

How do you like your equilibrium selection problems? Hard, or very hard?

Paul W. Goldberg*

Dept. of Computer Science
University of Liverpool,
Ashton Street, Liverpool L69 3BX, UK.
P.W.Goldberg@liverpool.ac.uk

Abstract. The PPAD-completeness of Nash equilibrium computation is taken as evidence that the problem is computationally hard in the worst case. This evidence is necessarily rather weak, in the sense that PPAD is only known to lie "between P and NP", and there is not a strong prospect of showing it to be as hard as NP. Of course, the problem of finding an equilibrium that has certain sought-after properties should be at least as hard as finding an unrestricted one, thus we have for example the NP-hardness of finding equilibria that are socially optimal (or indeed that have various efficiently checkable properties), the results of Gilboa and Zemel [6], and Conitzer and Sandholm [3]. In the talk I will give an overview of this topic, and a summary of recent progress showing that the equilibria that are found by the Lemke-Howson algorithm, as well as related homotopy methods, are PSPACE-complete to compute. Thus we show that there are no short cuts to the Lemke-Howson solutions, subject only to the hardness of PSPACE. I mention some open problems.

1 Overview

There are two ways to view any algorithm for computing Nash equilibria. First, simply as a way to find a Nash equilibrium, one that is hopefully fast in practice, even though it may take exponential time in the worst case. Second, as a criterion for *equilibrium selection*, i.e. choosing some equilibrium that is considered to be preferable to others, in some sense a more plausible outcome. The latter viewpoint is especially relevant if the algorithm in question is somehow simple or decentralized. We can consider the problem of computing an equilibrium that is found by some specified algorithm, noting that we are not restricted to using that particular algorithm in order to find the equilibrium.

By way of example, in [7] we analyzed the classical Lemke-Howson algorithm for bimatrix games, in this context. The computational challenge is: Given a bimatrix game, find one of the solutions that could be computed using Lemke-Howson. Note that we are not asking about the complexity of the algorithm itself, which is already well-known to take exponential time in the worst case [9].

* Currently supported by EPSRC Grant EP/G069239/1

Just consider the complexity of finding the Lemke-Howson solutions. Of course, the problem is PPAD-hard, simply due to the PPAD-completeness of finding any equilibrium, but in fact we show that restricting to the Lemke-Howson solutions makes the problem PSPACE-complete, and thus in a sense even harder to compute than the restricted equilibria of [6, 3].

Homotopy methods The survey paper [8] discusses homotopy methods in detail, including the Lemke-Howson algorithm, which falls within this framework. The general idea is that in trying to solve a game \mathcal{G} , start by constructing a “starting game” \mathcal{G}_0 which is a version of \mathcal{G} where the numerical payoffs have been changed so that there is some “obvious” Nash equilibrium. Then consider a continuum of games that lie between \mathcal{G}_0 and \mathcal{G} . (The usual choice of a continuum of intermediate games has games whose payoffs are weighted averages of those in \mathcal{G}_0 and \mathcal{G} .) Within these games, there exists a continuous path of Nash equilibria that starts at the one for \mathcal{G}_0 and ends at an equilibrium of \mathcal{G} . Thus, we have specified a unique equilibrium of \mathcal{G} , and implicitly a natural path-following algorithm for finding it.

As an equilibrium selection theory, homotopy methods are attractive since the starting game \mathcal{G}_0 can be considered as representing some kind of “prior belief” about the behaviour of the other player(s). However, we showed in [7] that the equilibrium identified by this procedure is PSPACE-complete to compute, and moreover, the result extends to Lemke-Howson, in which the choice of initially dropped label corresponds to a particular choice of \mathcal{G}_0 .

Path-following algorithms The PSPACE-completeness result for Lemke-Howson solutions uses, quite intensively, the ideas developed in [4]. The problem END OF THE LINE that is used to characterise the complexity class PPAD, has a PSPACE-complete version in which you are required to compute the end-of-line obtained by following the path that begins at the known starting vertex of the graph. (Notice that this version is no longer (apparently) in NP since there is no obvious efficient test that a solution is correct.) This problem, called in [7] OTHER END OF THIS LINE, or OEOTL for short, is the one we reduce from. The proof proceeds by showing that a homotopy method —itself a path-following algorithm— captures the path-following approach to solving OEOTL.

The extension to Lemke-Howson requires us to design the game in such a way that all alternative solutions that might be produced, share features that efficiently encode a solution to generic instances of OEOTL.

Fictitious play We conclude with an interesting open problem that is inspired by the PSPACE-completeness results described above. Fictitious Play is a good algorithm to consider in the capacity of equilibrium selection. This simple and intuitive procedure (see e.g. [5] or indeed Wikipedia) is known to converge to Nash equilibrium under certain sufficient conditions, although for some games it fails to converge. See [2] for a detailed discussion of why it is a natural and appealing algorithm to consider. The following problem looks natural and simple to state:

Input. A bimatrix game \mathcal{G}

Question. Does Fictitious Play converge, when applied to \mathcal{G} ? Assume both players start at their first strategies.

Further question. Compute the equilibrium, if indeed FP converges.

The results of [1] indicate that a naive simulation of FP requires exponential time, but it does not rule out the possibility of “short cuts” alluded to in the abstract to this talk, in the context of Lemke-Howson. On the other hand, I know of no upper bound for the above questions, not even an exponential one.

References

1. Felix Brandt, Felix Fischer and Paul Harrenstein, On the Rate of Convergence of Fictitious Play, 3rd SAGT (these proceedings), (2010)
2. V. Conitzer, Approximation Guarantees for Fictitious Play, Proceedings of the 47th Annual Allerton Conference on Communication, Control, and Computing (Allerton-09), (2009) 636-643.
3. V. Conitzer and T. Sandholm, Complexity results about Nash equilibria, 18th International Joint Conference on Artificial Intelligence (IJCAI), (2003)
4. C. Daskalakis, P.W. Goldberg and C.H. Papadimitriou, The Complexity of Computing a Nash Equilibrium, *SIAM Journal on Computing*, 39(1), (2009), 195-259
5. D. Fudenberg and D.K. Levine, *The Theory of Learning in Games*, MIT Press (1998)
6. I. Gilboa and E. Zemel, Nash and correlated equilibria: Some complexity considerations, *Games and Economic Behavior*, 1 (1989), 80-93
7. P.W. Goldberg, C.H. Papadimitriou and R. Savani, The Complexity of the Homotopy Method, Equilibrium Selection, and Lemke-Howson Solutions, Arxiv technical report 1006.5352 (2010)
8. P.J.-J. Herings and R. Peeters, Homotopy methods to compute equilibria in game theory, *Economic Theory*, 42(1) (2010) 119-156
9. R. Savani and B. von Stengel, Hard-to-Solve Bimatrix Games, *Econometrica*, 74(2) (2006) 397-429