Minimum Support Size of the Defender's Strong Stackelberg Equilibrium Strategies in Security Games

Jiarui Gan
University of Chinese Academy of Sciences
The Key Lab of Intelligent Information Processing, ICT, CAS
Beijing 100190, China
ganjr@ics.ict.ac.cn

Bo An
Nanyang Technological University
Singapore 639798
boan@ntu.edu.sg

Abstract

Stackelberg security games have been applied to address challenges in security resource allocation of real-world infrastructure protection tasks. The key to such an application is to efficiently compute the defender's optimal strategy in consideration of the attacker's surveillance capability and best response. Experimental results show that the defender's optimal strategy often uses only a small subset of pure strategies, as compared with the entire pure strategy set which can be exponentially large. A number of algorithms in the literature have already exploited this small support size observation. This paper analyzes a number of widely studied security games and provides bounds on the minimum support size of the defender's Strong Stackelberg Equilibrium (SSE) strategies in security games.

Introduction

Stackelberg security games have been used to model many real-world scenarios where a defender commits to a strategy and an attacker makes its attacking decision with knowledge of the defender's commitment. Systems applying Stackelberg game models to assist with randomized resource allocation decisions have been developed and are currently in use, such as: ARMOR, developed at the Los Angeles International Airport (LAX) to randomize checkpoints on the roadways entering the airport and canine patrol routes within the airport terminals (Pita et al. 2008); IRIS, used by the Federal Air Marshals Service (FAMS) as a scheduler for randomized deployment (Tsai et al. 2009); PROTECT, used by the US Coast Guard (USCG) to randomize patrolling at the port of Boston (Shieh et al. 2012); and GUARDS, used by the United States Transportation Security Administration (TSA) to scheduling resources to protect airports in the USA (An et al. 2011b; Pita et al. 2011).

The core of the above applications is computing of the defender's optimal strategy, where difficulty lies in the large scale of the defender's strategy space. Specifically, a defender plays a mixed strategy which is a probabilistic distribution over a set of pure strategies. The pure strategy set can be very large due to combinatorial explosion. For example, in a security game where the defender allocates security resources to cover a set of n targets, the defender could have 2^n pure strategies. How to optimize the defender's strategy over such a large scale pure strategy space is thus a big challenge. Although algorithms have been designed to speed up the computing process or to scale up the model (Paruchuri et al. 2008; Kiekintveld et al. 2009; Jain et al. 2010; Korzhyk, Conitzer, and Parr 2010; An et al. 2011c; Jain et al. 2011) it still remains as an open research area to design more efficient algorithms to deal with scalability and uncertainty (An et al. 2011a; 2012; Brown et al. 2012; Yang, Ordonez, and Tambe 2012; Yin and Tambe 2012; An et al. 2013).

It has been noticed that only pure strategies with non-zero probabilities in a mixed strategy contribute to the implementation of the mixed strategy while the others are unused. We call the set of pure strategies with non-zero probabilities the *support* of the mixed strategy. Experimental results show that the defender's optimal strategy often has a small support, as compared with the entire pure strategy set which can be exponentially large (Shieh et al. 2012). A number of algorithms in the literature have already exploited this small support size observation, and particularly avoided enumerating the entire set of pure strategies while computing the defender's optimal strategy. For example, a column generation method adds one best pure strategy into the support each time, until the solution cannot be improved (e.g., (Jain et al. 2010)); and a double-oracle based approach (Jain et al. 2011; Jain, Conitzer, and Tambe 2013), where the players play a series of games iteratively using a subset of pure strategies until convergence. One interesting question lies behind these applications that is yet not answered to the best of our knowledge is how small the support of the defender's optimal strategy can be.

To answer the above question, this paper analyzes the structure of the defender's strategy in a number of widely studied security games and provides bounds on the minimum support size of the defender's Strong Stackelberg Equilibrium (SSE) strategies in security games. The rest of the paper is organized as follows. We first introduce the Stackelberg game and Stackelberg equilibria. Then starting with a typical Stackelberg security game which can be compactly represented, we present the bound of the minimum support size of the defender's SSE strategies for a general Stackelberg game and, more generally, for a Bayesian Stackelberg game

which allows for multiple types of attackers. In the end, we apply the bound to other types of specialized Stackelberg security games.

Stackelbeg Games and Equilibria

A Stackelberg game is played by two players: a leader and a follower. The leader acts first and the follower observes the leader's strategy before taking an action. Each player has a set of pure strategies, denoted as $\mathcal{S}$ and $\mathcal{T}$ for the leader and the follower, respectively. Let s_j be the j^{th} pure strategy in $\mathcal{S}$ and t_i be the i^{th} pure strategy in $\mathcal{T}$. When the leader plays s_j and the follower plays t_i, they receive payoffs $U_l^{j,i}$ and $U_f^{j,i}$, respectively. Furthermore, the leader commits a mixed strategy $\boldsymbol{x} = \langle x_j \rangle$ which is a probabilistic distribution over the set $\mathcal{S}$ of pure strategies. Similarly, the attacker commits a mixed strategy $\boldsymbol{y} = \langle y_i \rangle$. In this case, the expected payoff for the leader and the follower are defined respectively as follows:

$$U_l(\boldsymbol{x}, \boldsymbol{y}) = \sum_{i=1}^{|\mathcal{T}|} y_i \sum_{j=1}^{|\mathcal{S}|} x_j U_l^{j,i} \tag{1}$$

$$U_f(\boldsymbol{x}, \boldsymbol{y}) = \sum_{i=1}^{|\mathcal{T}|} y_i \sum_{j=1}^{|\mathcal{S}|} x_j U_f^{j,i} \tag{2}$$

Stackelberg Equilibria In a Stackelberg equilibrium, the follower observes the leader's strategy $\boldsymbol{x}$ and responds with strategy $f(\boldsymbol{x}) : \boldsymbol{x} \rightarrow \boldsymbol{y}$ that is optimal with respect to his expected payoff, i.e., $U_f(\boldsymbol{x}, f(\boldsymbol{x})) \geq U_f(\boldsymbol{x}, \boldsymbol{y})$ for all other feasible strategies $\boldsymbol{y}$; the leader, knowing the payoff-driven behavior of the follower, chooses an optimal strategy $\boldsymbol{x}^*$ that maximizes his payoff, i.e., $U_l(\boldsymbol{x}^*, f(\boldsymbol{x}^*)) \geq U_l(\boldsymbol{x}, f(\boldsymbol{x}))$ for all other feasible strategies $\boldsymbol{x}$. Typically, there might be more than one optimal strategy for the follower, and each of these strategies may induce a different payoff for the leader. Two types of Stackelberg equilibrium, the Strong Stackelberg Equilibrium (SSE) and the Weak Stackelberg Equilibrium (WSE), were introduced (Leitmann 1978; Breton, Alj, and Haurie 1988), where SSE assumes the follower breaks ties in favor of the defender and chooses the one that is optimal for the leader, while WSE assumes that the follower chooses the worst one for the leader. Most exiting work takes SSE as the solution concept as most literatures in this research filed did, since 1) SSE exists in all Stackelberg games, while WSE may not (Basar et al. 1995); 2) SSE can always be induced by the leader through deviating infinitesimal from the optimal strategy (Von Stengel and Zamir 2004).

Definition 1. *A pair of strategies $(\boldsymbol{x}, f(\boldsymbol{x}))$ forms a Strong Stackelberg Equilibrium (SSE) if it satisfies the following:*

1. *The leader plays a best-response: $U_l\big(\boldsymbol{x}, f(\boldsymbol{x})\big) \geq U_l\big(\boldsymbol{x}', f(\boldsymbol{x}')\big)$, for all the leader strategies $\boldsymbol{x}'$.*
2. *The follower plays a best-response: $U_f\big(\boldsymbol{x}, f(\boldsymbol{x})\big) \geq U_f(\boldsymbol{x}, \boldsymbol{y}')$, for all the follower strategies $\boldsymbol{y}'$.*
3. *The follower breaks ties in favor of the leader: $U_l\big(\boldsymbol{x}, f(\boldsymbol{x})\big) \geq U_l(\boldsymbol{x}, \boldsymbol{y}')$, for all optimal follower strategies $\boldsymbol{y}'$.*

Specially, we restrict the attacker's strategy to pure strategies, i.e., $f(\boldsymbol{x}) : \boldsymbol{x} \rightarrow t$, $t \in \mathcal{T}$, because according to Eq. (2) if one mixed strategy is optimal for the attacker, then all the pure strategies with non-zero probabilities should be optimal, so that the attacker always has a best pure-strategy response. It follows that the payoff functions defined in Eqs. (1) and (2) can be rewritten as (assuming the attacker attacks t_i):

$$U_l(\boldsymbol{x}, t_i) = \sum_{j=1}^{|\mathcal{S}|} x_j U_l^{ji}, \tag{3}$$

$$U_f(\boldsymbol{x}, t_i) = \sum_{j=1}^{|\mathcal{S}|} x_j U_f^{ji}. \tag{4}$$

Bound the Minimum Support Size of the Leader's SSE Strategies

In this section, we exploit the structure of a leader strategy and presents an upper bound on the size of the minimum support (defined below) of the leader's SSE strategies. For ease of description, we use Φ to represent the upper bound on the minimum support size of the leader's SSE strategies. Denote $\|\boldsymbol{x}\|$ as the support size of a mixed strategy $\boldsymbol{x}$, we have

$$\min_{\boldsymbol{x} \in \mathcal{X}^*} \|\boldsymbol{x}\| \leq \Phi,$$

where $\mathcal{X}^*$ is the set of leader's SSE strategies. In other words, the leader always has an SSE strategy that can be implemented by no more than Φ pure strategies.

Definition 2. *A pure strategy is a **support strategy** of a mixed strategy if it is assigned with a non-zero probability by this mixed strategy. The set of all support strategies is called the **support** of the mixed strategy.*

We start from a special Stackelberg security game model presented below, where a leader's strategy can be represented compactly as a coverage vector, since it provides an easy-to-follow example of how large the pure strategy set could be and how tight the support could be bounded. We call this game model *compact game model* so as to distinguish it from other types of Stackelberg security games. Then with a similar idea that Φ for a compact game is obtained, we generalize it and present Φ for a general Stackelberg game and furthermore a Bayesian Stackelberg game.

Compact Game Model

The compact game model applies to many security domains (e.g., Pita et al. 2008; Tsai et al. 2009). In this game model, the leader is a *defender* who protects a set of targets, and the follower is an *attacker* who wants to attack a target. Let the set of targets be $1, \ldots, n$. The defender allocates security resources to protect the targets, and his pure strategy is an allocation of the recourses. Typically, in the compact game model, a target is either *covered* or *uncovered* by security resources, and when it is covered, adding more resources to

it makes no difference. In this case, a rational defender assigns at most one resource to a target, and a pure strategy of the defender can thus be defined as a 0/1 coverage vector with the j^{th} pure strategy $\boldsymbol{s}_j = \langle s_{ji} \rangle \in \{0,1\}^n$, where $s_{ji} = 1$ represents that target i is covered and $s_{ji} = 0$ uncovered. Correspondingly, the attacker's pure strategy is to choose one target to attack. Let t_i be the attacker's pure strategy of attacking target i. When the attacker plays t_i and target i is uncovered, he receives utility $U_a^0(t_i)$, and the defender receives utility $U_d^0(t_i)$. Similarly, when the attacker plays t_i and target i is covered, he receives utility $U_a^1(t_i)$, and the defender receives $U_d^1(t_i)$. It follows that when the defender plays a mixed strategy $\boldsymbol{x}$, the targets are covered with probabilities $\boldsymbol{c}(\boldsymbol{x}) = \langle c_i \rangle = \sum_{\boldsymbol{s}_j \in \mathcal{S}} x_j \boldsymbol{s}_j$ with c_i for target i, and are uncovered with $\mathbf{1} - \boldsymbol{c}$. We refer to $\boldsymbol{c}$ as the *coverage vector*. For a strategy profile $\langle \boldsymbol{x}, t_i \rangle$, the expected utility for the defender and the attacker can be defined respectively as follows:

$$U_d\big(\boldsymbol{c}, t_i\big) = c_i U_d^1(t_i) + \big(1 - c_i\big) U_d^0(t_i), \tag{5}$$

$$U_a\big(\boldsymbol{c}, t_i\big) = c_i U_a^1(t_i) + \big(1 - c_i\big) U_a^0(t_i). \tag{6}$$

Generally, some resource restrictions, such as scheduling constraints, can be enforced on the defender's pure strategy set $\mathcal{S}$. However, even if in the presence of such restrictions, the size of $\mathcal{S}$ may still be exponentially large in terms of the number of targets. For example, when there are m available resources, the size of $\mathcal{S}$ is at least $\binom{n}{m}$.

Φ for a Compact Game According to Eqs. (5) and (6), two defender strategies $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$ results in the same attacker response and moreover the same utility for each player, if they induce the same coverage vector, i.e., $\boldsymbol{c}(\boldsymbol{x}_1) = \boldsymbol{c}(\boldsymbol{x}_2)$. We utilize this observation to seek equivalent defender strategies with smaller supports. For example, when there are three targets and the defender plays a mixed strategy where pure strategies $(1,0,0)$, $(1,1,0)$, $(0,1,0)$ and $(0,1,1)$ are assigned with probabilities 0.25 for each, a coverage vector $(0.5, 0.5, 0.5)$ is induced. In this case, it is also possible to use only pure strategies $(1,0,0)$ and $(0,1,1)$ with probability 0.5 for each and all the others with probability 0, which induces the same coverage vector. In fact, a defender strategy first induces a coverage vector and then affects the game through this coverage vector. Therefore, rather than specifying an exact mixed strategy, the defender could first calculate an optimal coverage vector that is implementable by his feasible mixed strategies, i.e.,

$$\boldsymbol{c}^* \in \mathbf{P} = \left\{ \sum_{\boldsymbol{s}_j \in \mathcal{S}} x_j \boldsymbol{s}_j \;\middle|\; \boldsymbol{x} \succeq \mathbf{0}, \mathbf{1}^\mathsf{T} \boldsymbol{x} = 1 \right\},$$

such that

$$U_d\big(\boldsymbol{c}^*, f(\boldsymbol{c}^*)\big) \geq U_d\big(\boldsymbol{c}', f(\boldsymbol{c}^*)\big), \forall \boldsymbol{c}' \in \mathbf{P},$$

and then implement $\boldsymbol{c}^*$ with a mixed strategy which is optimal for the defender in this case. Note that the implementable coverage vector set $\mathbf{P}$ is a convex hull in an n-dimensional space defined by points in $\mathcal{S}$. According to the Carathéodory's theorem (Danninger-Uchida 2001), any point, in particular $\boldsymbol{c}^*$, in the convex hull lies in an r-simplex with vertices in $\mathcal{S}$, where $r \leq n$, and any point in the r-simplex can be represented as a convex combination of the $r+1$ vertices of the simplex. Therefore, $\boldsymbol{c}^*$ can always be implemented by no more than $n+1$ pure strategies corresponding to the vertices of the simplex where $\boldsymbol{c}^*$ lies in, and an upper bound $\Phi = n+1$ is obtained for the compact game model (Corollary 2).

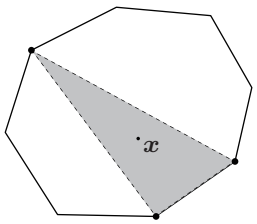

Figure 1: Carathéodory's theorem: in a plane (i.e., a 2-dimensional space), any point $\boldsymbol{x}$ in a convex hull also lies in a 2-simplex.

Theorem 1 (Carathéodory's theorem (Danninger-Uchida 2001))**.** *If a point $\boldsymbol{x} \in \mathbb{R}^d$ lies in the convex hull of a point set $\mathcal{P}$, there is a subset $\mathcal{P}'$ of $\mathcal{P}$ consisting of $d+1$ or fewer points such that $\boldsymbol{x}$ lies in the convex hull of $\mathcal{P}'$, i.e., $\boldsymbol{x}$ lies in an r-simplex with vertices in $\mathcal{P}$, where $r \leq d$ (Figure 1).*

Corollary 2. *The minimum support size of the defender's SSE strategies in a compact game is $n+1$, where n is the number of targets.*

Φ for a General Stackelberg Game

The compact game model is a special case of Stackelberg security games. In other application of Stackelberg game models, the defender's strategies over the targets may be more complex than being simply either covering or not covering. For example, in some real-world scenarios, it makes a difference when different number or different types of security resources are assigned to a target (Pita et al. 2011; Shieh et al. 2012); and in a network-based game (Washburn and Wood 1995; Tsai et al. 2010; Jain et al. 2011; Jain, Conitzer, and Tambe 2013), the defender places security resources on the edges (e.g., streets, roads) that lead to the target, instead of directly on the targets (more concrete examples are presented following the definitions of the models in the next section). In these cases, the defender strategy cannot be represented as coverage over the targets, and Corollary 2 is thus not applicable. In this section, we present Φ for a general Stackelberg game using a similar idea that Φ for a compact game is obtained. This general Φ applies to all security games derived from Stackelberg games, which we exemplify in the next section.

According to Eqs. (3) and (4), if two leader strategies $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$ satisfy $U_f(\boldsymbol{x}_1, t_i) = U_f(\boldsymbol{x}_2, t_i)$ and $U_l(\boldsymbol{x}_1, t_i) = U_l(\boldsymbol{x}_2, t_i)$, $\forall\, i = 1, \ldots, |\mathcal{T}|$, the follower would respond with the same pure strategy according to U_f when they are played (in particular, when there are multiple optimal pure

strategies for the follower, he refers to U_l and also has the same response) and receive the same payoff according to U_l; and the follower's response would induce the same leader payoff; $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$ are thus equivalent for both players in terms of the payoffs induced. We write the above conditions as follows:

$$\sum_{j=1}^{|\mathcal{S}|} x_{1j}\boldsymbol{u}_j = \sum_{j=1}^{|\mathcal{S}|} x_{2j}\boldsymbol{u}_j,$$

where $\boldsymbol{u}_j = (U_l^{j1}, \ldots, U_l^{j|\mathcal{T}|}, U_f^{j1}, \ldots, U_f^{j|\mathcal{T}|})^\mathsf{T}$ is a $2|\mathcal{T}|$-dimensional vector. Therefore, if some strategy profile $\langle \boldsymbol{x}^*, t_{i^*} \rangle$ forms SSE, another strategy profile $\langle \boldsymbol{x}', t_{i^*} \rangle$, such that $\sum_{j=1}^{|\mathcal{S}|} x'_j \boldsymbol{u}_j = \sum_{j=1}^{|\mathcal{S}|} x^*_j \boldsymbol{u}_j$, also forms SSE. This is equivalent to implementing the $2|\mathcal{T}|$-dimensional vector $\sum_{j=1}^{|\mathcal{S}|} x^*_j \boldsymbol{u}_j$ with the set of points $\{\boldsymbol{u}_1, \ldots, \boldsymbol{u}_{|\mathcal{S}|}\}$ as their convex combination. Similar to the analysis in the last section, according to the Carathéodory's theorem, there is always an implementation where no more than $2|\mathcal{T}|+1$ points have non-zero weights, which corresponds to an SSE strategy for the leader with a support of less than $2|\mathcal{T}| + 1$ pure strategies, namely, $\Phi = 2|\mathcal{T}| + 1$ for a general Stackelberg game.

A Tighter Φ Deeper analysis indicates that $\Phi = 2|\mathcal{T}| + 1$ can be even tighter. To present the tighter Φ, we first assume that there is only one optimal strategy for the follower, and we relax this assumption later. In this case, for two leader strategies $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$, the follower responds with the same optimal strategy if $U_f(\boldsymbol{x}_1, t_i) = U_f(\boldsymbol{x}_2, t_i), \forall\, i = 1, \ldots, |\mathcal{T}|$ (note that the condition $U_l(\boldsymbol{x}_1, t_i) = U_l(\boldsymbol{x}_2, t_i)$, $\forall\, i = 1, \ldots, |\mathcal{T}|$ is not necessary since there is no tie by assumption); and the leader receives the same payoff if $U_l(\boldsymbol{x}_1, t^*) = U_l(\boldsymbol{x}_2, t^*)$, where t^* is the follower's optimal strategy. Now given the SSE strategy profile $\langle \boldsymbol{x}^*, t_{i^*} \rangle$, any strategy profile $\langle \boldsymbol{x}', t_{i^*} \rangle$, such that

$$\sum_{j=1}^{|\mathcal{S}|} x'_j \boldsymbol{u}^\diamond_j = \sum_{j=1}^{|\mathcal{S}|} x^*_j \boldsymbol{u}^\diamond_j, \tag{7}$$

also forms SSE, where $\boldsymbol{u}^\diamond_j = (U_f^{j1}, \ldots, U_f^{j|\mathcal{T}|}, U_l^{ji^*})^\mathsf{T}$. The rest is the same with the proof of the previous Φ, so that $\Phi = |\mathcal{T}| + 2$ given the assumption that there is only one optimal strategy for the follower.

Next, we relax the assumption on the follower's optimal strategies and show that the above conclusion still holds. We show that a strategy profile $\langle \boldsymbol{x}', t_{i^*} \rangle$ with $\boldsymbol{x}'$ satisfying Eq. (7) still forms SSE. When the assumption is relaxed and the leader plays $\boldsymbol{x}'$, the follower may choose another pure strategy $t_{i'}$ that is also optimal for him. There are following cases: 1) $\langle \boldsymbol{x}', t_{i'} \rangle$ gives the leader the same payoff as $\langle \boldsymbol{x}^*, t_{i^*} \rangle$, then it is also optimal for the leader and satisfies the SSE condition; 2) $\langle \boldsymbol{x}', t_{i'} \rangle$ gives the leader a different payoff, which is after all impossible because if it gives the leader a higher payoff, then it is better than $\langle \boldsymbol{x}^*, t_{i^*} \rangle$, which contradicts that $\langle \boldsymbol{x}^*, t_{i^*} \rangle$ forms SSE; if it gives the leader a lower payoff, then the follower should not choose it because he break ties in favor of the leader. Therefore, $\boldsymbol{x}'$ is the leader's optimal strategy, a tighter Φ of $|\mathcal{T}| + 2$ is obtained (Corollary 3).

Corollary 3. *The minimum support size of the leader's SSE strategies in a general Stackelberg game is $|\mathcal{T}| + 2$, where $\mathcal{T}$ is the set of follower's pure strategies.*

Generalize Attacker Types: Φ for a Bayesian Stackelberg Game

A Bayesian Stackelberg game allows multiple types of leaders and followers. Typically, the leader type is restricted to one for the security game interest. The Bayesian Stackelberg games arise in scenarios where the leader has uncertain knowledge about different types of followers she may face (Paruchuri et al. 2008; Jain, Kiekintveld, and Tambe 2011). Although Corollary 3 can be applied to a Bayesian Stackelberg game by using the Harsanyi transformation to transform multiple followers to a single follower (Harsanyi and Selten 1972), the single follower's pure strategy space is the cross product of each follower type's pure strategy set, which can be exponentially large. The Φ obtained may thus be meaningless. In the follows we present a tighter Φ for a Bayesian Stackelberg game, not applying Corollary 3 directly but using the same core idea.

Let there be Λ types of followers. A follower of type λ occurs with probability p_λ and has a set $\mathcal{T}_\lambda$ of pure strategies indexed by $\mathcal{I}_\lambda = \{1, \ldots, |\mathcal{T}_\lambda|\}$. Given a strategy profile $\langle \boldsymbol{x}, I \rangle$, where $\boldsymbol{x}$ is the leader's mixed strategy, and $I = \langle i_\lambda \rangle \in \mathcal{I}_1 \times \cdots \times \mathcal{I}_\Lambda$ represents the indices of the followers' pure strategies, the expected payoffs for follower λ and the leader are defined respectively as follows:

$$U_\lambda(\boldsymbol{x}, I) = \sum_{j=1}^{|\mathcal{S}|} x_j U_\lambda^{ji_\lambda}, \tag{8}$$

$$U_l^\Lambda(\boldsymbol{x}, I) = \sum_{\lambda \in \Lambda} p_\lambda \sum_{j=1}^{|\mathcal{S}|} x_j U_l^{ji_\lambda}, \tag{9}$$

where $U_\lambda^{ji_\lambda}$ and $U_l^{ji_\lambda}$ are respectively the payoffs for follower λ and the leader when they choose pure strategies $s_j \in \mathcal{S}$ and $t_{i_\lambda} \in \mathcal{T}_\lambda$. Given an SSE strategy profile $\langle \boldsymbol{x}^*, I^* \rangle$, where $I^* = \langle i^*_\lambda \rangle$, $\langle \boldsymbol{x}', I^* \rangle$ also forms SSE if $\boldsymbol{x}'$ satisfy the following equation which is a variant of Eq. (10):

$$\sum_{j=1}^{|\mathcal{S}|} x'_j \boldsymbol{u}^\diamond_{\lambda j} = \sum_{j=1}^{|\mathcal{S}|} x^*_j \boldsymbol{u}^\diamond_{\lambda j}, \forall\, \lambda \in \Lambda, \tag{10}$$

where $\boldsymbol{u}^\diamond_{\lambda j} = (U_\lambda^{j1}, \ldots, U_\lambda^{j|\mathcal{T}_\lambda|}, U_l^{ji^*_\lambda})^\mathsf{T}$ is specified for each type of followers. Obviously, each follower makes identical responses under $\boldsymbol{x}^*$ and $\boldsymbol{x}'$ according to Eq. (8), and the leader receives the same expected payoff according to Eq. (9). Combining $\boldsymbol{u}^\diamond_{\lambda j}$ for all $\lambda \in \Lambda$ makes a vector of dimension $\sum_{\lambda \in \Lambda}(|\mathcal{T}_\lambda| + 1)$. By applying the Carathéodory's theorem, $\Phi = \sum_{\lambda \in \Lambda}(|\mathcal{T}_\lambda| + 1) + 1$ is then obtained for a Bayesian Stackelberg game (Corollary 4).

Corollary 4. *The minimum support size of the leader's SSE strategies in a Bayesian Stackelberg game is $\sum_{\lambda \in \Lambda}(|\mathcal{T}_\lambda| + 1) + 1$, where $\mathcal{T}_\lambda$ is the pure strategy set of follower type λ.*

Φ for Other Stackelberg Security Games

In this section, we review other types of Stackelberg security game models discussed in the literature and apply Corollary 3 to these security game models.

Multiple Protection Types for a Target

In some real-world scenarios, it makes a difference when different number or types of security resources are assigned to a target. For example, an attacker would be more likely captured at a target protected by ten security guards than one that is protected by only one guard. The defender is able to execute a variety of security activities on each target. Each activity π requires m_π security resources and provides different payoffs $U_d^\pi(t_i)$ and $U_a^\pi(t_i)$ when it is executed on target i and the attacker attacks this target. For example, as shown in Table 1, the defender can execute three activities $\{\pi_0, \pi_1, \pi_2\}$, where π_k assigns k security guards on a target. Assigning more security guards on a target provides a higher/lower payoff for the defender/attacker. The defender's pure strategy is thus an assignment of all m available security resources to execute activities on the n targets, which is more complex than a 0/1 vector. Furthermore, we can construct a counter example as follows, which implies that the game cannot be represented through a coverage vector. Suppose the attacker attacks target 1, and the defender executes π_0 and π_2 on target 1 with probability 0.5 for each, the expected payoffs for both players are 10 and 5, respectively; however, by executing π_1 on target 1 with probability 1.0, the same coverage rate can be attained, while the payoffs for the players, being 20 and -15 respectively, are different from the previous case.

	π_0	π_1	π_2
t_1	-10, 30	20, -15	30, -20
t_2	-15, 40	20, -10	30, -15

Table 1: Multiple protection types for each target (each entry shows the defender's/attacker's payoff).

Despite of the more complex defender strategy, the number of attacker's pure strategies, being equal to the number n of targets, remains unchanged as compared with the compact game model. Applying Corollary 3, we obtain $\Phi = n + 2$.

A Network-based Security Game

In a network-based game model (Washburn and Wood 1995; Tsai et al. 2010; Jain et al. 2011; Jain, Conitzer, and Tambe 2013), a defender takes action on a graph $G = (\mathcal{V}, \mathcal{E})$ with a set $\mathcal{V}$ of nodes and a set $\mathcal{E}$ of edges. The attacker is able to start at a node $s \in S \subseteq \mathcal{V}$ and travels through a path P in an attempt to reach one of the targets $t \in T \subseteq \mathcal{V}$. The defender places security resources on the edges, instead of directly on the targets, to capture the attacker. If an edge is covered by a resource, and the attacker travels through this edge, the attacker is captured; otherwise, the attacker travels through this edge successfully. Therefore, the defender's pure strategy is an allocation of resources on the edges denoted by the set L of covered edges; the attacker's pure strategy is an s-t path P from a start node to a target node. Denote the set of feasible defender and attacker pure strategies as $\mathcal{L}$ and $\mathcal{P}$, respectively.

(a) Coverage over targets (b) Coverage over edges

Figure 2: Counter examples where coverage vectors do not make sense.

The network-based security game cannot be represented through a coverage vector over the targets. A counter example is shown in Figure 2(a). If the defender has one resource and plays a mixed strategy with 0.5 probability for each of the pure strategies $\{0, 1\}$ (i.e., placing one resource on e_1 and no resource on e_2) and $\{1, 0\}$, then a coverage of 0 is induced since the target is not covered by neither of the pure strategies, but the attacker has 0.50 chance of being caught no matter which path he chooses. Furthermore, a coverage vector over the edges does not make sense, either. This is shown by a counter example in Figure 2(b). If the defender has two resources and plays $\{1, 0, 1, 0\}$ and $\{0, 1, 0, 1\}$ with 0.5 probability for each, a coverage vector of $(0.5, 0.5, 0.5, 0.5)^\intercal$ over the edges is induced, which is the same with playing $\{1, 1, 0, 0\}$ and $\{0, 0, 1, 1\}$ with 0.50 probability for each. However, the attacker will always be caught under the former defender strategy while he has 0.50 chance of not being caught under the latter one no matter which path he chooses.

By applying Corollary 3, $\Phi = |\mathcal{P}| + 2$ is obtained for a network-based security game. Specifically, in a network-based security game, $|\mathcal{P}|$ can be very large, such that $|\mathcal{P}| + 2 > |\mathcal{L}|$, which corresponds to the case where the vertex number (i.e., $|\mathcal{L}|$) of a convex hull is even smaller than the dimension (i.e., $|\mathcal{P}| + 2$) of the space. Since obviously the support size cannot be larger than the number of pure strategies, a more accurate Φ is $\min\{|\mathcal{P}| + 2, |\mathcal{L}|\}$. Note that this special case is not discussed when we obtain the previous Φs because the pure strategy set of the defender is more likely to be much larger than that of the attacker.

Conclusion

In this paper, we analyze the structure of the defender's strategy in a number of widely studied security games and provides Φ for these games. $\Phi = |\mathcal{T}| + 2$ is obtained for a general form two-player Stackberg game, and more generally $\Phi = \sum_{\lambda \in \Lambda}(|\mathcal{T}_\lambda| + 1) + 1$ for a Bayesian Stackelberg game. Apparently, Φ depends only on the follower's pure strategy space, so that the support of the defender's SSE strategy can be bounded very tightly when the pure strategy space of the attacker is drastically smaller than that of the defender. This happens in the compact game model, and the model with multiple protection types for each target. When

the attacker's pure strategy space is even larger than the defender's, which may happen in a network-based game, Φ degenerates to the total number of defender's pure strategies.

References

An, B.; Jain, M.; Tambe, M.; and Kiekintveld, C. 2011a. Mixed-initiative optimization in security games: A preliminary report. In *AAAI Spring Symposium: Help Me Help You: Bridging the Gaps in Human-Agent Collaboration*.

An, B.; Pita, J.; Shieh, E.; Tambe, M.; Kiekintveld, C.; and Marecki, J. 2011b. GUARDS and PROTECT: Next generation applications of security games. *ACM SIGecom Exchanges* 10(1):31–34.

An, B.; Tambe, M.; Ordonez, F.; Shieh, E. A.; and Kiekintveld, C. 2011c. Refinement of strong Stackelberg equilibria in security games. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence (AAAI'11)*, 587–593.

An, B.; Kempe, D.; Kiekintveld, C.; Shieh, E.; Singh, S.; Tambe, M.; and Vorobeychik, Y. 2012. Security games with limited surveillance. *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI'12)* 1241–1248.

An, B.; Brown, M.; Vorobeychik, Y.; and Tambe, M. 2013. Security games with surveillance cost and optimal timing of attack execution. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'13)*, 223–230.

Basar, T.; Olsder, G. J.; Clsder, G.; Basar, T.; Baser, T.; and Olsder, G. J. 1995. *Dynamic noncooperative game theory*, volume 200. SIAM.

Breton, M.; Alj, A.; and Haurie, A. 1988. Sequential Stackelberg equilibria in two-person games. *Journal of Optimization Theory and Applications* 59(1):71–97.

Brown, M.; An, B.; Kiekintveld, C.; Ordóñez, F.; and Tambe, M. 2012. Multi-objective optimization for security games. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'12)*, 863–870.

Danninger-Uchida, G. 2001. Carathéodory theorem. In Floudas, C., and Pardalos, P., eds., *Encyclopedia of Optimization*. Springer US. 236–237.

Harsanyi, J. C., and Selten, R. 1972. A generalized Nash solution for two-person bargaining games with incomplete information. *Management Science* 18(5-Part-2):80–106.

Jain, M.; Kardes, E.; Kiekintveld, C.; Ordónez, F.; and Tambe, M. 2010. Security games with arbitrary schedules: A branch and price approach. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI'10)*, 792–797.

Jain, M.; Korzhyk, D.; Vaněk, O.; Conitzer, V.; Pěchouček, M.; and Tambe, M. 2011. A double oracle algorithm for zero-sum security games on graphs. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'11)*, 327–334.

Jain, M.; Conitzer, V.; and Tambe, M. 2013. Security scheduling for real-world networks. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'13)*, 215–222.

Jain, M.; Kiekintveld, C.; and Tambe, M. 2011. Quality-bounded solutions for finite Bayesian Stackelberg games: scaling up. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'11)*, 997–1004.

Kiekintveld, C.; Jain, M.; Tsai, J.; Pita, J.; Ordóñez, F.; and Tambe, M. 2009. Computing optimal randomized resource allocations for massive security games. In *Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'09)*, 689–696.

Korzhyk, D.; Conitzer, V.; and Parr, R. 2010. Complexity of computing optimal Stackelberg strategies in security resource allocation games. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI'10)*, 805–810.

Leitmann, G. 1978. On generalized stackelberg strategies. *Journal of Optimization Theory and Applications* 26(4):637–643.

Paruchuri, P.; Pearce, J. P.; Marecki, J.; Tambe, M.; Ordonez, F.; and Kraus, S. 2008. Playing games for security: an efficient exact algorithm for solving Bayesian Stackelberg games. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'12)*, 895–902.

Pita, J.; Jain, M.; Marecki, J.; Ordóñez, F.; Portway, C.; Tambe, M.; Western, C.; Paruchuri, P.; and Kraus, S. 2008. Deployed ARMOR protection: the application of a game theoretic model for security at the Los Angeles International Airport. In *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'08)*, 125–132.

Pita, J.; Tambe, M.; Kiekintveld, C.; Cullen, S.; and Steigerwald, E. 2011. GUARDS: game theoretic security allocation on a national scale. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'11)*, 37–44.

Shieh, E. A.; An, B.; Yang, R.; Tambe, M.; Baldwin, C.; DiRenzo, J.; Maule, B.; and Meyer, G. 2012. PROTECT: An application of computational game theory for the security of the ports of the United States. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'12)*, 2173–2179.

Tsai, J.; Kiekintveld, C.; Ordonez, F.; Tambe, M.; and Rathi, S. 2009. Iris-a tool for strategic security allocation in transportation networks. In *Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'09)*, 82–90.

Tsai, J.; Yin, Z.; Kwak, J.-y.; Kempe, D.; Kiekintveld, C.; and Tambe, M. 2010. Urban security: game-theoretic resource allocation in networked physical domains. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAAI'10)*, 881–886.

Von Stengel, B., and Zamir, S. 2004. Leadership with commitment to mixed strategies.

Washburn, A., and Wood, K. 1995. Two-person zero-

sum games for network interdiction. *Operations Research* 43(2):243–251.

Yang, R.; Ordonez, F.; and Tambe, M. 2012. Computing optimal strategy against quantal response in security games. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'12)*, 847–854.

Yin, Z., and Tambe, M. 2012. A unified method for handling discrete and continuous uncertainty in bayesian stackelberg games. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'12)*, 855–862.