

Computing Shapley's Saddles

FELIX BRANDT, MARKUS BRILL, FELIX FISCHER, PAUL HARRENSTEIN, JAN HOFFMANN

Institut für Informatik, Ludwig-Maximilians-Universität München

Categories and Subject Descriptors: F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems; I.2.11 [Distributed Artificial Intelligence]: Multiagent Systems; J.4 [Computer Applications]: Social and Behavioral Sciences—*Economics*

General Terms: Theory, Algorithms, Economics

Additional Key Words and Phrases: Game Theory, Solutions Concepts, Shapley's Saddles, Computational Complexity

1. INTRODUCTION

Game-theoretic solution concepts, such as Nash equilibrium, are playing an ever increasing role in the study of systems of autonomous agents. A common criticism of Nash equilibrium is that its existence relies on the possibility of *randomizing* over actions, which in many cases is deemed unsuitable, impractical, or even infeasible.

In work dating back to the early 1950s Lloyd Shapley proposed ordinal set-valued solution concepts for zero-sum games that he refers to as *saddles* [Shapley, 1964]. Based on the elementary notions of dominance and stability, saddles are intuitively appealing, they always exist, and are unique in important classes of games. In this note, we survey recent results concerning the computational complexity of Shapley's saddles and identify some open problems [Brandt et al., 2009a,b].

2. PRELIMINARIES

A (finite) *game in normal-form* is a tuple $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$, where $N = \{1, 2, \dots, n\}$ is a nonempty finite set of *players* and for each player $i \in N$, A_i is a nonempty finite set of *actions* available to player i , and $p_i : (\prod_{i \in N} A_i) \rightarrow \mathbb{R}$ is a function mapping each action profile (i.e., combination of actions) to a real-valued *payoff* for player i .

A two-player game is called a *zero-sum* or *matrix game*, and can be represented by a single matrix M that contains the payoffs for the first player, if $p_2(a, b) = -p_1(a, b)$ for all $(a, b) \in A_1 \times A_2$. Γ_M will be used to denote the matrix game with matrix M . A two-player game is called *symmetric* if $A_1 = A_2$ and $p_1(a, b) = p_2(b, a)$ for all $a, b \in A_1$. Observe that Γ_M is symmetric if and only if M is *skew symmetric*, i.e., if $M^T = -M$. We will assume that games are given explicitly, i.e., as a table containing the payoffs for every possible action profile.

A *saddle point* of a matrix game Γ_M is a pair (a, b) of actions $a \in A_1, b \in A_2$ such that entry $M(a, b)$ is maximal in column b and minimal in row a of M . Saddle points

Authors' addresses: {brandt,brill,fischerf,harrenst,hoffmann}@tcs.ifi.lmu.de. This material is based on work supported by the Deutsche Forschungsgemeinschaft under grants BR-2312/6-1 and BR 2312/3-2.

are one of the earliest solution concepts proposed in game theory. They happen to coincide with the optimal outcome both players can guarantee in the worst case and thus enjoy a very strong normative foundation. Unfortunately, however, not every matrix game possesses a saddle point (the well-known Matching Pennies game is a counterexample). To remedy this situation, von Neumann considered *mixed*, i.e., randomized, strategies and proved that every matrix game has a mixed saddle point (or *equilibrium*) that moreover maintains the appealing normative properties of saddle points. The existence result was later generalized to arbitrary games by Nash, at the expense of its normative foundation. As mentioned earlier, requiring randomization in order to reach a stable outcome has been criticized on various grounds.

Shapley showed that existence of saddles (and even uniqueness in the case of matrix games) can also be guaranteed by moving to *minimal sets* of actions rather than randomizations over them [Shapley, 1964]. Let $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ be a game in normal-form and $S = (S_1, S_2, \dots, S_n)$ with $S_i \subseteq A_i$ for all $i \in N$. Shapley's saddles generalize saddle points by requiring that for every action a of a player $i \in N$ that is *not* included in S_i , there should be a reason for its exclusion, namely an action in S_i that dominates a with respect to $S_{-i} = \prod_{j \neq i} S_j$. To formalize this idea, we first define the notions of strict and weak dominance. For a player $i \in N$ and two actions $a, b \in A_i$,

a *strictly dominates* b with respect to S_{-i} if $p_i(a, s_{-i}) > p_i(b, s_{-i})$ for all $s_{-i} \in S_{-i}$, and

a *weakly dominates* b with respect to S_{-i} if $p_i(a, s_{-i}) \geq p_i(b, s_{-i})$ for all $s_{-i} \in S_{-i}$, with at least one strict inequality.

Then, S is a *generalized saddle point (GSP)* of Γ if for each player $i \in N$ and each $a \in A_i \setminus S_i$ there exists $s \in S_i$ such that s strictly dominates a with respect to S_{-i} . A *strict saddle* is a GSP that contains no other GSP. Similarly, S is a *weak generalized saddle point (WGSP)* of Γ if for each player $i \in N$ and each $a \in A_i \setminus S_i$, there exists $s \in S_i$ such that s weakly dominates a with respect to S_{-i} . A *weak saddle* is a WGSP that contains no other WGSP.

Strict saddles may be considered a “refinement” of iterated strict dominance as all strict saddles of a normal-form game are contained in the subgame that one obtains by iterated elimination of strictly dominated actions. Since strict dominance implies weak dominance, every strict saddle is a WGSP and thus contains a weak saddle. Consider for example the matrix game Γ_M with

$$M = \begin{pmatrix} 4 & 2 & 3 & 5 \\ 2 & 4 & 5 & 3 \\ 2 & 2 & 3 & 6 \\ 1 & 3 & 1 & 4 \\ 2 & 1 & 6 & 1 \end{pmatrix}$$

and actions r_1 through r_5 for the row player and c_1 through c_4 for the column player. The pair $S = (\{r_1, r_2, r_3\}, \{c_1, c_2\})$ is a strict saddle and a WGSP. Since both r_1 and r_2 weakly dominate r_3 with respect to $\{c_1, c_2\}$ and c_2 and c_1 weakly (and even strictly) dominate c_3 and c_4 , respectively, with respect to $\{r_1, r_2\}$, the pair $S' = (\{r_1, r_2\}, \{c_1, c_2\})$ is also a WGSP. Indeed, S' does not contain a smaller

WGSP and therefore is a weak saddle. Some reflection reveals that S and S' are in fact the *unique* strict and weak saddle of this game.

It is easy to see that every normal-form game has a strict and a weak saddle. By definition, $(A_1, A_2, \dots, A_n)$ is a GSP. Furthermore, every GSP that is not a saddle must contain a GSP that is strictly smaller. Finiteness implies that there exists a minimal GSP, i.e., a strict saddle. An analogous argument applies to the weak saddle. We finally note that both strict and weak saddle are *ordinal* solution concepts, i.e., they are invariant under order-preserving transformations of the payoff functions. This is in contrast to mixed-strategy Nash equilibrium, for which invariance holds only under positive *affine* transformations.

3. RESULTS AND OPEN PROBLEMS

Shapley [1964] has shown that every matrix game possesses a unique strict saddle by pointing out that the set of GSPs in such games is closed under intersection. He also describes an algorithm, attributed to Harlan Mills, to compute this saddle. Uniqueness ceases to hold in general two-player games. It turns out, however, that this does not have any serious consequences from a computational perspective: Mills' algorithm can be generalized to efficiently compute *all* strict saddles of an arbitrary n -player game, of which there can be only polynomially many.

THEOREM 1 [Brandt et al., 2009a]. *All strict saddles of an n -player game can be computed in polynomial time.*

The computation of weak saddles turns out to be significantly more complicated, a situation reminiscent of that for iterated weak and strict dominance. Somewhat surprisingly, it is not even known whether weak saddles can be computed efficiently in matrix games. A polynomial-time algorithm can however be given for the subclass of symmetric matrix games in which the two players get the same payoff if and only if they play the same action. Formally, a symmetric matrix game Γ_M with action set A is called a *confrontation game* if for all $a, b \in A$, $M(a, b) = 0$ if and only if $a = b$. Duggan and Le Breton [1996] have shown that confrontation games contain a unique weak saddle, which moreover contains the same actions for both players. It turns out that this weak saddle can be computed in polynomial time.

THEOREM 2 [Brandt et al., 2009a]. *The weak saddle of a confrontation game can be computed in polynomial time.*

In general, symmetric matrix games can have an exponential number of weak saddles. It follows immediately that even for this restricted class of games, computing *all* weak saddles requires exponential time in the worst case.

For general games, including those with only two players, a number of hardness results suggests that even a single weak saddle cannot be found in polynomial time. The first of these results states that deciding whether a particular action is contained in some weak saddle is hard for parallel access to NP, and thus not even in NP, unless the polynomial hierarchy collapses.

THEOREM 3 [Brandt et al., 2009b]. *Deciding whether a given action is contained in some weak saddle of a game is Θ_2^P -hard.*

A gap remains between the lower bound of Θ_2^p and the best known upper bound of Σ_2^p for the same problem.

The same techniques that lead to the previous theorem can be used to show that it is NP-complete to decide whether a game has a nontrivial weak saddle, i.e., a weak saddle that does not contain all actions.

THEOREM 4 [Brandt et al., 2009b]. *Deciding whether a game has a nontrivial weak saddle is NP-complete.*

This result has a number of interesting consequences.

COROLLARY 1. *Deciding whether a game has a unique weak saddle is coNP-hard. Recognizing a weak saddle is coNP-complete. Finding a weak saddle is NP-hard under polynomial-time Turing reductions.*

The last result leads to an interesting open problem: No stronger lower bound can be derived using the same technique, by reducing from an NP-complete decision problem, yet it is not clear how a weak saddle should be found in nondeterministic polynomial time.

Analogous computational problems can also be defined for *very weak saddles*, which use an even weaker version of dominance that does not require a strict inequality. While these problems are not connected in an obvious way to their counterparts for weak saddles, most of the above results can be shown to apply to very weak saddles as well [Brandt et al., 2009b]. Similarly, some of the above results have been extended to *mixed saddles*, which are defined using dominance by mixed strategies [Duggan and Le Breton, 2001, Brandt et al., 2009a].

REFERENCES

- F. Brandt, M. Brill, F. Fischer, and P. Harrenstein. Computational aspects of Shapley's saddles. In *Proceedings of the 8th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 209–216, 2009a.
- F. Brandt, M. Brill, F. Fischer, and J. Hoffmann. The computational complexity of weak saddles. In M. Mavronicolas and V. G. Papadopoulos, editors, *Proceedings of the 2nd International Symposium on Algorithmic Game Theory (SAGT)*, volume 5814 of *Lecture Notes in Computer Science (LNCS)*, pages 238–249. Springer-Verlag, 2009b.
- J. Duggan and M. Le Breton. Dutta's minimal covering set and Shapley's saddles. *Journal of Economic Theory*, 70:257–265, 1996.
- J. Duggan and M. Le Breton. Mixed refinements of Shapley's saddles and weak tournaments. *Social Choice and Welfare*, 18(1):65–78, 2001.
- L. Shapley. Some topics in two-person games. In M. Dresher, L. S. Shapley, and A. W. Tucker, editors, *Advances in Game Theory*, volume 52 of *Annals of Mathematics Studies*, pages 1–29. Princeton University Press, 1964.