Bribery in Multiple-Adversary Path-Disruption Games is Hard for the Second Level of the Polynomial Hierarchy[*]

(Extended Abstract)

Adrian Marple
Stanford University
Stanford, CA 94305, USA
adrian.marple@gmail.com

Anja Rey and Jörg Rothe
Heinrich-Heine-Universität Düsseldorf
40225 Düsseldorf, Germany
{rey, rothe}@cs.uni-duesseldorf.de

ABSTRACT

Path-disruption games, a class of cooperative games introduced by Bachrach and Porat [1], model situations where the players, sitting on the vertices of a given graph, try to prevent – by blocking all possible paths – their adversaries from traveling from a set of source vertices to a set of target vertices. Rey and Rothe [3] studied bribery in these games and showed that when costs are assigned to the vertices, the corresponding problem is NP-complete in the single-adversary case, and is in $\Sigma_2^p = \text{NP}^{\text{NP}}$, the second level of the polynomial hierarchy, in the multiple-adversary case. They left open whether the latter problem is Σ_2^p-complete. In this note, we solve this open question in the affirmative.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—Multiagent Systems; J.4 [**Computer Applications**]: Social and Behavioral Sciences—Economics

General Terms

Economics, Theory

Keywords

Game theory, path-disruption games, bribery, complexity

1. INTRODUCTION

Bachrach and Porat [1] introduced path-disruption games and studied their properties. Informally stated, a *single-adversary path-disruption game* is a cooperative game played on a graph whose vertices represent the players and some adversary seeks to reach some target vertex from some source vertex. To prevent this from happening, the players try to form coalitions that block all paths for the adversary, and if some coalition of players is successful, it wins the game. A *multiple-adversary path-disruption game* refers to the same setting, yet with multiple adversaries, each seeking to reach some target vertex from some source vertex.

Rey and Rothe [3] (see also [4]) introduced the notion of bribery in path-disruption games and studied the complexity of the corresponding problems. In this setting, the adversaries try to bribe

some of the players such that (a) they do not exceed their budget, and (b) no blocking coalition will be formed. Rey and Rothe [3] show that this problem is NP-complete for a single adversary. For the case of multiple adversaries, they have shown that this problem is in $\Sigma_2^p = \text{NP}^{\text{NP}}$, the second level of the polynomial hierarchy [2, 5], and they have suspected that it is even complete for this class. In this note, we solve this open question in the affirmative by establishing a Σ_2^p-hardness lower bound for the bribery problem in multiple-adversary path-disruption games with costs.

2. OUR RESULT

We consider path-disruption games (PDGs) with multiple adversaries and costs (PDGC-MULTIPLE), the most general case of PDGs that Bachrach and Porat [1] introduced. Such games are played on an undirected graph $G = (V, E)$ with $n = |V|$ vertices, where $v_i \in V$ represents player $i \in N = \{1, \ldots, n\}$. Each of the m adversaries is associated with a pair (s_j, t_j) of source and target vertices in V, the jth adversary seeking to travel from s_j to t_j. A coalition $C \subseteq N$ *blocks a path from s_j to t_j* if there is no path from s_j to t_j in the subgraph $G|_{V \smallsetminus \{v_i \mid i \in C\}}$ of G induced by $V \smallsetminus \{v_i \mid i \in C\}$ (or if s_j or t_j are not in $V \smallsetminus \{v_i \mid i \in C\}$). Moreover, given a cost function $c : V \to \mathbb{R}_{>0}$ and a reward $r \in \mathbb{R}_{\geq 0}$, we define the coalitional function $v : 2^N \to \mathbb{R}$ of the game as follows: For $C \subseteq N$, let $\tilde{v}(C) = 1$ if C blocks every path from s_j to t_j for each j, $1 \leq j \leq m$, and $\tilde{v}(C) = 0$ otherwise; let $\mu(C) = \min\{\sum_{i \in B} c(v_i) \mid B \subseteq C \wedge \tilde{v}(B) = 1\}$, and $\mu(C) = \infty$ otherwise; and let $v(C) = r - \mu(C)$ if $\mu(C) < \infty$, and $v(C) = 0$ otherwise. Define bribery in PDGC-MULTIPLE by:

PDGC-MULTIPLE-BRIBERY	
Given:	A PDGC-MULTIPLE instance (N, v) on a graph G, a price function $\pi : V \to \mathbb{Q}_{\geq 0}$, and a budget $K \in \mathbb{Q}_{\geq 0}$.
Question:	Is there a coalition $B \subseteq N$ such that $\sum_{i \in B} \pi(v_i) \leq K$, and no coalition $C \subseteq N \smallsetminus B$ has a value $v(C) > 0$?

THEOREM 1. *PDGC-MULTIPLE-BRIBERY is Σ_2^p-complete.*

PROOF. The reduction is from the well-known Σ_2^p-complete problem QBF$_2$, which asks whether a given quantified boolean formula $F = (\exists X)(\forall Y) f(X, Y)$ is valid, where X is a set of p boolean variables, Y is a set of q boolean variables, and $f(X, Y)$ is a disjunction of k implicants, $f(X, Y) = \bigvee_{i=1}^{k} (u_i \wedge v_i \wedge w_i)$, and u_i, v_i, and w_i, $1 \leq i \leq k$, are literals over $X \cup Y$. That is, we ask whether there exists an assignment to the variables of X such that for all variable assignments to Y, f evaluates to true. Note that we assume that every implicant has exactly three literals. The graph G for the game that is part of the PDGC-MULTIPLE-BRIBERY instance to be constructed from F, is built from the three graphs, G_1, G_2, and G_3,

(a) Graph G_1 for X variables   (b) Graph G_2 for Y variables

(c) Graph G_3 for implicant $(u_i \wedge v_i \wedge w_i)$ in f

Figure 1: Three graphs for the reduction proving Theorem 1

shown in Figure 1. In particular, G is constructed from G_1, G_2, and G_3 by identifying, for each occurrence of a literal u_i, v_i, or w_i in f, the vertex in G_3 representing this literal with the vertex representing the corresponding variable ($x \in X$ or $y \in Y$) or its negation ($\neg x$ or $\neg y$) in G_1 or G_2. The players on vertices in G_1 labeled with X variables or their negations are bribable for a price of 1 but have 0 cost (they are free to participate in a coalition if not bribed), sources s_j and targets t_j have a cost of $6k + q + 1$ and a price of $p + 1$, and all other vertices have cost 1 and a price of $p + 1$. Let $K = p$ be the briber's budget and $r = 6k + q + 1$ be the reward. Thus, the briber can bribe up to p players having price 1 (note that all other players – those not on vertices in G_1 labeled with X variables or their negations – are too expensive to bribe). A coalition C can only have a positive value if it contains a successful subcoalition with at most $6k + q$ players other than those on vertices labeled $x \in X$ or $\neg x$ that are not bribed, and with no player on a source or target vertex (i.e., $\mu(C) < r$). Intuitively, the purpose of graph G_1 in Figure 1(a) is to enforce consistency on the part of the briber, the purpose of G_2 in Figure 1(b) is for consistency of the coalition, and the purpose of G_3 in Figure 1(c) is to enforce the implicants.

Consider the case where the briber does not play consistently, i.e., either plays (1) neither x nor $\neg x$ or (2) both x and $\neg x$, for an $x \in X$. In case (2), since the number of bribable players is limited to p, there is some other $x' \in X$ such that case (1) holds. In case (1), a coalition C consisting of the players corresponding to x and $\neg x$, to any consistent assignment to the variables in Y, and to vertices a_i, a_i', b_i, b_i', c_i, and c_i' for all i, $1 \le i \le k$, can form to block all paths from s_j to t_j for $1 \le j \le 7$. Since this sums up to $\mu(C) = 6k + q < r$ (i.e., $v(C) > 0$), inconsistent bribery must fail.

Now assume that the briber plays consistently. In this case, the coalition must include all players on vertices d_i, $1 \le i \le 2k$; other-

wise, there would be a path from s_1 to t_1. Note that, by construction, for all i, $1 \le i \le k$, at least two players on vertices in $\{a_i, b_i, c_i\}$ must be part of the blocking coalition, since otherwise there is either a path from s_4 to t_4 or from s_5 to t_5. Likewise, at least two players on vertices in $\{a_i', b_i', c_i'\}$ must participate in a blocking coalition. Also notice that if the player on a_i is not in the blocking coalition, the ones on b_i' and c_i' must be (and again symmetric statements can be made for a_i', b_i, etc.). Furthermore, for each $y \in Y$, either the player on y or $\neg y$ must be part of the blocking coalition; otherwise, there is a path from s_2 to t_2. Altogether, this forces the coalition to include $6k + q$ vertices with cost 1 each, leaving for each i two players on one of $\{a_i, a_i'\}$, $\{b_i, b_i'\}$, or $\{c_i, c_i'\}$ out of the blocking coalition, and for each $y \in Y$, a player on one of y or $\neg y$ out as well. Therefore, the blocking coalition represents a consistent assignment to the variables in Y and, for each implicant, cancels the effect of two out of three of its literals. In this case, the briber can only succeed if for some implicant all three literals are satisfied, no matter what assignment the coalition chooses. This, of course, happens if and only if the original quantified boolean formula F is valid. □

3. CONCLUSIONS

Confirming a conjecture of Rey and Rothe [3], we have shown that bribery in multiple-adversary path-disruption games with costs is Σ_2^p-hard. On the one hand, this result completes the picture of the complexity of bribery problems in PDGs; on the other hand, it provides a Σ_2^p-completeness result for a natural problem in game theory, which are far rarer than NP-completeness results in this area. One recent other such result is due to Woeginger [7] who showed that recognizing core stability in additive hedonic games is Σ_2^p-complete as well.

Our computational hardness result may be interpreted as providing protection against bribery attacks on multi-adversary PDGs with costs, which is a positive result in light of the motivation of these games in terms of network security issues [1]. One may wonder to what extent this result provides more insight into the computational complexity of PDGC-MULTIPLE-BRIBERY, given that its NP-hardness follows from the single-adversary case with costs [3]. As pointed out by Woeginger [6], however, Σ_2^p-hardness indeed provides a much better protection than merely NP-hardness, since most of the common methods used to circumvent NP-hardness – such as approximation, fixed-parameter tractability, or typical-case analyses – are far less applicable to circumvent Σ_2^p-hardness.

4. REFERENCES

[1] Y. Bachrach and E. Porat. Path disruption games. In *Proc. AAMAS'10*, pages 1123–1130. IFAAMAS, 2010.

[2] A. Meyer and L. Stockmeyer. The equivalence problem for regular expressions with squaring requires exponential space. In *Proc. FOCS'72*, pages 125–129, 1972.

[3] A. Rey and J. Rothe. Bribery in path-disruption games. In *Proc. ADT'11*, pages 247–261. Springer-Verlag *LNAI #6992*, 2011.

[4] A. Rey and J. Rothe. Probabilistic path-disruption games. In *Proc. ECAI'12*, pages 923–924. IOS Press, 2012. Extended version in *Proc. STAIRS'12*, pages 264–269, IOS Press, 2012.

[5] L. Stockmeyer. The polynomial-time hierarchy. *Theoretical Computer Science*, 3(1):1–22, 1976.

[6] G. Woeginger. Coalitions in hedonic games, 2012. Invited talk at *COMSOC'12*, Kraków, Poland.

[7] G. Woeginger. A hardness result for core stability in additive hedonic games. *Mathematical Social Sciences*, 65(2):101–104, 2013.