Methods for Finding Leader-Follower Equilibria with Multiple Followers

(Extended Abstract)

Nicola Basilico
University of Milan
Milan, Italy
nicola.basilico@unimi.it

Stefano Coniglio
University of Southampton
Southampton, UK
s.coniglio@soton.ac.uk

Nicola Gatti
Politecnico di Milano
Milan, Italy
ngatti@elet.polimi.it

ABSTRACT

Leader-follower (LF) equilibria play a central role in several applications of game theory. In spite of this, the literature only presents sporadic results for the case with two or more followers. In this work, we address the problem of computing LF equilibria in this setting, assuming that the followers play a Nash Equilibrium after the leader's commitment.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Multi-agent systems

Keywords

Game Theory (cooperative and non-cooperative); Equilibrium computation; Multi-Follower Stackelberg Games; Experimentation

1. INTRODUCTION

Leader-Follower equilibria (LFEs) have been much studied for the case of a single follower who, in an equilibrium, is known to play w.l.o.g. a pure strategy [6]. In this case, computing an LFE is easy with complete information, whereas it becomes $\mathcal{FNP}$-hard for Bayesian games [2]. In the case of multiple followers, the literature shows that, if the followers play a correlated equilibrium, an LFE can be found in polynomial time [1] whereas, if they play sequentially, the problem is $\mathcal{FNP}$-hard [2].

In this work, we address the fundamental case where, after the leader's commitment, the followers play a Nash Equilibrium (NE). It is relevant in many applications, such as *social planning* (e.g., urban traffic plans) or *monetary economics* (e.g., monetary policies of central banks). We refer to the corresponding LFE as *Leader-Follower Nash Equilibrium* (LFE-N). Depending on the leader's assumption on the followers' compliance, we define an *optimistic* and a *pessimistic* variant of the problem where the followers select an NE which either maximizes or minimizes the leader's utility.

2. PROBLEM ANALYSIS

We outline few hardness results of LFE-Ns computation:

Appears in: *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016), J. Thangarajah, K. Tuyls, C. Jonker, S. Marsella (eds.), May 9–13, 2016, Singapore.*

PROPOSITION 1. *Computing an optimistic or pessimistic LFE-N is $\mathcal{FNP}$-hard and it is not in Poly-$\mathcal{APX}$ unless $\mathcal{P} = \mathcal{NP}$, even for polymatrix games.*

PROPOSITION 2. *Deciding whether one of the leader's actions is played with strictly positive probability in an optimistic LFE-N is $\mathcal{NP}$-hard.*

In general, the computation of an LFE-N amounts to solving a *bilevel program*. In the first level, we look for a leader's strategy δ while, in the second level and for the given δ, we look for followers' strategies $\rho_1, \rho_2, \ldots, \rho_{n-1}$ forming an NE which either maximizes (optimistic case) or minimizes (pessimistic case) the leader's utility. If we assume the convexity of the second level problem, bilevel programs can be cast as (compact) single-level programs by substituting for the second level their KKT conditions [4]. This is not needed for optimistic LFE-Ns, as we can turn the second level problem into one of pure feasibility over which the leader has full control: he looks for a strategy δ and, given δ, also for an NE in the followers' game such that his utility is maximized. This allows for the exact solution of the problem via (nonlinear) mathematical programming. For pessimistic LFE-Ns, we cannot get rid of the second level objective function as the leader cannot control which NE the followers choose. Moreover, KKT conditions do not yield a compact reformulation, as even the sole feasible region of the second level (the set of NEs of a game parameterized by δ) is highly nonconvex. For this case, we investigate the use of (heuristic) black-box optimization techniques, which only assure a lower bound on the leader's utility in an optimal LFE-N.

3. COMPUTING LFE-NS

In the following, we consider both Normal-Form (NF) and PolyMatrix (PM) games, randomly generated with GAMUT.

3.1 Optimistic variant

To compute optimistic LFE-Ns, we propose different exact mathematical programming formulations, tailored for both NF and PM games, which we test with two global optimization solvers, SCIP and BARON. Let us first introduce our notation. Let $N = \{1, \ldots, n\}$ be the set of agents and, for each $i \in N$, let A_i be the corresponding set of actions, with $m_i = |A_i|$. For each agent $i \in N$, let $x_i \in [0,1]^{m_i}$, with $e^T x_i = 1$ (where e is the all-one vector), be his *strategy*, where each component x_i^a represents the probability that action $a \in A_i$ is played. Let $U_i \in \mathbb{R}^{m_1 \times \ldots \times m_n}$ denote, for each agent $i \in N$, his multidimensional payoff matrix, where each component $U_i^{a_1,\dots,a_n}$ is the utility of agent i when the agents play actions $a_1,\dots,a_n$. Let agent n be the leader, whom we relabel as ℓ, and $F = N\backslash\{\ell\}$ be the set of followers. While we assume $n = 3$ ($F = \{1, 2\}$) and $m_\ell = m_1 = m_2$, the generalization to $n > 3$ and general m_is is straightforward. For all $f \in F$, we adopt the notation $f' := F \setminus \{f\}$ and denote x_ℓ, x_1, x_2 by δ, ρ_1, ρ_2, respectively. The following formulation is the one which turned out to be the most efficient (when solved with SCIP):

$$\max_{\substack{\delta,\rho_1,\rho_2,v_f\geq 0\\ u_f,r_f,y_f,z\geq 0\\ s_f\in\{0,1\}}} \quad \sum_{i\in A_\ell}\sum_{j\in A_1}\sum_{k\in A_2} z^{ijk}U_\ell^{ijk} \quad \text{s.t.} \tag{1}$$

$$u_f^j = \sum_{i\in A_\ell}\sum_{k\in A_{f'}} y_{f'}^{ik}U_f^{ijk} \quad \forall f\in F, j\in A_f \tag{2}$$

$$y_f^{ij} = \delta^i\rho_f^j \;\; \forall i\in A_\ell, f\in F, j\in A_f \tag{3}$$

$$z^{ijk} = y_1^{ij}\rho_2^k \;\; \forall i\in A_\ell, j\in A_1, k\in A_2 \tag{4}$$

$$\sum_{i\in A_\ell}\sum_{j\in A_f} y_f^{ij} = 1, y_f\geq 0 \quad \forall f\in F \tag{5}$$

$$\sum_{i\in A_\ell}\sum_{j\in A_1}\sum_{k\in A_2} z^{ijk} = 1, z\geq 0 \tag{6}$$

$$e^T\delta = 1, e^T\rho_1 = 1, e^T\rho_2 = 1 \tag{7}$$

$$v_f \geq u_f^j \quad \forall f\in F, j\in A_f \tag{8}$$

$$r_f^j = v_f - u_f^j \quad \forall f\in F, j\in A_f \tag{9}$$

$$\rho_f^j \leq 1 - s_f^j \quad \forall f\in F, j\in A_f \tag{10}$$

$$r_f^j \leq M_f s_f^j \quad \forall f\in F, j\in A_f. \tag{11}$$

Overall, the formulation contains $2m^2+m^3$ quadratic constraints and a linear objective function. Solving it with SCIP, we obtain average optimality gaps $< 30\%$ ($< 35\%$ in the worst-case) for $m \leq 40$. This formulation may thus constitute a valid (empirical) approximation algorithm, yielding a $\frac{6.5}{10}$-approximation.

The average (over 10 random instances) computing times, which we report as a function of n and m for both NF and PF games in the following two tables, show that we can tackle instances comparable, in size, to the largest ones used in [5] for NE finding algorithms, in spite of our problem being more general. They also show how, in the PF case, the formulation (which contains fewer nonlinearities in this case) allows for a substantial reduction of the computing times.

Normal-form games

	$m=2$	$m=3$	$m=4$	$m=5$
$n=3$	0.06	0.20	0.92	23.79
$n=4$	0.19	8.274	142.66	1304.45
$n=5$	278.06	409.78	2016.97	–
$n=6$	172.90	2350.95	2212.95	–

Polymatrix games

	$m=5$	$m=6$	$m=7$	$m=8$	$m=9$	$m=10$
$n=3$	0.24	2.17	1.87	7.31	24.45	194.71
$n=4$	4.84	10.85	121.57	247.84	622.72	1947.54
$n=5$	7.51	90.83	332.04	1982.77	2396.01	2175.29
$n=6$	10.31	1169.50	2062.75	–	–	–

3.2 Pessimistic variant

For the pessimistic case, we introduce a Black-Box approach (BB) based on a Radial Basis Function (RBF) estimation, relying on the solver RBFOpt [3]. The idea is of exploring the leader's strategy space (variables δ) with a direct search which, iteratively, builds an RBF approximation of the objective function, relying on the solution of an *oracle formulation* for the objective function evaluation. Given an incumbent value $\hat{\delta}$, the oracle solves the (NF or PM) pessimistic second level problem (where the leader's utility is minimized) with the following exact mathematical programming formulation, with δ fixed to $\hat{\delta}$:

$$\min_{\substack{\rho_1,\rho_2,v_f\geq 0\\ u_f,r_f,y\geq 0\\ s_f\in\{0,1\}}} \quad \sum_{i\in A_\ell}\sum_{j\in A_1}\sum_{k\in A_2} \hat{\delta}^i y^{jk}U_\ell^{ijk} \quad \text{s.t.} \tag{12}$$

$$u_f^j = \sum_{i\in A_\ell}\sum_{k\in A_{f'}} \hat{\delta}^i\rho_{f'}^k U_f^{ijk} \quad \forall f\in F, j\in A_f \tag{13}$$

$$y^{jk} = \rho_1^j\rho_2^k \quad \forall j\in A_1, k\in A_2 \tag{14}$$

$$\sum_{j\in A_1}\sum_{k\in A_2} y_f^{jk} = 1 \tag{15}$$

$$\text{Constraints (7)–(11).} \tag{16}$$

In Fig. 1 (a), we compare BB (also in the optimistic setting) to the exact (optimistic) formulation. On the instances with $m \leq 5$ (for which, being solved to optimality by SCIP, we know the optimal solution value), we see that the gap between optimistic and pessimistic LFE-Ns is rather small, suggesting that, in the pessimistic case, the leader could force the followers to play a strategy providing him with a utility not dramatically smaller than that which he would obtain in the optimistic setting. The figure also shows the limit of the approach, indicating that, for $m \geq 15$, the time needed to solve the oracle formulation becomes too high for the method to be practical.

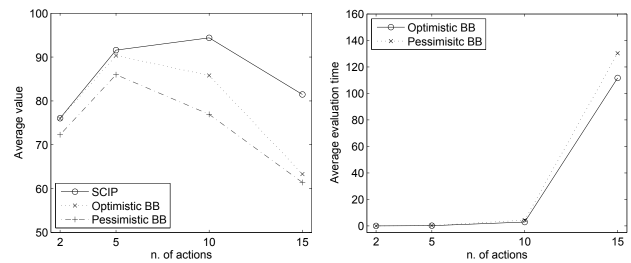

(a) Average objective (b) Average oracle time

Figure 1: Black-Box approach.

REFERENCES

[1] V. Conitzer and D. Korzhyk. Commitment to correlated strategies. In *Proc. of the 25th AAAI Conference on Artificial Intelligence AAAI*, 2011.

[2] V. Conitzer and T. Sandholm. Computing the optimal strategy to commit to. In *ACM Conference on Electronic Commerce*, pages 82–90, 2006.

[3] A. Costa, G. Nannicini, T. Schroepfer, and T. Wortmann. Black-box optimization of lighting simulation in architectural design. In *Complex Systems Design & Management Asia*, pages 27–39. 2015.

[4] S. Dempe. *Foundations of bilevel programming.* Springer Science & Business Media, 2002.

[5] R. Porter, E. Nudelman, and Y. Shoham. Simple search methods for finding a nash equilibrium. *Games and Economic Behavior*, 63(2):642–662, 2008.

[6] B. von Stengel and S. Zamir. Leadership games with convex strategy sets. *Games and Economic Behavior*, 69:446–457, 2010.