

Complexity and Approximations in Robust Coalition Formation via Max-Min k -Partitioning

Extended Abstract

Anisse Ismaili
RIKEN AIP Center for
Advanced Intelligence
Project
Tokyo, Japan
anisse.ismaili@riken.jp

Noam Hazon
Ariel University
Ariel, Israel
noamh@ariel.ac.il

Emi Watanabe
Makoto Yokoo
Kyushu University
Fukoka, Japan
watanabe@agent.yokoo@
inf.kyushu-u.ac.jp

Sarit Kraus
Bar-Ilan University
Ramat-Gan, Israel
sarit@cs.biu.ac.il

ABSTRACT

Coalition formation is beneficial to multi-agent systems, especially when the value of a coalition depends on the relationship among its members. However, an attack can significantly damage a coalition structure by disabling agents. Therefore, getting prepared in advance for such an attack is particularly important. We study a robust k -coalition formation problem modeled by max-min k -partition of a weighted graph. We show that this problem is Σ_2^P -complete, which holds even for $k = 2$ and arbitrary weights, or $k = 3$ and non-negative weights. We also propose the Iterated Best Response (IBR) algorithm which provides a run-time absolute bound for the approximation error and can be generalized to the max-min optimization version of any Σ_2^P -complete problem. We tested IBR on fairly large instances of both synthetic graphs and real life networks, yielding near optimal results in a reasonable time.

KEYWORDS

Coalition Formation; k -Partition; Robustness; Complexity

ACM Reference Format:

Anisse Ismaili, Noam Hazon, Emi Watanabe, Makoto Yokoo, and Sarit Kraus. 2019. Complexity and Approximations in Robust Coalition Formation via Max-Min k -Partitioning. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019*, IFAAMAS, 3 pages.

1 INTRODUCTION

In many situations agents regroup into coalitions toward working together, in order to complete a set of k tasks [7, 24]. While any coalition structure with less or more than k coalitions will not complete all the tasks or includes redundant activity, another important aspect is the robustness of the solution. Indeed, an attack may decrease the system's performance by removing agents. Therefore, the coalition structure should be carefully selected in order to be prepared for the worst anticipated loss from such malicious action. In this paper we take the view of a defender that would like to find a coalition structure with k coalitions that maximizes the social welfare, given that an attacker will try to minimize it by removing up

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada. © 2019 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved. ...\$ACM ISBN 978-x-xxxx-xxxx-x/YY/MM
...\$15.00

to m agents. We use a concise representation based on a weighted graph: every agent is a node, and the additional value agent i has for agent j being with him is the weight of link $\{i, j\}$. Our problem is thus to find a max-min k -partition of the given graph.

Related Work. Coalition structure generation [22] aims at partitioning a set of agents into coalitions to maximize some system-wide performance measure. Traditionally, the input consists of a black-box characteristic function that returns the value for each coalition [23]. State-of-the-art algorithms can solve problem instances with 25 agents within 100 seconds. Also, several concise characteristic function representations have been proposed, for example, marginal contribution nets (MC-nets) [13], synergy coalition group [8], and coalition resource game [31]. The graphical representation [2, 6, 10, 26, 27] which we use is one of the simplest. There is vast literature on analyzing Stackelberg security games [25, 28] and additional works that study coalition formation problems with other models of failures [1, 3–5, 12, 19]. The closest work is Okimoto et al. [20], with the differences that we constrain the number of coalitions (vs. unconstrained) and use a compact representation (vs. oracle). Moreover, our algorithm for the defender has a provable approximation bound and a practical running time.

Model. A max-min k -partition instance is a tuple $\langle N, L, w, k, m, \theta \rangle$.

- (N, L, w) is a weighted undirected graph. $N = [n]$, where $n \in \mathbb{N}$ is a set of nodes.¹ The set of links $L \subseteq \binom{N}{2}$ consists of unordered node pairs. Link $\ell = \{i, j\}$ maps to weight $w_{ij} \in \mathbb{Z}$. Equivalently, $w : N^2 \rightarrow \mathbb{Z}$ satisfies for any $(i, j) \in N^2$ that $w(i, i) = 0$, $w(i, j) = w(j, i)$ and $w(i, j) \neq 0 \Rightarrow \{i, j\} \in L$.
- k is the size of a partition, $2 \leq k < n$.
- $m \in \mathbb{N}$ is the number of nodes that could be removed.
- $\theta \in \mathbb{Z}$ is a threshold value.

Let π denote a k -partition of N , which is a collection of node-subsets $\{S_1, \dots, S_k\}$, such that for each $i \in [k]$, $S_i \subseteq N$, and $\forall S_i, S_j \in \pi$, where $i \neq j$, $S_i \cap S_j = \emptyset$ holds. We say that a k -partition π is complete when $\bigcup_{i \in [k]} S_i = N$ holds (otherwise, it is incomplete). For node $i \in N$, $\pi(i)$ is the node-subset to which it belongs. For any $S \subseteq N$, we define

$$W(S) = \sum_{\{i, j\} \subseteq S} w(i, j).$$

Then, let $W(\pi)$ denote $\sum_{S \in \pi} W(S)$. We require that no node-subset be empty; hence, if some node-subset is empty, we set $W(\pi) = -\infty$.

¹Given $n \in \mathbb{N}$, $[n]$ is shorthand of $\{1, \dots, n\}$.

Given a k -partition $\pi = \{S_1, \dots, S_k\}$ and a set $M \subseteq N$, the remaining incomplete partition π_{-M} after removing M is defined as $\{S'_1, \dots, S'_k\}$, where $S'_i = S_i \setminus M$. Let $W_{-m}(\pi)$ denote the minimum value after removing at most m nodes, i.e., it is defined as:

$$W_{-m}(\pi) = \min_{M \subseteq N, |M| \leq m} \{W(\pi_{-M})\}.$$

To obtain $W_{-m}(\pi) \neq -\infty$, every $S \in \pi$ needs to contain at least $m + 1$ nodes, so that no node-subset of π_{-M} is emptied.

2 COMPLEXITY OF MAX-MIN-K-PARTITION

We studied computational complexity for the max-min defender's problem and found two intricate results that are detailed in [14]. The standard verification problem itself turns out to be coNP-complete, which intricates one more level in the polynomial hierarchy (PH). We indeed show that MAX-MIN- k -PARTITION (given a max-min k -partition instance, does a k -partition π s.t. $W_{-m}(\pi) \geq \theta$ exist?) is complete for class Σ_2^P , even in two cases:

- (a) when $k = 2$ for arbitrary link weights $w \leq 0$, or
- (b) when $k = 3$ for non-negative link weights $w \geq 0$.

Though we don't know for $k = 2$ and $w \geq 0$, these results match what is known on MAXCUT [15] (amounts to MIN-2-CUT with $w \leq 0$, and NP-complete) and MIN-3-CUT [9] (NP-complete when one node is fixed in each node-subset), but one level higher in PH.

3 ITERATED BEST RESPONSE ALGORITHM

Usually, in the second level of PH, instances become very quickly intractable (e.g., above $n \geq 20$ in [20]). Thus, we introduce an algorithm that we call Iterated Best Response algorithm (IBR) for solving a max-min k -partition problem. The idea is to start from a random (new) k -partition π , and then iterate the following loop:

- (1) Attacker response: $M \leftarrow \arg \min_{M' \subseteq N, |M'| \leq m} W(\pi_{-M'})$.
- (2) Defender response: find $\hat{\pi}$, optimal k -partition of $N \setminus M$ s.t. $\text{df}(\hat{\pi}) \leq |M|$ (where $\text{df}(\hat{\pi}) = \sum_{C \in \hat{\pi}} \min\{|C| - m - 1, 0\}$), and complete $\hat{\pi}$ into a k -partition π of N s.t. $\forall C \in \pi, |C| \geq m + 1$.

An outer loop may run this best-response dynamics several times.

This algorithm provides a run-time absolute bound for the approximation error. Let lb be the maximum value $W_{-m}(\pi)$ found so far for any k -partition π , that is the value of the currently best known solution. Let ub be the minimum value $W(\hat{\pi})$ found so far. Then the solution returned by the algorithm is within an additive $ub - lb$ of the optimum. Denoting $\text{OPT} = \max_{\pi} \{W_{-m}(\pi)\}$, it means:

$$\text{OPT} - W_{-m}(\pi) \leq ub - lb.$$

Consider the example in Figure 1. Assume $k = 2$ and $m = 1$. Due to negative links, $\{(1, 2, 5), (3, 4, 6)\}$ and $\{(1, 4, 5), (2, 3, 6)\}$ are the only meaningful 2-partitions. Also, removing 1 or 3 is always better than removing other nodes. For these meaningful actions of defender/attacker, a payoff matrix is given as the table in Figure 1 ($M = \emptyset$ means no attack). Assume $\pi = \{(1, 2, 5), (3, 4, 6)\}$ is chosen at first. Then, the best response of the attacker is $M = \{3\}$. Then, lb is updated to 13. Then, the defender chooses $\hat{\pi}_{-M} = \{(1, 4, 5), (2, 6)\}$, which is an optimal partition of $\{1, 2, 4, 5, 6\}$. Then, $W(\hat{\pi}_{-M}) = 20$ is used to update ub , i.e., as long as the attacker chooses $\{3\}$, the value of the defender is at most 20. Next, the defender chooses $\{(1, 4, 5), (2, 3, 6)\}$, which subsumes $W(\hat{\pi}_{-M})$. The best response of the attacker is $M = \{1\}$ and lb is updated to 18. The defender

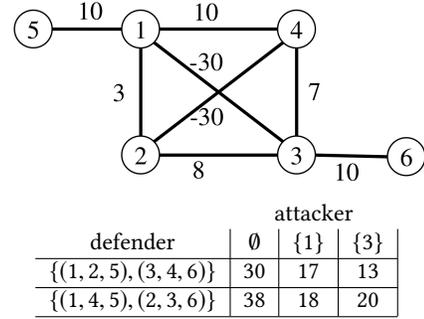


Figure 1: Example ($n = 6$, $k = 2$, and $m = 1$) and payoff matrix

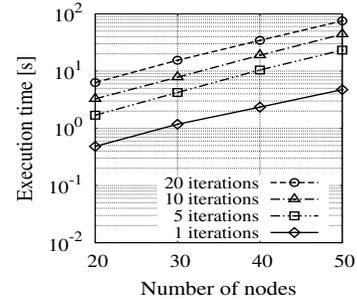


Figure 2: Evaluation results (real life network)

chooses $\hat{\pi}_{-M} = \{(4, 5), (2, 3, 6)\}$, which is an optimal partition of $\{2, 3, 4, 5, 6\}$. Then, $W(\hat{\pi}_{-M}) = 18$ is used to update ub . Now, $lb = ub$ holds and IBR terminates.

4 EXPERIMENTAL EVALUATION

We experimentally evaluate the performance of IBR. All the tests were run on a machine: an Intel Xeon E5-2680v4 CPU @ 2.40GHz processor with 125.8GB RAM, Ubuntu 16.40 LTS, and a mixed integer programming package Gurobi version 7.5.0. We show experiments based on a real life network called Wikipedia Requests for Adminship (RfA) network [30]. This is a network among Wikipedia users where each link (i, j) has a weight corresponding to the vote of user i towards user j to become an administrator. The weight of a link is given based on the intensity of the sentiment expressed in the vote [17]. The original graph is directed. For a pair of nodes i and j , we create an undirected link with weight $w(i, j) + w(j, i)$. The original graph has about 10,000 nodes and 100,000 links. Based on this original graph, we select a subgraph with n nodes by randomly choosing a root node, then by adding neighboring nodes in a breadth-first manner. In an obtained graph, the probability that a link exists is about 20% (about 90% of them have positive weights). We set the number of removed nodes m to $n/10$.

Figure 2 shows the computation time of IBR by varying n with $k = 3$. We show the results by varying the number of iterations for *outer-loop*. Each data point is an average of 100 problem instances. We can see that IBR can solve problem instances with $n = 50$ within 5 seconds when the number of iterations is 1.

Acknowledgements. This work was partially supported by the Ministry of Science and Technology, Israel, the Japan Science and Technology Agency (JST) Strategic International Collaborative Research Program, and JSPS KAKENHI Grant Number JP17H00761.

REFERENCES

- [1] Yoram Bachrach, Ian Kash, and Nisarg Shah. 2012. Agent failures in totally balanced games and convex games. In *Proceedings of the 8th International Workshop on Internet and Network Economics (WINE-2012)*, 15–29.
- [2] Yoram Bachrach, Pushmeet Kohli, Vladimir Kolmogorov, and Morteza Zadimoghaddam. 2013. Optimal Coalition Structure Generation in Cooperative Graph Games. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence (AAAI-2013)*, 81–87.
- [3] Yoram Bachrach, Reshef Meir, Michal Feldman, and Moshe Tennenholtz. 2011. Solving cooperative reliability games. In *Proceedings of the 27th Conference on Uncertainty in Artificial Intelligence (UAI-2011)*, 27–34.
- [4] Yoram Bachrach, Rahul Savani, and Nisarg Shah. 2014. Cooperative max games and agent failures. In *Proceedings of the 13th International Conference on Autonomous Agents and Multi-agent Systems (AAMAS-2014)*, 29–36.
- [5] Yoram Bachrach and Nisarg Shah. 2013. Reliability weighted voting games. In *Proceedings of the 6th International Symposium on Algorithmic Game Theory (SAGT-2013)*, 38–49.
- [6] Simina Brânzei and Kate Larson. 2009. Coalitional Affinity Games. In *Proceedings of the 8th International Joint Conference on Autonomous Agents and Multi-agent Systems (AAMAS-2009)*, 1319–1320.
- [7] Georgios Chalkiadakis, Edith Elkind, and Michael Wooldridge. 2011. *Computational Aspects of Cooperative Game Theory*. Morgan-Claypool.
- [8] Vincent Conitzer and Tuomas Sandholm. 2006. Complexity of constructing solutions in the core based on synergies among coalitions. *Artificial Intelligence* 170, 6-7 (2006), 607–619.
- [9] Elias Dahlhaus, David Johnson, Christos H. Papadimitriou, Paul D. Seymour, and Mihalis Yannakakis. 1992. The Complexity of Multiway Cuts (Extended Abstract). In *Proceedings of the 24th Annual ACM Symposium on Theory of Computing (STOC-1992)*, 241–251.
- [10] Xiaotie Deng and Christos H Papadimitriou. 1994. On the complexity of cooperative solution concepts. *Mathematics of Operations Research* 19, 2 (1994), 257–266.
- [11] Drew Fudenberg and Jean Tirole. 1991. *Game Theory*. The MIT Press, Cambridge, Chapter 3: Extensive-Form Games.
- [12] Qingyu Guo, Bo An, Yevgeniy Vorobeychik, Long Tran-Thanh, Jiarui Gan, and Chunyan Miao. 2016. Coalitional security games. In *Proceedings of the 15th International Conference on Autonomous Agents and Multi-agent Systems (AAMAS-2016)*, 159–167.
- [13] Samuel Ieong and Yoav Shoham. 2005. Marginal contribution nets: a compact representation scheme for coalitional games. In *Proceedings of the 6th ACM conference on Electronic commerce (EC-2005)*, 193–202.
- [14] Anisse Ismaili. 2019. The Complexity of Max-Min k -Partitioning. In *arXiv preprint arXiv:1902.06812*.
- [15] Richard M. Karp. 1972. Reducibility among combinatorial problems. In *Complexity of computer computations*. Springer, 85–103.
- [16] Ker-I Ko and Chih-Long Lin. 1995. *On the Complexity of Min-Max Optimization Problems and their Approximation*. Springer US, Boston, MA, 219–239.
- [17] Srijan Kumar, Francesca Spezzano, V. S. Subrahmanian, and Christos Faloutsos. 2016. Edge Weight Prediction in Weighted Signed Networks. In *Proceedings of IEEE 16th International Conference on Data Mining (ICDM-2016)*, 221–230.
- [18] John Neumann and Oskar Morgenstern. 2004. *Theory of games and economic behavior*. Princeton University Press, Princeton.
- [19] Tenda Okimoto, Nicolas Schwind, Maxime Clement, Tony Ribeiro, Katsumi Inoue, and Pierre Marquis. 2015. How to form a task-oriented robust team. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2015)*, 395–403.
- [20] Tenda Okimoto, Nicolas Schwind, Emir Demirović, Katsumi Inoue, and Pierre Marquis. 2018. Robust Coalition Structure Generation. In *Proceedings of the 21st International Conference on Principles and Practice of Multi-Agent Systems (PRIMA-2018)*, 140–157.
- [21] Christos H Papadimitriou and Mihalis Yannakakis. 1991. Shortest paths without a map. *Theoretical Computer Science* 84, 1 (1991), 127–150.
- [22] Talal Rahwan, Tomasz P Michalak, Michael Wooldridge, and Nicholas R Jennings. 2015. Coalition structure generation: A survey. *Artificial Intelligence* 229 (2015), 139–174.
- [23] Tuomas Sandholm, Kate Larson, Martin Andersson, Onn Shehory, and Fernando Tohmé. 1999. Coalition structure generation with worst case guarantees. *Artificial Intelligence* 111, 1-2 (1999), 209–238.
- [24] Onn Shehory and Sarit Kraus. 1998. Methods for task allocation via agent coalition formation. *Artificial intelligence* 101, 1 (1998), 165–200.
- [25] Arunesh Sinha, Fei Fang, Bo An, Christopher Kiekintveld, and Milind Tambe. 2018. Stackelberg Security Games: Looking Beyond a Decade of Success. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI-2018)*, 5494–5501.
- [26] Liat Sless, Noam Hazon, Sarit Kraus, and Michael Wooldridge. 2014. Forming coalitions and facilitating relationships for completing tasks in social networks. In *Proceedings of the 13th International Conference on Autonomous Agents and Multi-agent Systems (AAMAS-2014)*, 261–268.
- [27] Liat Sless, Noam Hazon, Sarit Kraus, and Michael Wooldridge. 2018. Forming k coalitions and facilitating relationships in social networks. *Artificial Intelligence* 259 (2018), 217–245.
- [28] Milind Tambe. 2011. *Security and game theory: algorithms, deployed systems, lessons learned*. Cambridge University Press.
- [29] Suguru Ueda, Makoto Kitaki, Atsushi Iwasaki, and Makoto Yokoo. 2011. Concise characteristic function representations in coalitional games based on agent types. In *Proceedings of the 10th International Conference on Autonomous Agents and Multi-agent Systems (AAMAS-2011)*, 1271–1272.
- [30] Robert West, Hristo S Paskov, Jure Leskovec, and Christopher Potts. 2014. Exploiting Social Network Structure for Person-to-Person Sentiment Analysis. *Transactions of the Association for Computational Linguistics* 2 (2014), 297–310.
- [31] Michael Wooldridge and Paul E Dunne. 2006. On the computational complexity of coalitional resource games. *Artificial Intelligence* 170, 10 (2006), 835–871.