A Note on the Computational Hardness of Evolutionary Stable Strategies
Draft

Noam Nisan

Abstract

We present a very simple reduction that when given a graph G and an integer k produces a game that has an evolutionary stable strategy if and only if the maximum clique size of G is not exactly k. Formally this shows that existence of evolutionary stable strategies is hard for a complexity class called $co - D^p$, slightly strengthening (and greatly simplifying) the known NP-hardness and co-NP-hardness. En route we show that even recognizing an evolutionary stable strategy is co-NP complete.

1 Introduction

The concept of evolutionary stable strategy (ESS) is a refinement of Nash equilibrium that attempts capturing the notion that small deviations from it must be strictly inferior to it. The idea is that such an equilibrium will be "evolutionary stable", since any small "mutation" will die out. This concept, introduced in [SP73], has been since heavily used in explaining evolutionary phenomena – see the classic [Smi82].

Given is a square payoff matrix u, where $u(i, j)$ denotes the "payoff" for type i when facing type j. We are interested in an ESS in the *mixed game*, i.e. we will define when a probability distribution x on the types is an ESS (both the rows and the columns of u are indexed by the types). For probability distributions x and y on the types, we denote the expected payoff of x against y by $u(x, y) = \sum_i \sum_j x_i y_j u(i, j)$.

Definition 1 *A mixed ESS of the payoff matrix $u(\cdot, \cdot)$ is probability distribution x such that:*

- **Symmetric Nash Equilibrium:** *x is a best response to itself, i.e. for every y, $u(x, x) \geq u(y, x)$.*
- **Maynard Smith's 2nd condition**: *For every $y \neq x$ such that $u(y, x) = u(x, x)$ we have that $u(y, y) < u(x, y)$.*

Two questions of computational complexity naturally present themselves:

1. What is the computational complexity of recognizing whether a given x is an ESS of a given matrix u?
2. What is the computational complexity of finding an ESS of a given matrix u, or even of determining whether u has an ESS?

Unlike the situation with regular Nash equilibria, an ESS does not exist for every game, nor is it clear how to efficiently recognize one. Once one observes – as for completeness we spell out in the appendix – that the quantification over x and y need only be over vectors with rational coefficients of reasonable size, the definition directly implies that recognizing an ESS is in $co - NP$, and thus that finding an ESS is in Σ_2 (or, more precisely, reasonable decision variants of finding an ESS are in Σ_2). Etessami and Lochbihler [EL04] studied the second question and showed, surprisingly, that deciding whether a given matrix u has an ESS is both NP-hard and $co - NP$-hard. In this paper we give a very simple and elementary proof of a slightly stronger result, and en-route settle the first question.

We give a very simple reduction that when given a graph G and an integer k produces a payoff matrix u that has an ESS if and only if the maximum clique size of G is not exactly k. The problem of having maximum clique size *exactly* k is known to be complete for a complexity class called D^p defined in [PY82], we thus get:

Theorem 1: Existence of a mixed ESS is $co - D^p$-hard.

As D^p contains both NP and co-NP, hardness for these classes is directly implied. Alternatively, readers who wish to remain ignorant of the definition of D^p, may directly convince themselves that the problem of the maximum clique size being exactly k is both NP-hard and $co - NP$-hard[1], and directly get:

Corollary 1 [EL04]: Existence of a mixed ESS is NP-hard and $co - NP$-hard.

Our proof also settles the first question:

Theorem 2: Recognizing a mixed ESS is $co - NP$-complete.

This result may be viewed as somewhat critical of the whole notion of ESS as even recognizing one can not be done effectively. The main open problem still remains the determination of the exact complexity of finding an ESS.

2 The reduction

Given a graph $G = (V, E)$, $|V| = n$, and an integer $1 < k < n$, we will build a $(n + 1) \times (n + 1)$ payoff matrix u that will have an ESS if and only if the largest clique size of G, is not exactly k. u's rows and columns will be the vertices of V, named $1..n$, with an additional row and column named 0.

- for $1 \le i \ne j \le n$: $u(i, j) = 1$ if $(i, j) \in E$, and $u(i, j) = 0$ if $(i, j) \notin E$.
- for $1 \le i \le n$: $u(i, i) = 1/2$.
- for $1 \le i \le n$: $u(0, i) = u(i, 0) = a = 1 - 1/(2k)$.

[1] $co-NP$-hardness is easy by simply adding a k-clique to a "at-most-k-clique" problem instance. For NP-hardness, it turns out that the standard reduction of max-clique from 3-SAT also implies this since in it $(k+1)$-cliques can never occur anyway. (Each of the m clauses of the $3-CNF$ is converted to 7 vertices – one for each satisfying assignment of the clause – and "compatible" clauses are connected by an edge. An m-clique corresponds to a satisfying assignment, and no $(m + 1)$-cliques are possible since vertices within a single clause are never inter-connected.)

In the rest of the paper, x (and y) will always be probability distributions on $\{0...n\}$, i.e. vectors of length $n+1$, with $x_i \geq 0$ for all $0 \leq i \leq n$ and $\sum_i x_i = 1$.

Before we proceed with the proof, we will require a slight variant of a lemma of Motzkin and Straus, whose original version was used also in [EL04]:

Lemma 1 *(Essentially [MS65][2]): For every x, with $x_0 = 0$, $u(x,x) \leq 1 - 1/(2k')$, where k' is the size of the maximum clique in G. Equality is achieved if and only if x is uniform over a k'-clique.*

For completeness we give the simple proof.

Proof: Simple calculations will reveal that $1-1/(2k')$ is indeed achieved for a uniform distribution on a k'-clique. We can prove the lemma by induction on the number of non-edges between vertices in the support of x. The base of the induction is when the support of x is a k''-clique ($k'' \leq k'$), and then $u(x,x) = 1 - \sum_i x_i^2/2 \leq 1 - 1/(2k'') \leq 1 - 1/(2k')$, with equality holding only if x is uniform over a k'-clique. (The inequality is since for x with $||x||_1 = 1$, the value of $||x||_2$ is minimized exactly when all components are equal.)

For the induction step consider two vertices i_1, i_2 in the support of x with $(i_1, i_2) \notin E$. Denote $p = \sum_{(j,i_1)\in E} x_j$ and $q = \sum_{(j,i_2)\in E} x_j$, and without loss of generality assume $p \geq q$. We can now define a new vector x' by moving all the probability from i_2 to i_1 (i.e. $x'_{i_1} = x_{i_1} + x_{i_2}$ and $x'_{i_2} = 0$, and other indices unchanged.) Note that i_2 is no longer in the support of x' and thus the induction hypothesis applies to it. It may be easily verified that $u(x',x') = u(x,x) + x_{i_2}(p-q) + x_{i_1}x_{i_2} > u(x,x)$ and this completes the induction step. ∎

We will now prove the properties of the reduction:

Claim 1 *If C is a maximal clique of G of size $k' > k$, and x is the uniform distribution on C, then x is an ESS.*

Proof: Notice that $u(x,x) = 1 - 1/(2k')$, however, $u(0,x) = a < u(x,x)$, and that since C is maximal, for every $i \notin C$, $u(i,x) \leq 1 - 1/k' < u(x,x)$. Thus x is best response to itself, and all best responses y must be supported on C. We now need to verify the second condition for $y \neq x$'s that are supported on C. But for such y, using the lemma, $u(y,y) < 1 - 1/(2k')$ while $u(x,y) = 1 - 1/(2k')$. ∎

Claim 2 *If G contains no clique of size k then the pure strategy 0 is an ESS.*

Proof: The first condition is trivial since for every i, $u(i,0) = a = u(0,0)$. For the second condition, it suffices to verify that for every y which is not the pure 0 strategy, $u(y,y) < u(0,y)$.

We first note that wlog we can assume that $y_0 = 0$. Denote y^* to be the vector y with the 0'th coordinate zeroed, renormalized. I.e. for $i \neq 0$, $y^*_i = y_i/(1-y_0)$. We have that $u(y,y) = (2y_0 - y_0^2)a + (1-y_0)^2 u(y^*, y^*)$, and $u(0,y) = a$, thus the second condition holds for y iff it holds for y^*, hence we can assume wlog that $y_0 = 0$.

Now we apply the lemma, and since G has no cliques of size k, $u(y,y) \leq 1 - 1/2(k-1) < a = u(0,y)$. ∎

[2] In [MS65], the diagonal entries $u(i,i)$ were defined to be 0, and consequently the bound was $1 - 1/k'$, and could be achieved even for distributions whose support was not a clique.

Claim 3 *If the maximum size clique of G is at least k then the pure strategy 0 is not an ESS.*

Proof: While the first condition does trivially hold, the second condition fails for y which is a uniform distribution on a maximum size clique of G of size $k' \geq k$. In this case $u(y, y) = 1 - 1/(2k') \geq a = u(0, y)$. ∎

Claim 4 *If the maximum size clique of G is at most k then u has no ESS except, perhaps, the pure 0 strategy.*

Proof: Let x which is not the pure strategy 0, satisfy the first condition. As before let x^* be x with the 0'th coordinate zeroed and then renormalized. I.e. for $i \neq 0$, $x_i^* = x_i/(1 - x_0)$. Now, $u(x, x) = (2x_0 - x_0^2)a + (1 - x_0)^2 u(x^*, x^*)$. Using the lemma we bound $u(x^*, x^*) \leq 1 - 1/(2k) = a$ and thus $u(x, x) \leq a$. Now the second condition fails for y which is the pure strategy 0, which is also a best reply to x, but $u(y, y) = a = u(x, y)$. ∎

We can now easily conclude the proofs of the theorems. Claims 1 and 2 show that if the maximal clique size of G is not exactly k then an ESS exists. Claims 3 and 4 show that if the maximal clique size of G is exactly k then no ESS exists. This proves theorem 1. Claims 2 and 3 show that the pure strategy 0 is an ESS if and only if the maximal clique size of G is less than k, this gives $co - NP$-hardness and thus theorem 2.

Acknowledgements

I thank Nati Linial, Sergiou Hart, and Alex Samorodnitski for helpful discussions and Michal Feldman for comments on a previous draft.

Bibliography

[SP73] J. Maynard Smith and G. P. Price. *The logic of animal conflict*, Nature, 246:15-18, 1973.

[MS65] T. S. Motzkin and E. G. Straus. *Maxima for graphs and a new proof of a theorem of Turan.* Canadian Journal of Mathematics, 17:533-540, 1965.

[Smi82] J. Maynard Smith, *Evolution and the theory of games.* Cambridge university press, 1982.

[EL04] K. Etessami and A. Lochbihler. *The computational complexity of evolutionary stable strategies.* ECCC technical report TR04-055, 2004.

[PY82] C. H. Papadimitriou and M. Yannakakis. *The complexity of facets (and some facets of complexity)*, 14th STOC, pages 255-260, 1982.

Appendix: Rationality of x and y

While in the definition of an ESS, an ESS could have irrational coefficients, we show why for any payoff matrix u with rational coefficients, all coefficients of every ESS are rational numbers, with size (bit-length of the numerators and denominators) which is polynomially bounded in the size of the input coefficients. Similarly, we show that the quantification over y in the second condition need only range over y's with small rational coefficients as a counter example, if it exists, can be found there. In fact, in both cases we show that the values of x and of y can be determined efficiently once their support is known – and the solution is obtained by solving a set of linear equations.

Let us first consider choosing x. For x to be a symmetric Nash equilibrium (i.e. to satisfies the first condition), with support exactly the set S of rows (and columns), it most be a solution of the linear program that specifies that all pure strategies $i \in S$ are best response to x, i.e. $u(i,x) \geq u(j,x)$ for every $i \in S$ and j (as well as that x is a probability distribution, i.e. $x_i \geq 0$ and $\sum_i x_i = 1$). If the solution to this LP is unique then it can be found efficiently and is obtained at a vertex of the polytope and is thus rational. If the solution is not unique, say $x \neq y$ are both solutions, then x can not be an ESS since the second condition will fail as y is a best response both to x and to itself.

Now consider the situation that we have rational u and x and need to find a y that is a counter example to the second condition. If such a y exists then clearly the "worst case" is at the maximum of $u(y,y) - u(x,y)$ over the y's that satisfy $u(y,x) = u(x,x)$. Now finding this maximum is computationally difficult, but once we know the support S of y, it can be done efficiently. First the support must be a subset of the set of pure best replies to x (this is equivalent to $u(y,x) = u(x,x)$). Second, notice that we are looking here for the maximum of a quadratic form. Once the support is known to be exactly S, it means that it is not obtained at the border (i.e. at locations where the inequalities of the form $y_i \geq 0$ for $i \in S$ are saturated). Thus the maximum must satisfy the condition that all partial derivatives (relative to y_i for $i \in S$) are zero (the equality $\sum_i y_i = 1$ should first be eliminated by setting $y_n = 1 - \sum_{i<n} y_i$). Since the derivatives are linear, this is just a solution of linear equations.