

Brief Announcement: Efficient Best Response Computation for Strategic Network Formation under Attack

Tobias Friedrich, Sven Ihde, Christoph Keßler,
Pascal Lenzner, Stefan Neubert and David Schumann
Algorithm Engineering Group, Hasso Plattner Institute, Potsdam, Germany

ABSTRACT

Inspired by real world examples, e.g. the Internet, researchers have introduced an abundance of strategic games to study natural phenomena in networks. Unfortunately, almost all of these games have the conceptual drawback of being computationally intractable, i.e. computing a best response strategy or checking if an equilibrium is reached is NP-hard. Thus, a main challenge in the field is to find tractable realistic network formation models.

We address this challenge by establishing that the recently introduced model by Goyal et al. [WINE'16], which focuses on robust networks in the presence of a strong adversary, is a rare exception which is both realistic and computationally tractable. In particular, we sketch an efficient algorithm for computing a best response strategy, which implies that deciding whether the game has reached a Nash equilibrium can be done efficiently as well. Our algorithm essentially solves the problem of computing a minimal connection to a network which maximizes the reachability while hedging against severe attacks on the network infrastructure.

KEYWORDS

Strategic Network Formation, Best Response Computation, Network Robustness, Network Reachability

1 INTRODUCTION

Many of today's important networks, most prominently the Internet, are essentially the outcome of an unsupervised decentralized network formation process among many selfish entities [20]. Creating links in such networks yields a connectivity benefit but also harbors the risk of collateral damage if a neighbor is attacked.

Required features of any Internet-like communication network are reachability and robustness. Such networks have to ensure that even in case of cascading edge or node failures caused by technical defects or malicious attacks, e.g. DDoS-attacks or viruses, most participating nodes can still communicate. This focus on robustness has long been neglected and is now a very recent endeavor in the strategic network formation community, see e.g. [6, 14, 16, 18].

Our contribution is the proof that the recently introduced model by Goyal et al. [13, 14] is one of the few examples of a tractable realistic model for strategic network formation, which answers an open question by these authors. In particular, we provide an

efficient algorithm for computing a best response strategy for two variants of their model which enables large-scale simulations to analyze phenomena in real-world networks.

Related Work: The model by Goyal et al. [13, 14] essentially augments the well-known reachability model by Bala & Goyal [2] with robustness considerations. In particular, different types of adversaries are introduced which attack (and destroy) a node of the network. This attack then spreads virus-like to neighboring nodes and destroys them as well. Besides deciding which links to form, players also decide whether they want to buy immunization against eventual attacks. The model is the first model which incorporates network formation and immunization decisions at the same time.

On the one hand, the authors of [13, 14] provide beautiful structural results, e.g. they show that equilibrium networks are much more diverse than in the non-robust version, that the amount of edge overbuilding due to robustness concerns is small and that all equilibrium networks achieve very high social welfare. On the other hand, the authors raise the intriguing open problem of settling the complexity of computing a best response strategy in their model¹.

Computing a best response in network formation games can be done efficiently for the reachability model [2] and if the allowed strategy changes are simple [17]. However, these examples are exceptions. For almost all related network formation models, e.g. [4–6, 8, 10, 11, 15, 19], NP-hardness has been shown.

To the best of our knowledge, besides the model by Goyal et al. [13, 14] there are only a few other models which combine selfish network formation with robustness considerations and all of them consider a much weaker adversary. The earliest are models by Bala & Goyal [3] and Kliemann [16] and both augment the model [2] with single edge failures. Other related models are by Meirom et al. [18] and Chauhan et al. [6]. Both consider players striving for centrality but who at the same time want to protect themselves against single edge failures. The complexity status for computing a best response was only settled for the model by Chauhan et al. [6] where this problem was proven to be NP-hard.

Also vaccination games, e.g. [1, 7, 21], are related. There the network is fixed and the selfish nodes only have to decide if they want to immunize or not. Computing a best response in these models is trivial but pure Nash equilibria may not exist.

2 MODEL

We work with the strategic network formation model proposed by Goyal et al. [13, 14] and mostly use their notation. In this model the n nodes of a network $G = (V, E)$ correspond to individual players $v_1, \dots, v_n$. The edge set E is determined by the players'

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SPAA '17, July 24-26, 2017, Washington DC, USA

© 2017 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-4593-4/17/07.

<https://doi.org/http://dx.doi.org/10.1145/3087556.3087594>

¹This question was raised in versions 1-4 of [13] and is replaced in later versions and in [14] with a reference to a preprint [12] of the present paper.

strategic behavior as follows. Each player $v_i \in V$ can decide to buy undirected edges to a subset of other players, paying $\alpha > 0$ per edge, where α is some fixed parameter of the model.

If player v_i decides to buy the edge to node v_j , then we say that the edge $\{v_i, v_j\}$ is owned and paid for by player v_i . Buying an undirected edge entails connectivity benefits and risks for both participating endpoints. In order to cope with these risks, each player can also decide to buy immunization against attacks at a cost of $\beta > 0$, which is also a fixed parameter. We call a player *immunized* if this player decides to buy immunization, and *vulnerable* otherwise.

The strategy $s_i = (x_i, y_i)$ of player v_i consists of the set $x_i \subseteq V \setminus \{v_i\}$ of the nodes to buy an edge to, and the immunization choice $y_i \in \{0, 1\}$, where $y_i = 1$ if and only if player v_i decides to immunize. The strategy profile $\mathbf{s} = (s_1, \dots, s_n)$ of all players then induces an undirected graph $G(\mathbf{s}) = (V, \bigcup_{v_i \in V} \bigcup_{v_j \in x_i} \{v_i, v_j\})$. The immunization choices $y_1, \dots, y_n$ in $\mathbf{s}$ partition V into the set of immunized players $\mathcal{I} \subseteq V$ and vulnerable players $\mathcal{U} = V \setminus \mathcal{I}$. The components in the induced subgraph $G[\mathcal{U}]$ are called *vulnerable regions* and the set of those regions will be denoted by $\mathcal{R}_{\mathcal{U}}$. The vulnerable region of any vulnerable player $v_i \in \mathcal{U}$ is $\mathcal{R}_{\mathcal{U}}(v_i)$. *Immunized regions* $\mathcal{R}_{\mathcal{I}}$ are defined analogously as the components of the induced subgraph $G[\mathcal{I}]$.

After the network $G(\mathbf{s})$ is built, we assume that an adversary attacks one vulnerable player according to a strategy known to the players. We consider mostly the maximum carnage adversary [13, 14] which tries to destroy as many nodes of the network as possible. To achieve this, the adversary chooses a vulnerable region of maximum size and attacks some player in that region. If there is more than one such region with maximum size, then one of them is chosen uniformly at random. If a vulnerable player v_i is attacked, then v_i will be destroyed and the attack spreads to all vulnerable neighbors of v_i , eventually destroying all players in v_i 's vulnerable region $\mathcal{R}_{\mathcal{U}}(v_i)$. Let $t_{\max} = \max_{R \in \mathcal{R}_{\mathcal{U}}} \{|R|\}$ be the number of nodes in the vulnerable region of maximum size and $\mathcal{T} = \{v_i \in \mathcal{U} \mid |\mathcal{R}_{\mathcal{U}}(v_i)| = t_{\max}\}$ is the corresponding set of nodes which may be targeted by the adversary. The set of targeted regions is $\mathcal{R}_{\mathcal{T}} = \{R \in \mathcal{R}_{\mathcal{U}} \mid |R| = t_{\max}\}$, and $\mathcal{R}_{\mathcal{T}}(v_i)$ is the targeted region of a player $v_i \in \mathcal{T}$. Thus, if $v_i \in \mathcal{T}$ is attacked, then all players in the region $\mathcal{R}_{\mathcal{T}}(v_i)$ will be destroyed.

The *utility* of a player v_i in network $G(\mathbf{s})$ is defined as the expected number of nodes reachable by v_i after the adversarial attack on network $G(\mathbf{s})$ (zero in case v_i was destroyed) less v_i 's expenditures for buying edges and immunization. More formally, let $CC_i(t)$ be the connected component of v_i after an attack to node $v_t \in \mathcal{T}$ and let $|CC_i(t)|$ denote its number of nodes. Then the utility (or profit) $u_i(\mathbf{s})$ of v_i in the strategy profile $\mathbf{s}$ is

$$u_i(\mathbf{s}) = \frac{1}{|\mathcal{T}|} \left(\sum_{v_t \in \mathcal{T}} |CC_i(t)| \right) - |x_i| \cdot \alpha - y_i \cdot \beta.$$

Fixing the strategies of all other players, the *best response* of a player v_i is a strategy $s_i^* = (x_i^*, y_i^*)$ which maximizes v_i 's utility $u_i((s_1, \dots, s_{i-1}, s_i^*, s_{i+1}, \dots, s_n))$. We will call the strategy change to s_i^* a best response for player v_i in the network $G(\mathbf{s})$, if changing from strategy $s_i \in \mathbf{s}$ to strategy s_i^* is the best possible strategy for player v_i if no other player changes her strategy.

A best response is calculated for one arbitrary but fixed player v_a , which we call the *active player*. Furthermore let $\mathcal{C}$ be the set of connected components which exist in $G(\mathbf{s}) \setminus v_a$. Let $\mathcal{C}_{\mathcal{U}} = \{C \in \mathcal{C} \mid C \cap \mathcal{I} = \emptyset\}$, $\mathcal{C}_{\mathcal{I}} = \mathcal{C} \setminus \mathcal{C}_{\mathcal{U}}$ and $\mathcal{C}_{inc} = \{C \in \mathcal{C} \mid \exists u \in C : \{u, v\} \in E\}$, where $\mathcal{C}_{\mathcal{U}}$ is the set of components in which all vertices are vulnerable, $\mathcal{C}_{\mathcal{I}}$ is the set of components which contain at least one immunized vertex and $\mathcal{C}_{inc}$ is the set of components to which player v_a is connected via edges bought by some other player.

3 THE BEST RESPONSE ALGORITHM

A naive approach to calculate the best response for player v_a would consider all 2^n possible strategies and select one that yields the best utility. This is clearly infeasible for a larger number of players. Our algorithm exploits three observations to reduce the complexity from exponential to polynomial:

Observation 1: The network $G(\mathbf{s}) \setminus v_a$ may consist of several connected components that can be dealt with independently for most decisions. As long as the set of possible targets of the adversary does not change, the best response of v_a can be constructed by first choosing components to which a connection is profitable and then choosing for each of those components an optimal set of nodes within the respective component to build edges to.

Observation 2: Homogeneous components in $G(\mathbf{s}) \setminus v_a$, which consist of only vulnerable or only immunized nodes, provide the same benefit no matter whether v_a connects to them with one or with more than one edge. Thus the connection decision is a binary decision for those components.

Observation 3: Mixed components in $G(\mathbf{s}) \setminus v_a$, which contain both immunized and vulnerable nodes, consist of homogeneous regions that again have the property that at most one edge per homogeneous region can be profitable. Merging those regions into block nodes forms an auxiliary tree, called Meta Tree, which we use in an efficient dynamic programming algorithm to compute the most profitable subset of regions to connect with.

The Algorithm: Due to space constraints we introduce our algorithm, called BESTRESPONSECOMPUTATION, informally. A schematic overview can be found in Fig. 1. We refer to [12] for details.




Figure 1: Schematic overview of the best response algorithm.

Our algorithm solves the problem of finding a best response strategy by considering both options of buying or not buying immunization and computing for both cases the best possible set of edges to buy. Thus, the first step of `BESTRESPONSECOMPUTATION` is to drop the current strategy of the active player v_a and to replace it with the empty strategy $s_0 = (\emptyset, 0)$ in which player v_a does not buy any edge and does not buy immunization. Then the resulting strategy profile s' and the set of connected components C_U and C_I with respect to network $G(s') \setminus v_a$ is considered.

The subroutine `SUBSETSELECT` determines the optimal sets of components of C_U to connect to if v_a does not immunize by solving an adjusted Knapsack problem which involves only small numbers. Two such sets of components are computed, depending on player v_a becoming targeted or not by connecting to these components. Additionally, `GREEDYSELECT` computes a best possible subset of components of C_U to connect to in case v_a buys immunization. It does so by greedily connecting to all profitable components in C_U .

The challenging part of the problem is to cope with the connected components in C_I which also contain immunized nodes. For such components our algorithm detects and merges equivalent nodes and thereby simplifies these components. An auxiliary tree structure, which we call the Meta Tree, is constructed (see Fig. 2). This tree is




Figure 2: A mixed component; after merging the indicated nodes (middle); another merge yields the Meta Tree (right).

then used in a dynamic programming fashion to efficiently compute the best possible set of nodes to buy edges towards within the respective component. Thus, we handle components in C_I by first performing a data-reduction similar to kernelization approaches in the realm of Parameterized Algorithmics [9] and then solving the reduced problem via dynamic programming.

The subroutine `POSSIBLESTRATEGY` obtains the best set of nodes in components in C_I . As this set depends on the number of targeted regions, it has to be determined for several cases independently. These cases are v_a not being immunized and not being targeted, v_a not being immunized but being targeted, and v_a being immunized. For each case, the subroutine `POSSIBLESTRATEGY` first chooses an arbitrary single edge to buy into the previously selected components from C_U . Then the best set of edges into components in C_I to buy is computed independently for each component $C \in C_I$ via the subroutines `PARTNERSETSELECT`, `METATREECONSTRUCT` and `METATREESELECT`. The union of the obtained sets is then returned. Finally, the algorithm compares the empty strategy and the individually obtained best possible strategies for the above mentioned cases and selects the one which maximizes player v_a 's utility.

The run time of our best response algorithm heavily depends on the size of the largest obtained Meta Tree and we achieve a worst-case run time of $O(n^4 + k^5)$ for the maximum carnage adversary and $O(n^4 + nk^5)$ for the random attack adversary, where n is the number of nodes in the network and k is the number of blocks in the largest

Meta Tree. This yields a run time in $O(n^5)$ and $O(n^6)$, respectively. To contrast this upper bound, we also provide empirical results showing that k is usually much smaller than n , which emphasizes the effectiveness of our data-reduction.

4 CONCLUSION

For most models of strategic network formation computing a utility maximizing strategy is known to be NP-hard. In this paper, we have proven that the model by Goyal et al. [13, 14] is a notable exception to this rule. The presented efficient algorithm for computing a best response for a player circumvents a combinatorial explosion essentially by simplifying the given network and thereby making it amenable to a dynamic programming approach. An efficient best response computation is the key ingredient for using the model in large scale simulations and for analyzing real world networks. Moreover, our algorithm can be adapted to a significantly stronger adversary and we are confident that further modifications for coping with other variants of the model are possible. One such variant is the directed version which would more accurately model the differences in risk and benefit which depend on the flow direction.

REFERENCES

- [1] James Aspnes, Kevin Chang, and Aleksandr Yampolskiy. 2006. Inoculation strategies for victims of viruses and the sum-of-squares partition problem. *J. Comput. System Sci.* 72, 6 (2006), 1077–1093.
- [2] Venkatesh Bala and Sanjeev Goyal. 2000. A Noncooperative Model of Network Formation. *Econometrica* 68, 5 (2000), 1181–1229.
- [3] Venkatesh Bala and Sanjeev Goyal. 2003. A strategic analysis of network reliability. In *Networks and Groups*. Springer, 313–336.
- [4] Davide Bilò, Luciano Gualà, Stefano Leucci, and Guido Proietti. 2014. Locality-based network creation games. In *SPAA'14*. 277–286.
- [5] Davide Bilò, Luciano Gualà, and Guido Proietti. 2015. Bounded-Distance Network Creation Games. *ACM TEAC* 3, 3, Article 16 (2015), 20 pages.
- [6] Ankit Chauhan, Pascal Lenzner, Anna Melnichenko, and Martin Münn. 2016. On Selfish Creation of Robust Networks. In *SAGT'16*. 141–152.
- [7] Po-An Chen, Mary David, and David Kempe. 2010. Better vaccination strategies for better people. In *EC'10*. 179–188.
- [8] Andreas Cord-Landwehr and Pascal Lenzner. 2015. Network Creation Games: Think Global - Act Local. In *MFCS'15*. 248–260.
- [9] Rodney G. Downey and Michael R. Fellows. 2013. *Fundamentals of Parameterized Complexity*. Springer.
- [10] Shayan Ehsani, Saber Shokat Fadaee, Mohammad Amin Fazli, Abbas Mehrabian, Sina Sadeghian Sadeghbad, Mohammad Ali Safari, and Morteza Saghafian. 2015. A Bounded Budget Network Creation Game. *ACM TALG* 11, 4 (2015), 34.
- [11] Alex Fabrikant, Ankur Luthra, Elitza N. Maneva, Christos H. Papadimitriou, and Scott Shenker. On a network creation game. In *PODC'03*. 347–351.
- [12] Tobias Friedrich, Sven Ihde, Christoph Keßler, Pascal Lenzner, Stefan Neubert, and David Schumann. 2016. Efficient Best-Response Computation for Strategic Network Formation under Attack. *CoRR* abs/1610.01861 (2016). <http://arxiv.org/abs/1610.01861>
- [13] Sanjeev Goyal, Shahin Jabbari, Michael Kearns, Sanjeev Khanna, and Jamie Morgenstern. 2015. Strategic Network Formation with Attack and Immunization. *CoRR* abs/1511.05196 (2015). <http://arxiv.org/abs/1511.05196>
- [14] Sanjeev Goyal, Shahin Jabbari, Michael Kearns, Sanjeev Khanna, and Jamie Morgenstern. 2016. Strategic Network Formation with Attack and Immunization. In *WINE'16*. 429–443.
- [15] Bernd Kawald and Pascal Lenzner. 2013. On dynamics in selfish network creation. In *SPAA'13*. 83–92.
- [16] Lasse Kliemann. 2011. The Price of Anarchy for Network Formation in an Adversary Model. *Games* 2, 3 (2011), 302–332.
- [17] Pascal Lenzner. Greedy Selfish Network Creation. In *WINE'12*. 142–155.
- [18] Eli A. Meir, Shie Mannor, and Ariel Orda. 2015. Formation games of reliable networks. In *INFOCOM'15*. 1760–1768.
- [19] Mátúš Mihalák and Jan Christoph Schlegel. 2010. The Price of Anarchy in Network Creation Games Is (Mostly) Constant. In *SAGT'10*. 276–287.
- [20] Christos H. Papadimitriou. 2001. Algorithms, games, and the internet. In *STOC'01*. 749–753.
- [21] Sudip Saha, Abhijin Adiga, and Anil Kumar S Vullikanti. Equilibria in Epidemic Containment Games. In *AAAI'14*. 777–783.