Inducible Equilibrium for Security Games (Extended Abstract)

Qingyu Guo[1], Jiarui Gan[2], Fei Fang[3], Long Tran-Thanh[4], Milind Tambe[5], Bo An[1]
[1]School of Computer Science and Engineering, Nanyang Technological University, {qguo005, boan}@ntu.edu.sg
[2]Department of Computer Science, University of Oxford, Jiarui.gan@cs.ox.ac.uk
[3]School of Computer Science, Carnegie Mellon University, feifang@cmu.edu
[4]Department of Electronics and Computer Science, University of Southampton, ltt08r@ecs.soton.ac.uk
[5]Center for Artificial Intelligence in Society, University of Southern California, tambe@usc.edu

ABSTRACT

Strong Stackelberg equilibrium (SSE) is the standard solution concept of Stackelberg security games. The SSE assumes that the follower breaks ties in favor of the leader and this is widely acknowledged and justified by the assertion that the defender can *often* induce the attacker to choose a preferred action by making an infinitesimal adjustment to her strategy. Unfortunately, in security games with resource assignment constraints, the assertion might not be valid. To overcome this issue, inspired by the notion of inducibility and the pessimistic Stackelberg equilibrium [20, 21], this paper presents the *inducible Stackelberg equilibrium* (ISE), which is guaranteed to exist and avoids overoptimism as the outcome can *always* be induced with infinitesimal strategy deviation. Experimental evaluation unveils the significant overoptimism and sub-optimality of SSE and thus, verifies the advantage of the ISE as an alternative solution concept.

KEYWORDS

Security games; inducible Stackelberg equilibrium; utility guarantee

1 INTRODUCTION

The past decade has witnessed the huge success of game theoretic reasoning in complex security domains [1, 2, 4, 12, 15, 17]. Various applications based on the Stackelberg security game (SSG) model have been deployed to protect airports, ports, wildlife and so on [8, 17]. The standard solution concept in Stackelberg games is Stackelberg equilibrium [11], which assumes that both players are rational and have no incentive to deviate in the equilibrium. The *strong* form of the Stackelberg equilibrium, called Strong Stackelberg Equilibrium (SSE) assumes that the follower will always break ties in favor of the defender and is the most commonly adopted concept in related literature [7, 14, 20] and in most security game applications [17]. In essence, researchers have implicitly or explicitly claimed or assumed that the defender can always induce the favorable strong equilibrium by selecting a strategy arbitrarily close to the equilibrium [3, 6, 10, 19].

However, the assertion that the defender can always induce SSE may break in security domains with resource assignment constraints, e.g., protecting flights with air marshals (FAMS) [10, 18], protection ports [16], protecting targets with externalities [9]. Unfortunately, existing research has failed to realize the potential impossibility to induce SSE in such domains. If the desired SSE cannot be induced, the results claimed in existing works may be overly optimistic. Such overoptimism is highly problematic since these results may be used in making security resource acquisition decisions [13], and the SSE strategy recommended may not be the optimal one, thus failing in the primary mission of security games, which is to optimize the use of limited security resources.

In this paper, we offer remedies for this shortcoming. First, we formalize the notion of overoptimism by defining the utility guarantee of the defender's strategies, and show with a motivating example that the utility claimed to be guaranteed by the SSE is much higher than the actually guaranteed utility. Inspired by the notion of inducible strategy [20] and the pessimistic Stackelberg equilibrium [21], we propose a new solution concept for security games called *inducible Stackelberg equilibrium* (ISE) based on a novel tie-breaking rule. ISE possesses nice properties that it is guaranteed to exist and avoids overoptimism as it offers the defender the highest guaranteed utility. Second, we prove that the problem of computing an ISE polynomially reduces to that of computing an SSE and thus, introducing the ISE does not invalidate existing algorithmic results. We also provide algorithmic implementation for computing the ISE and conduct experiments to evaluate our results; our experiments unveil the significant overoptimism and sub-optimality of the SSE, which suggests the practical significance of the ISE solution.

2 MOTIVATING EXAMPLE

This paper focuses on security games with arbitrary schedules (SPARS) model [10]. Consider the SPARS instance shown in the following figure where there are four targets, i.e., $T = \{t_1, t_2, t_3, t_4\}$. The defender has one resource $R = \{r\}$. For a target $t \in T$, the defender's payoff for an uncovered attack is denoted by $U_d^u(t)$ and for a covered attack $U_d^c(t)$. Similarly, $U_a^u(t)$ and $U_a^c(t)$ are attacker's payoffs respectively. These payoffs are depicted in the figure. We first consider the scenario without resource assignment constraints, which has a unique SSE with coverage strategy $\mathbf{c} = \langle \frac{4}{15}, \frac{1}{5}, \frac{4}{15}, \frac{4}{15} \rangle$, where c_t represents the marginal probability that t is covered by a defender resource. In SSE, all targets in T have the same expected utility for the attacker and thus form a tie, denoted as $\Gamma(\mathbf{c}) = T$. The tie-breaking rule in SSE indicates that the attacker will break the tie $\Gamma(\mathbf{c}) = T$ by attacking t_2. This can be induced by decreasing the coverage on t_2 with infinitesimal amount and increasing the coverage on other targets, making t_2 be strictly preferred.

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

https://doi.org/doi

However, with resource assignment constraints, the defender will not be able to decrease the coverage on one target arbitrarily while simultaneously not decreasing coverage on all other targets. Suppose $J = \{s_1, s_2, s_3, s_4\}$ as shown in the figure. (There is only one resource.) The game still has a unique SSE where the defender plays $\mathbf{x} = \langle \frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0 \rangle$ and the attacker is assumed to attack t_2, bringing the defender an expected utility of $-\frac{1}{3}$. Such outcome is explicitly or implicitly considered with previous mentioned infinitesimal strategy deviation in security game literature [10]. Unfortunately, there exists no strategy arbitrarily close to $\mathbf{x}$ which makes t_2 be strictly preferred by the attacker. If x_1 is decreased, the attacker will prefer t_1 over t_2; otherwise t_3 or t_4 will be attacked. That is to say, any infinitesimal strategy deviation will cause the attacker to attack t_1, t_3 or t_4. The best induced outcome for the defender is only arbitrarily close to $-\frac{2}{3}$, achieved by decreasing x_1 with infinitesimal amount and the attacker is induced to attack t_1.

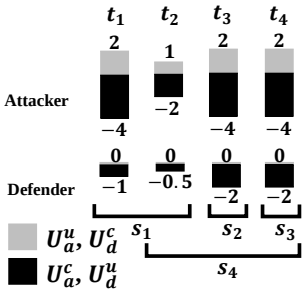

3 INDUCIBLE STACKELBERG EQUILIBRIUM

The above example reveals a failure of the attempt to induce the desired SSE outcome by playing a strategy arbitrarily close to the SSE strategy. To formalise this situation, we propose the notion called *utility guarantee*. Let $\mathcal{X}$ be the strategy space of the defender, and consider only the pure strategy for the attacker. Let $U_* : \mathcal{X} \times T \to \mathbb{R}$ be the expected utility function for player $*$ being the attacker (a) and the defender (d) respectively. Let $\Gamma(\mathbf{x}) = \arg\max_{t \in T} U_a(\mathbf{x}, t)$ denote the *attack set* w.r.t. $\mathbf{x}$.

Definition 1 (**Utility Guarantee**). *We say an expected utility v can be guaranteed by defender's mixed strategy $\mathbf{x}$ iff: $\forall \epsilon > 0, \forall \delta > 0, \exists \mathbf{x}' \in \mathcal{X}$ such that $\|\mathbf{x} - \mathbf{x}'\| \le \delta$ and $U_d(\mathbf{x}', f^W(\mathbf{x}')) \ge v - \epsilon$, where $f^W : \mathcal{X} \to T$ satisfies $f^W(\mathbf{x}') \in \arg\min_{t \in \Gamma(\mathbf{x}')} U_d(\mathbf{x}', t)$ for all $\mathbf{x}' \in \mathcal{X}$. Let $U^g(\mathbf{x}) \subset \mathbb{R}$ be the set of utilities guaranteed by $\mathbf{x}$. $\sup U^g(\mathbf{x})$ is called the utility guarantee of $\mathbf{x}$.*

In other words, if a strategy $\mathbf{x}$ guarantees a utility value v, the defender can obtain an expected utility at least arbitrarily close to v by playing a strategy arbitrarily close to $\mathbf{x}$, *regardless* of how the attacker actually breaks the tie (the spirit of "guarantee"). As shown in the motivating example, the utility of an SSE strategy might fail to be guaranteed. This results in overoptimism and a suboptimal solution.

To remedy the overoptimism, we propose a new solution concept called *Inducible Stackelberg Equilibrium (ISE)*, based on the notion of *inducibility* [20].

Definition 2 (**Inducible Target**). *A target t is inducible iff there exists at least one defender's mixed strategy $\mathbf{x} \in \mathcal{X}$ such that t is the unique best response target against $\mathbf{x}$.*

We denote by $T^i = \{t \in T | \exists \mathbf{x} \in \mathcal{X} : \Gamma(\mathbf{x}) = \{t\}\}$ the set of inducible targets. ISE is a profile $\langle \mathbf{x}, f^I(\mathbf{x}) \rangle$ where $f^I : \mathcal{X} \to T$ satisfies that $f^I(\mathbf{x}) \in \arg\max_{t \in \Gamma(\mathbf{x}) \cap T^i} U_d(\mathbf{x}, t)$, and $\mathbf{x} \in \arg\max_{\mathbf{x}' \in \mathcal{X}} U_d(\mathbf{x}', f^I(\mathbf{x}'))$. Comparing ISE with SSE, we notice that in ISE, the attacker also breaks the ties in favor of the defender, and the only difference is that the attacker is restricted to attack the inducible targets in ties. It turns out that the inducible targets characterize the highest utility guarantee achieved by any mixed strategy of defender.

One may notice that the ISE strategy coincides with the so called *pessimistic leader-follower equilibrium* [5, 21]. In fact, ISE is a generalization of the pessimistic Stackelberg equilibrium in the context of security games where the notion of utility guarantee is proposed to formalise the inducibility issue we observed, and the tie-breaking rule f^I is provided to make the solution consistent with SSEs in security games literature. Following the similar analysis in [5, 21], we prove several nice properties of ISE, namely, existence and the optimality w.r.t. the utility guarantee, making ISE an appealing alternative of SSE to overcome the potential overoptimism.

4 EXPERIMENTAL EVALUATION

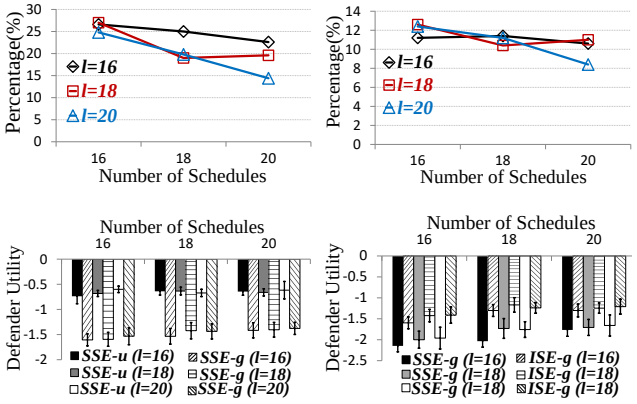

Figure 1: Overoptimism and sub-optimality of SSE.

We conduct experiments to evaluate the overoptimism (providing a non-guaranteed utility) and the sub-optimality (w.r.t. the utility guarantee) of SSE. The rewards and penalties are all integers randomly drawn from $[0, 5]$ and $[-5, 0]$ respectively. The game has 200 targets, 1 resource, number of schedules $|S| \in \{16, 18, 20\}$ and number of targets covered by each schedule $l \in \{16, 18, 20\}$. As shown in Figure 1, SSE suffers from overoptimism and sub-optimality since a significant proportion of random instances are spotted with overoptimistic and/or sub-optimal SSEs. Moreover, once the overoptimism and sub-optimality occur on an SSE, the actual utility guarantee ("SSE-g") is significantly less than the provided amount by SSE ("SSE-u"), and the guaranteed utility of SSE is significantly less than the optimal utility in ISE ("ISE-g").

ACKNOWLEDGMENTS

This research was supported by MURI grant W911NF-17-1-0370 and the National Research Foundation, Prime Minister's Office, Singapore under its IDM Futures Funding Initiative. Long Tran-Thanh was supported by the EPSRC funded project STRICT (EP/N02026X/1).

REFERENCES

[1] Maria-Florina Balcan, Avrim Blum, Nika Haghtalab, and Ariel D. Procaccia. 2015. Commitment Without Regrets: Online Learning in Stackelberg Security Games. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, (EC'15)*. 61–78.

[2] Nicola Basilico, Andrea Celli, Giuseppe De Nittis, and Nicola Gatti. 2017. Coordinating Multiple Defensive Resources in Patrolling Games with Alarm Systems. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems (AAMAS'17)*. 678–686.

[3] Nicola Basilico, Nicola Gatti, and Francesco Amigoni. 2009. Leader-follower strategies for robotic patrolling in environments with arbitrary topologies. In *8th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'09)*. 57–64.

[4] Avrim Blum, Nika Haghtalab, and Ariel D. Procaccia. 2014. Learning Optimal Commitment to Overcome Insecurity. In *Advances in Neural Information Processing Systems (NIPS'14)*. 1826–1834.

[5] Stefano Coniglio, Nicola Gatti, and Alberto Marchesi. 2017. Pessimistic Leader-Follower Equilibria with Multiple Followers. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI'17*. 171–177.

[6] Vincent Conitzer. 2012. Computing Game-Theoretic Solutions and Applications to Security. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI'12)*. 2106–2112.

[7] Vincent Conitzer and Tuomas Sandholm. 2006. Computing the optimal strategy to commit to. In *Proceedings 7th ACM Conference on Electronic Commerce (EC'06)*. 82–90.

[8] Fei Fang, Thanh Hong Nguyen, Rob Pickles, Wai Y. Lam, Gopalasamy R. Clements, Bo An, Amandeep Singh, Milind Tambe, and Andrew Lemieux. 2016. Deploying PAWS: Field Optimization of the Protection Assistant for Wildlife Security. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI'16)*. 3966–3973.

[9] Jiarui Gan, Bo An, and Yevgeniy Vorobeychik. 2015. Security Games with Protection Externalities. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, (AAAI'15)*. 914–920.

[10] Manish Jain, Erim Kardes, Christopher Kiekintveld, Fernando Ordóñez, and Milind Tambe. 2010. Security Games with Arbitrary Schedules: A Branch and Price Approach. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI'10)*. 792–797.

[11] George Leitmann. 1978. On generalized Stackelberg strategies. *Journal of Optimization Theory and Applications* 26, 4 (1978), 637–643.

[12] Bo Li and Yevgeniy Vorobeychik. 2014. Feature Cross-Substitution in Adversarial Classification. In *Advances in Neural Information Processing Systems (NIPS'14)*. 2087–2095.

[13] Sara Marie McCarthy, Milind Tambe, Christopher Kiekintveld, Meredith L. Gore, and Alex Killion. 2016. Preventing Illegal Logging: Simultaneous Optimization of Resource Teams and Tactics for Security. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI'16)*. 3880–3886.

[14] Martin J Osborne and Ariel Rubinstein. 1994. *A Course in Game Theory*. MIT Press.

[15] Ariel Rosenfeld and Sarit Kraus. 2017. When Security Games Hit Traffic: Optimal Traffic Enforcement Under One Sided Uncertainty. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI'17)*. 3814–3822.

[16] Eric Shieh, Bo An, Rong Yang, Milind Tambe, Craig Baldwin, Joseph DiRenzo, Ben Maule, and Garrett Meyer. 2012. PROTECT: A deployed game theoretic system to protect the ports of the United States. In *International Conference on Autonomous Agents and Multiagent Systems, (AAMAS'12)*. 13–20.

[17] Milind Tambe. 2011. *Security and Game Theory - Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press.

[18] Jason Tsai, Christopher Kiekintveld, Fernando Ordonez, Milind Tambe, and Shyamsunder Rathi. 2009. IRIS-A tool for strategic security allocation in transportation networks. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems Industry Track (AAMAS'09)*. 37–44.

[19] Pradeep Varakantham, Hoong Chuin Lau, and Zhi Yuan. 2013. Scalable Randomized Patrolling for Securing Rapid Transit Networks. In *Proceedings of the Twenty-Fifth Innovative Applications of Artificial Intelligence Conference (IAAI'13)*. 1563–1568.

[20] Bernhard Von Stengel and Shmuel Zamir. 2004. Leadership with commitment to mixed strategies. *Technical Report LSE-CDAM-2004-01, CDAM Research Report*. (2004).

[21] Bernhard von Stengel and Shmuel Zamir. 2010. Leadership games with convex strategy sets. *Games and Economic Behavior* 69, 2 (2010), 446–457.