and Pastink [13]. Further, we develop a new technique based on solving zero-sum games $(-R, R)$ and $(C, -C)$ instead.

3.1 Zero-sum games $(R, -R)$ and $(-C, C)$

Let (x^*, y^*) and $(\hat{x}, \hat{y})$ be Nash equilibria in the zero-sum games $(R, -R)$ and $(-C, C)$, respectively. Let $v_R = (x^*)^T R y^*$ (which is the value of the zero-sum game $(R, -R)$), and let $v_C = \hat{x}^T C \hat{y}$ (which is minus the value of the zero-sum game $(-C, C)$). Assume, without loss of generality, that $v_R \geq v_C$.

LEMMA 2. *[13, 6] The strategy profile $(\hat{x}, y^*)$ is an additive v_R-WSNE.*

Let j be a pure best response to the strategy x^* of the row player, and let r be a pure best response to the strategy j of the column player.

LEMMA 3. *[6] The strategy profile $(\frac{1}{2-v_R}x^* + \frac{1-v_R}{2-v_R}e_r, e_j)$ is an additive $\frac{1-v_R}{2-v_R}$-NE.*

3.2 Zero-sum games $(-R, R)$ and $(C, -C)$

Let (x^*, y^*) and $(\hat{x}, \hat{y})$ be NE in the zero-sum games $(-R, R)$ and $(C, -C)$, respectively. Let $u_R = (x^*)^T R y^*$ (which is minus the value of the zero-sum game $(-R, R)$), and let $u_C = \hat{x}^T C \hat{y}$ (which is the value of the zero-sum game $(C, -C)$). Assume, without loss of generality, that $u_R \geq u_C$.

LEMMA 4. *The strategy profile $(\hat{x}, y^*)$ is an additive $(1 - u_C)$-WSNE.*

4. NEW ALGORITHMS

Let (R, C) be a bimatrix game with all payoffs in $[0, 1]$. We consider the four zero-sum games $(R, -R)$, $(-C, C)$, $(-R, R)$, and $(C, -C)$, and we define the numbers v_R, v_C, u_R, and u_C in the way they were defined in Section 3.

4.1 Trading off between approximate NE and WSNE

In this section we give a polynomial-time algorithm that given an arbitrary bimatrix game as an input, outputs either an additive $\frac{1}{3}$-NE or an additive $\frac{1}{2}$-WSNE. This can be done by observing that if $v_R \leq \frac{1}{2}$ then Lemma 2 yields an additive $\frac{1}{2}$-WSNE, but if $v_R \geq \frac{1}{2}$ then Lemma 3 yields an additive $\frac{1}{3}$-NE.

THEOREM 5. *There is a polynomial-time algorithm that given a bimatrix game as input, outputs either an additive $\frac{1}{3}$-NE or an additive $\frac{1}{2}$-WSNE.*

Note that while the algorithm cannot guarantee which of the two kinds of approximate equilibria it will output, it either gives an output with a better approximation guarantee than the best algorithm for computing additive ε-NE [19], or better than the best algorithm for computing additive ε-WSNE [6].

4.2 Better approximate WSNE for symmetric payoff matrices

In this section we argue that additive $\frac{1}{2}$-WSNE can be computed in polynomial time in bimatrix games in which both payoff matrices are symmetric. Note that in the general case of arbitrary bimatrix games, the curently best approximation guarantee for additive ε-WSNE computable in polynomial time is $\varepsilon = 0.6528$ [11, 6].

The following theorem can be proved by considering four cases based on whether u_R and u_C are each at least $\frac{1}{2}$ or at most $\frac{1}{2}$, and applying Lemmas 2 or 4, as appropriate.

THEOREM 6. *If $u_R \geq v_R$ and $u_C \geq v_C$, then an additive $\frac{1}{2}$-WSNE can be found in polynomial time.*

The following lemma follows from the Minimax Theorem for zero-sum games.

LEMMA 7. *If the payoff matrix R is symmetric then $v_R = u_R$, and if the payoff matrix C is symmetric then $v_C = u_C$.*

Finally, the main result of this section follows by combining Lemma 7 and Theorem 6.

THEOREM 8. *There is a polynomial-time algorithm that given a bimatrix game in which both payoff matrices are symmetric as an input, computes an additive $\frac{1}{2}$-WSNE.*

REFERENCES

[1] Y. Anbalagan, H. Huang, S. Lovett, S. Norin, A. Vetta, and H. Wu. Large supports are required for well-supported Nash equilibria. In *Proceedings of the 18th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX'15)*, pages 78–84, 2015.

[2] Y. Anbalagan, S. Norin, R. Savani, and A. Vetta. Polylogarithmic supports are required for approximate well-supported Nash equilibria below 2/3. In *Proceedings of the 9th International Conference Web and Internet Economics (WINE'13)*, pages 15–23, 2013.

[3] Y. Babichenko, S. Barman, and R. Peretz. Simple approximate equilibria in large games. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC'14)*, pages 753–770, 2014.

[4] H. Bosse, J. Byrka, and E. Markakis. New algorithms for approximate Nash equilibria in bimatrix games. *Theoretical Computer Science*, 411(1):164–173, 2010.

[5] X. Chen, X. Deng, and S.-H. Teng. Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM*, 56(3), 2009.

[6] A. Czumaj, A. Deligkas, M. Fasoulakis, J. Fearnley, M. Jurdzinski, and R. Savani. Distributed methods for computing approximate equilibria. In *Proceedings of the 12th International Conference Web and Internet Economics (WINE'16)*, pages 15–28, 2016.

[7] A. Czumaj, M. Fasoulakis, and M. Jurdziński. Approximate well-supported Nash equilibria in symmetric bimatrix games. In *Proceedings of the 7th International Symposium on Algorithmic Game Theory (SAGT'14)*, pages 244–254, 2014.

[8] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.

[9] C. Daskalakis, A. Mehta, and C. Papadimitriou. Progress in approximate Nash equilibria. In *Proceedings of the 8th ACM Conference on Electronic Commerce (EC'07)*, pages 355–358, 2007.

[10] C. Daskalakis, A. Mehta, and C. Papadimitriou. A note on approximate Nash equilibria. *Theoretical Computer Science*, 410(17):1581–1588, 2009.

[11] J. Fearnley, P. W. Goldberg, R. Savani, and T. B. Sørensen. Approximate well-supported Nash equilibria below two-thirds. *Algorithmica*, 76(2):297–319, 2016.

[12] T. Feder, H. Nazerzadeh, and A. Saberi. Approximating Nash equilibria using small-support strategies. In *Proceedings of the 8th ACM Conference on Electronic Commerce (EC'07)*, pages 352–354, 2007.

[13] P. W. Goldberg and A. Pastink. On the communication complexity of approximate Nash equilibria. *Games and Economic Behaviour*, 85:19–31, 2014.

[14] S. Hémon, M. de Rougemont, and M. Santha. Approximate Nash equilibria for multi-player games. In *Proceedings of the 1st International Symposium on Algorithmic Game Theory (SAGT'08)*, pages 267–278, 2008.

[15] S. C. Kontogiannis, P. N. Panagopoulou, and P. G. Spirakis. Polynomial algorithms for approximating Nash equilibria of bimatrix games. In *Proceedings of the 2nd Workshop on Internet and Network Economics (WINE'06)*, pages 286–296, 2006.

[16] S. C. Kontogiannis and P. G. Spirakis. Well supported approximate equilibria in bimatrix games. *Algorithmica*, 57(4):653–667, 2010.

[17] R. Lipton, E. Markakis, and A. Mehta. Playing large games using simple strategies. In *Proceedings of the 4th ACM Conference on Electronic Commerce (EC'03)*, pages 36–41, 2003.

[18] A. Rubinstein. Setting the complexity of computing approximate two-player Nash equilibria. In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS'16)*, pages 258–265, 2016.

[19] H. Tsaknakis and P. G. Spirakis. An optimization approach for approximate Nash equilibria. *Internet Mathematics*, 5(4):365–382, 2008.