Individual Security and Network Design with Malicious Nodes

Extended Abstract

Tomasz Janus
Institute of Informatics, University of Warsaw
Warsaw, Poland
t.janus@mimuw.edu.pl

Mateusz Skomra
CMAP, Ecole Polytechnique, CNRS & INRIA
Palaiseau, France
mateusz.skomra@inria.fr

Marcin Dziubiński
Institute of Informatics, University of Warsaw
Warsaw, Poland
m.dziubinski@mimuw.edu.pl

ABSTRACT

Networks are beneficial to those being connected but can also be used as carriers of contagious hostile attacks. These attacks are often facilitated by exploiting corrupt network users. To protect against the attacks, users can resort to costly defense. The decentralized nature of such protection is known to be inefficient but the inefficiencies can be mitigated by a careful network design. Is network design still effective when not all users can be trusted? We propose a model of network design and defense with byzantine nodes to address this question. We study the optimal defended networks in the case of centralized defense and, for the case of decentralized defense, we show that the inefficiencies due to decentralization can be fully mitigated, despite the presence of the byzantine nodes.

KEYWORDS

Network design; Individual security; Byzantine players; Inefficiencies; Networks

ACM Reference Format:
Tomasz Janus, Mateusz Skomra, and Marcin Dziubiński. 2018. Individual Security and Network Design with Malicious Nodes. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018,* IFAAMAS, 3 pages.

1 INTRODUCTION

Game theoretic models of interdependent security have been used to study security of complex information and physical systems for more than a decade [11]. One of the key findings is that the externalities resulting from security decisions made by selfish agents lead to, potentially significant, inefficiencies. This motivates research on methods for improving information security, such as insurance [4] and network design [5, 6]. We study the effectiveness of network design for improving system security with malicious (or byzantine) players and strategic adversary.

Related work. There are two, overlapping, strands of literature that our work is related to: the interdependent security games [11] and multidefender security games [14, 15, 17]. Early research on interdependent security games assumed that the players only care about their own survival and that there are no benefits from being connected [1, 3, 7, 10, 12, 13, 18]. The authors of [3] study a setting in which the network is fixed, nodes care about their own survival only and both protection and contagion are perfect. They point out the high inefficiency of decentralized protection. For a comprehensive review of interdependent security games see an excellent survey [11]. Papers most related to our work are [5, 6, 8, 16]. The authors of [16] introduce malicious nodes to the model of [3] and show that their presence reduces the problem of underprotection. Works [5, 6] show that network design can mitigate inefficiencies of decentralized protection. Our work builds on [5, 6] by introducing malicious nodes to the model. The paper [8] shows that the inefficiencies caused by the decentralization of defense are relatively low under decentralized network formation.

2 THE MODEL

There are $(n+2)$ players: the designer (D), the nodes (V), and the adversary (A). Each of the nodes is either *genuine* or *byzantine*. There are at least $n = 3$ nodes and $n_B \geq 1$ of them are byzantine. The byzantine nodes cooperate with A, who knows their identity. All the nodes know their own type only. A has complete information about the game. He infects a subset of $n_{\mathrm{A}} \geq 1$ nodes. A *network* is modeled by an undirected graph $G = (V, E)$. The set of all networks over a set of nodes U is denoted by $\mathcal{G}(U)$. The game proceeds in four rounds (n, n_B, and n_{A} are fixed before the game):

(1) The types of the nodes are realized.
(2) D chooses $G \in \mathcal{G}(V)$.
(3) Nodes observe G and choose, simultaneously and independently, whether to protect or not. This determines the set of protected nodes Δ. The protection of the byzantine nodes is fake and, when attacked, such node gets infected and transmits the infection to all her neighbors.
(4) A observes the protected network (G, Δ) and chooses a set I of n_{A} nodes to infect. The infection spreads and eliminates all unprotected nodes reachable from I in G via a path that does not contain a genuine protected node. This leads to the residual network obtained from G by removing all the infected nodes.

Payoffs to the players are based on the residual network and costs of defense. The returns from a network are measured by a *network value function* Φ: $\bigcup_{U \subseteq V} \mathcal{G}(U) \to \mathbb{R}$. We consider the following family of network value functions: $\Phi(G) = \sum_{C \in C(G)} f(|C|)$, where $C(G)$ is the set of connected components of G. Moreover, the function $f : \mathbb{R}_{\geq 0} \to \mathbb{R}$ is increasing, strictly convex, $f(0) = 0$, and, for all $x \geq 1$, satisfies $f(3x) \geq 2f(2x)$ and $f(3x+2) \geq f(2x+2)+f(2x+1)$. Such form of network value function is in line with Metcalfe's law, where $f(x) = x^2$.

A and the byzantine nodes aim to minimize the value of the residual network. D aims to maximize the value of the residual

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden.

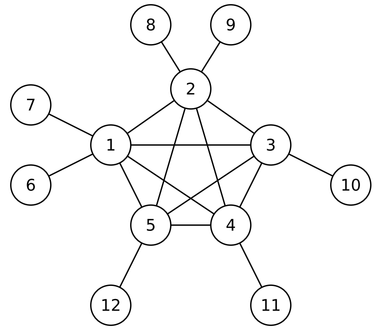

Figure 1: A generalized star with 12 nodes and core of size 5.

network minus the cost of defense. Genuine nodes aim to maximize an equal share of the value of their component minus the cost of protection $c \in \mathbb{R}_{>0}$. A and the byzantine nodes make choices that maximize their utility. D and the genuine nodes make choices that maximize the worst possible type realization (cf. [2]). The *pessimistic utility* of D from network G, the set of protected nodes Δ, and the set of infected nodes I, is denoted by $\hat{U}^{\mathrm{D}}(G, \Delta, I)$.

3 MAIN RESULTS

We divide the analysis into two parts. First, we consider the centralized defense model. Then, we use the results of this model to analyze the decentralized model and bound its price of anarchy.

3.1 Centralized defense

Fix the parameters n_B, n_{A} and suppose that the designer chooses both the network and the protection assignment. This leads to a two stage game where, in the first round, the designer chooses a protected network (G, Δ) and in the second round the adversary observes the protected network and nodes' types (recognizing the byzantine nodes) and chooses the nodes to attack. We are interested in subgame perfect equilibria of the game with pessimistic preferences of the designer. We call them centralized equilibria, for short. We start with the definition of a generalized star. We use $G[V']$ to denote the subnetwork of G induced by a set of nodes V'.

Definition 3.1 (Generalized k-star). Given a set of nodes V and $k \geq 1$, a *generalized k-star* over V is a network $G = (V, E)$ such that the set of nodes V can be partitioned into two sets, C (the core) of size $|C| = k$ and P (the periphery), in such a way that $G[C]$ is a clique, every node in P is connected to exactly one node in C, and every node in C is connected to $\lfloor n/k \rfloor - 1$ or $\lceil n/k \rceil - 1$ nodes in P.

An example of a generalized star is depicted in fig. 1. We are now ready to state the result characterizing equilibrium defended network and pessimistic equilibrium payoffs to the designer.

Proposition 3.2. *Let $n_B = n_{\mathrm{A}} = 1$, $n \geq 3$, $c > 0$. Let $k \geq 0$ be a number of nodes that is protected in some centralized equilibrium network. Then, there exists an equilibrium network (G, Δ) that has $|\Delta| = k$ protected nodes and the following structure:*

i) G has at most three connected components.

ii) If $k \geq 3$ and $n \bmod k \neq 1$, then G is a generalized k-star with protected core and unprotected periphery.

iii) If $k \geq 3$ and $n \bmod k = 1$, then G is composed of a generalized k-star of size $(n-1)$ with protected core and unprotected periphery and a single unprotected node.

iv) If $k = 0$ and $n \bmod 6 \neq 3$, then G has two connected components of size $\lfloor n/2 \rfloor$ and, if $n \bmod 2 = 1$, a single unprotected node.

v) If $k = 0$ and $n \bmod 6 = 3$, then G either has the structure given in item iv or G is composed of three components of size $n/3$.

vi) If $k = 2$, then G is composed of a generalized 2-star with protected core and unprotected periphery and at most two unprotected components.

The intuitions behind this result are as follows. When the cost of defense is high, then the designer is better off by not using any defense and partitioning the network into several components. Thanks to our assumptions on the component value function f, the number of such components is at most three.

When the cost of defense is sufficiently low, then it is profitable for the designer to protect some nodes. If the number of protected nodes is not smaller than 3, then, by choosing a generalized k-star with fully protected core (of optimal size $k \geq 3$ depending on the cost) and unprotected periphery, the designer knows that the strategic adversary is going to attack either the byzantine node (if she is among the core nodes) or any unprotected node (otherwise). Thus, in the worst case, a core node with the largest number of periphery nodes connected to her is byzantine. By distributing the core nodes evenly, the designer minimizes the impact of this worst case scenario.

3.2 Decentralized defense

Now we turn attention to the variant of the model where defense decisions are decentralized. Fix the parameters n_B, n_{A} and let $\mathcal{E}(n, c)$ denote the set of all equilibria of the game with n nodes and the cost of protection $c > 0$. Let $\hat{U}^{\mathrm{D}}_{\star}(n, c)$ denote the best payoff the designer can obtain in the centralized defense game (as discussed in section 3.1). The *price of anarchy* is the fraction of this payoff over the minimal payoff to the designer that can be attained in equilibrium of Γ (for the given cost of protection c), $\mathrm{PoA}(n, c) = \hat{U}^{\mathrm{D}}_{\star}(n, c) / \min_{\boldsymbol{e} \in \mathcal{E}(n,c)} \mathrm{E}\hat{U}^{\mathrm{D}}(\boldsymbol{e})$. Our main result provides asymptotic characterization of PoA (with a fixed cost c).

Theorem 3.3. *Suppose that for all $t \geq 0$ the function f satisfies $\lim_{n \to +\infty} f(n)/f(n-t) = 1$. Then, for any cost level $c > 0$ and any fixed parameters $n_B \geq 1$, $n_{\mathrm{A}} \geq 1$ we have $\lim_{n \to +\infty} \mathrm{PoA}(n, c) = 1$.*

Notice that the condition of theorem 3.3 is verified for $f(x) = x^a$ with $a \geq 2$. Hence, in the case of such functions f, the price of anarchy is 1, so the inefficiencies due to decentralization are fully mitigated by the network design. This is true, in particular, for Metcalfe's law.

The full version of this paper is available at [9].

ACKNOWLEDGMENTS

This work was supported by Polish National Science Centre through grant no 2014/13/B/ST6/01807. M. Skomra is supported by a grant from Région Ile-de-France.

REFERENCES

[1] D. Acemoglu, A. Malekian, and A. Ozdaglar. 2016. Network security and contagion. *J. Econom. Theory* 166 (2016), 536–585. https://doi.org/10.1016/j.jet.2016.09.009

[2] M. Aghassi and D. Bertsimas. 2006. Robust game theory. *Math. Program.* 107, 1 (2006), 231–273. https://doi.org/10.1007/s10107-005-0686-0

[3] J. Aspnes, K. Chang, and A. Yampolskiy. 2006. Inoculation strategies for victims of viruses and the sum-of-squares partition problem. *J. Comput. System Sci.* 72, 6 (2006), 1077–1093. https://doi.org/10.1016/j.jcss.2006.02.003

[4] R. Böhme and G. Schwartz. 2010. Modeling Cyber-Insurance: Towards A Unifying Framework. In *Proceedings of the 9th Workshop on the Economics of Information Security (WEIS 2010)*. http://www.econinfosec.org/archive/weis2010/papers/session5/weis2010_boehme.pdf

[5] D. Cerdeiro, M. Dziubiński, and S. Goyal. 2014. Individual security and network design. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC'14)*. ACM, New York, NY, 205–206. https://doi.org/10.1145/2600057.2602894

[6] D. Cerdeiro, M. Dziubiński, and S. Goyal. 2017. Individual security, contagion, and network design. *J. Econom. Theory* 170 (2017), 182–226. https://doi.org/10.1016/j.jet.2017.05.006

[7] H. Chan, M. Ceyko, and L. Ortiz. 2012. Interdependent Defense Games: Modeling Interdependent Security under Deliberate Attacks. In *Proceedings of the 28th Conference on Uncertainty in Artificial Intelligence (UAI 2012)*, N. de Freitas and K. Murphy (Eds.). AUAI Press, 152–162. http://auai.org/uai2012/papers/311.pdf

[8] S. Goyal, S. Jabbari, M. Kearns, S. Khanna, and J. Morgenstern. 2016. Strategic Network Formation with Attack and Immunization. In *Proceedings of the 12th Conference on Web and Internet Economics (WINE 2016) (Lecture Notes in Comput. Sci.)*, Y. Cai and A. Vetta (Eds.), Vol. 10123. Springer, Berlin, Heidelberg, 429–443. https://doi.org/10.1007/978-3-662-54110-4_30

[9] T. Janus, M. Skomra, and M. Dziubiński. 2018. Individual Security and Network Design with Malicious Nodes. (2018). arXiv:1804.07287

[10] H. Kunreuther and G. Heal. 2003. Interdependent Security. *J. Risk Uncertain.* 26, 2–3 (2003), 231–249. https://doi.org/10.1023/A:1024119208153

[11] A. Laszka, M. Felegyhazi, and L. Buttyán. 2015. A Survey of Interdependent Information Security Games. *Comput. Surveys* 47, 2, Article 23 (2015), 38 pages. https://doi.org/10.1145/2635673

[12] M. Lelarge and J. Bolot. 2008. A Local Mean Field Analysis of Security Investments in Networks. In *Proceedings of the 3rd International Workshop on Economics of Networked Systems (NetEcon'08)*. ACM, 25–30. https://doi.org/10.1145/1403027.1403034

[13] M. Lelarge and J. Bolot. 2008. Network Externalities and the Deployment of Security Features and Protocols in the Internet. *ACM SIGMETRICS Perform. Eval. Rev.* 36, 1 (2008), 37–48. https://doi.org/10.1145/1384529.1375463

[14] J. Lou, A. Smith, and Y. Vorobeychik. 2017. Multidefender Security Games. *IEEE Intelligent Systems* 32, 1 (2017), 50–60. https://doi.org/10.1109/MIS.2017.11

[15] J. Lou and Y. Vorobeychik. 2015. Equilibrium Analysis of Multi-Defender Security Games. In *Proceedings of the 24th International Conference on Artificial Intelligence (IJCAI'15)*. AAAI Press, 596–602.

[16] T. Moscibroda, S. Schmid, and R. Wattenhofer. 2009. The Price of Malice: A Game-Theoretic Framework for Malicious Behavior. *Internet Math.* 6, 2 (2009), 125–156. https://doi.org/10.1080/15427951.2009.10129181

[17] A. Smith, Y. Vorobeychik, and J. Letchford. 2014. Multi-Defender Security Games on Networks. *ACM SIGMETRICS Perform. Eval. Rev.* 41, 4 (2014), 4–7. https://doi.org/10.1145/2627534.2627536

[18] H. Varian. 2004. System Reliability and Free Riding. In *Economics of Information Security*, L. J. Camp and S. Lewis (Eds.). Adv. Inf. Secur., Vol. 12. Springer, Boston, MA, 1–15. https://doi.org/10.1007/1-4020-8090-5_1