Stackelberg Security Games: Looking Beyond a Decade of Success

Arunesh Sinha[1], **Fei Fang**[2], **Bo An**[3], **Christopher Kiekintveld**[4], **Milind Tambe**[5]
[1] University of Michigan, [2] Carnegie Mellon University, [3] Nanyang Technological University,
[4] University of Texas at El Paso, [5] University of Southern California
arunesh@umich.edu, feifang@cmu.edu, boan@ntu.edu.sg,
cdkiekintveld@utep.edu, tambe@usc.edu

Abstract

The Stackelberg Security Game (SSG) model has been immensely influential in security research since it was introduced roughly a decade ago. Furthermore, deployed SSG-based applications are one of most successful examples of game theory applications in the real world. We present a broad survey of recent technical advances in SSG and related literature, and then look to the future by highlighting the new potential applications and open research problems in SSG.

1 Introduction

Game theory has long held the promise of improving intelligent decision making for complex security problems. This promise has been partially realized with application of Stackelberg Security Games (SSG) to a variety of security problems since 2006. Starting with the application of allocating security resources to eight terminals of the Los Angeles International Airport, recent advances in SSGs have enabled solving problems of immense complexity such as protecting biodiversity in conservation areas that span over 2500 square kilometers [Fang and Nguyen, 2016] and screening 800 million airport passengers annually throughout USA [Brown *et al.*, 2016]. In this survey, we walk through the various innovations in the application of SSGs with an eye towards potential applications and research challenges in the near future.

Security is a critical concern around the world, manifesting in such examples as infrastructure and human life protection, wildlife protection, and protection against cyber attacks. In all of these domains, the defender has limited security resources that preclude full security coverage of important potential targets at all times. Thus, allocations of limited security resources must be intelligent, taking into account differences in priorities of targets requiring security coverage, the response of adversaries to the security posture based on knowledge gained from surveillance, and potential uncertainty over the types, capabilities, knowledge, and priorities of adversaries faced. Casting these problems as a Stackelberg game is the basis of a growing literature on models and solution methods for solving these games to yield randomized security policies.

The SSG model has flourished, with dozens of papers on SSGs every year in major AI conferences[1]. The applications have also spread across countries and application domains. As reported in conferences such IJCAI'17 and 2017 Conference on Decision and Game Theory for Security (GameSec'17), the SSGs framework is now being tested in Israel for traffic monitoring, and in Chile and Argentina for drug interdiction. SSGs have also found substantial applications in cybersecurity, auditing for privacy, drug design against viruses, traffic enforcement, software code testing, adversarial learning, and many others. We present a thematic view of Stackelberg game-based approaches to security. Since deployed applications of SSG are quite well-known, we focus our attention more on techniques and potential new applications. We also briefly mention other models and techniques for security that do not fit the SSG framework. Finally, we present some open problems in SSGs.

Related surveys: There have been a number of surveys on SSGs such as [Fang and Nguyen, 2016; Nguyen *et al.*, 2016], which are often not widely available (e.g., book chapters) and targeted at a specialized audience. We give a general survey of SSGs for an AI audience that focuses on recent developments, distills the underlying concepts across the varied applications of SSGs and provides a centralized presentation of open problems in this area. We hope that this introduction to the state-of-the-art will aid new researchers in this area.

2 The Basic SSG Model

In SSG, a defender must defend a set of targets T using a limited number of resources, whereas the attacker is able to observe and learn the defender's strategy and attack after careful planning. An action, or *pure strategy*, for the defender, represents deploying a set of resources R for defending the targets. The set of all defender pure strategies are given by a set of allocation constraints, referred to as *scheduling constraints* in literature. The pure strategy for an attacker is an attack at a target. The *mixed strategy* of the defender is a probability distribution over the pure strategies. Additionally, with each target, there is an associated set of payoff values that define the utilities for both the defender and the attacker

[1] Given the space constraint on references, we cannot cite all relevant papers even in the last few years, so we aim to cover a representative sample that outlines the scope of the field.

in case of a successful or a failed attack. A key assumption of many SSG models is that the payoff of an outcome depends only on the target attacked, and whether or not it is "covered" (protected) by the defender (with a notable exception in [Gan *et al.*, 2015]). This allows us to compactly represent the payoffs of a security game. Formally, if target t is attacked, the defender's utility is $U_d^c(t)$ if t is covered, or $U_d^u(t)$ if t is not covered. The attacker's utility is $U_a^c(t)$ if t is covered, or $U_a^u(t)$ if t is not covered. In SSGs, it is always better for the defender to cover the attacked target as compared to leaving it uncovered, whereas it is always better for the attacker to attack an uncovered target. This assumption is consistent with the payoff trends in the real world. A special case is *zero-sum games*, in which for each outcome the sum of utilities for the defender and attacker is zero, although general SSGs are not necessarily zero-sum.

Strong Stackelberg Equilibrium: The solution to an SSG is a *mixed* strategy for the defender that maximizes the expected utility of the defender, given that the attacker learns the mixed strategy of the defender and chooses a best response. The most commonly adopted version of this solution concept in related literature is called Strong Stackelberg Equilibrium (SSE) [Conitzer and Sandholm, 2006; von Stengel and Zamir, 2004]. In SSGs, the mixed strategy of the defender induces probabilities of covering every target t denoted by a vector $C = \{c_t\}$. Furthermore, it is enough to consider a pure strategy of the rational adversary [Conitzer and Sandholm, 2006], which is to attack a target t. The expected utility for defender for a strategy profile (C, t) is $U_d(t, C) = c_t U_d^c(t) + (1 - c_t)U_d^u(t)$, and a similar form for the adversary. An SSE for the basic SSG (non-Bayesian, rational adversary) is defined as follows:

Definition 1. *A pair of strategies (C^*, t^*) form a* Strong Stackelberg Equilibrium *(SSE) if they satisfy the following:*

1. *The attacker plays a best-response against C^*: $U_a(t^*, C^*) \geq U_a(t, C^*)$ for all targets t.*
2. *The attacker breaks ties in favor of the defender: $U_d(t^*, C^*) \geq U_d(t', C^*)$ for all targets t' such that $t' = \text{argmax}_t U_a(t, C^*)$*
3. *C^* is optimal for the defender, that is, $U_d(t^*, C^*) \geq U_d(t(C), C)$ for all defender's strategy C where $t(C) = \text{argmax}_t U_a(t, C)$ is the attacker's best response against C when breaking ties in favor of the defender.*

The assumption that the follower (attacker) will always break ties in favor of the leader (defender) in cases of indifference is reasonable because in most cases the leader can induce the favorable strong equilibrium by selecting a strategy arbitrarily close to the equilibrium that causes the follower to strictly prefer the desired strategy [von Stengel and Zamir, 2004]. Furthermore, an SSE exists in all Stackelberg games, which makes it an attractive solution concept compared to versions of Stackelberg equilibrium with other tie-breaking rules.

Many variations of this basic SSG model and solution concepts other than SSE have been studied to handle different types of adversaries. For example, one variation is when the adversary does not best respond, but samples the target to attack from a probability distribution over target $h(C)$ that is conditional on C. h is the response function of the bounded rational adversary and various forms of h have been studied in literature [Sinha *et al.*, 2016]. Sequential versions with repeated interactions have also been studied.

Target	Defender		Adversary	
	U_d^c	U_d^u	U_a^c	U_a^u
1	0	-1	-1	1
2	-1	-2	-4	2
3	0	-3	-4	4
4	-2	-10	-5	10

Table 1: Example SSG

Illustrative examples: Consider the defender and adversary utilities in Table 1 for four targets. The defender has only two security resources r_1, r_2, each of which can protect one target. The set of feasible pure strategies for the defender can be represented by a set of vectors $\mathcal{P} = \{P\}$ such that $\sum_{t=1}^{4} P_t = 2, P_t \in \{0, 1\}$, where $P_t = 1$ means a security resource is allocated to protect target t. The mixed strategy is completely specified as all vectors $\{c_t\}$ such that $\sum_{t=1}^{4} c_t = 2, c_t \geq 0, c_t \in [0, 1]$, which can be seen to be a relaxation of the constraints used to specify the pure strategies. The SSE is given by $\{0.5, 0.33, 0.5, 0.66\}$.

However, the mixed strategy cannot always be represented succinctly as above. Suppose the resources can defend subsets of two targets simultaneously, but r_1 can defend $(1, 2), (3, 4)$ only and r_2 can defend $(1, 4), (2, 3)$ only; such constraints are called scheduling constraints in literature. Let $P_{r,t} \in \{0, 1\}$ be an indicator variable for when resource r protects target t. The set of feasible pure strategies is given by $\mathcal{P} = \{P\}$ such that $P_t = \min\{1, P_{r_1,t} + P_{r_2,t}\}, P_t \in \{0, 1\}$ for all t. Now, the space of feasible mixed strategies cannot be represented by a simple relaxation of the pure strategy constraints. Rather, the mixed strategy space is given by all vectors $\sum_{P \in \mathcal{P}} p_P P$, where $\mathcal{P}$ is the set of all pure strategies and p_P denotes the probability of choosing pure strategy P.

Furthermore, in many SSGs, the adversary is bounded rational [Nguyen *et al.*, 2013]. In these cases, instead of a single target chosen as part of the best response by the adversary to the induced coverage $C = \{c_t\}$, the adversary's response $h(C)$ is to choose a target stochastically according to a probability q_t of choosing target t. One such probability function form is $q_t \propto e^{w_t c_t + a_t}$, where w_t and a_t are learnable constants. Intuitively, a_t denotes attractiveness of target t, that is, the adversary is more likely to consider attacking attractive targets. Also, w_t is always negative, that is, the adversary will likely not attack targets that have a higher chance of being defended. a_t can be a function of $U_a^c(t)$ and $U_a^u(t)$.

3 Deployed Applications

We briefly cover deployed applications of SSG to emphasize the practical importance of this domain. All references for these can be found in prior surveys [Nguyen *et al.*, 2016]. Figure 1 provides an illustration of some applications. SSGs started with applications in *infrastructure security* including

(a) Barrier free train station requires randomized inspection; SSGs have been used for such inspection.

(b) SSG has been tested for intelligent screening of passengers (as part of the TSA DARMS program).

(c) SSG is being used to design effective human patrols and drone usage in wildlife protection.

Figure 1: Deployed Applications

ARMOR at the Los Angeles Airport (LAX) deployed in 2007 to randomize checkpoints on the roadways entering the airport; which was followed by IRIS, a game-theoretic scheduler for randomized deployment of the U.S. Federal Air Marshal Service (FAMS) that has been in use since 2009; and PROTECT which is deployed for generating randomized patrol schedules for the U.S. Coast Guard in Boston, New York, Los Angeles, and other ports around the United States. The threat screening game (TSG) is currently being evaluated for screening airport passengers in the USA as part of the TSA DARMS program for an overhaul of screening procedures [Brown *et al.*, 2016].

Green security games [Fang and Nguyen, 2016; Fang *et al.*, 2015; Fang *et al.*, 2016] focus on defending against environmental crimes. These problems exhibit a spatial and temporal aspect that distinguishes work on these problems from infrastructure security. Moreover, the adversaries in such games cannot be expected to be completely strategic so behavioral adversary models play an important role. In particular, PAWS is a wildlife protection assistant system that has been extensively evaluated in Malaysia and the Queen Elizabeth National Park in Uganda and incorporated in operations. MIDAS is another application that was tested by the US Coast Guard for protecting fisheries against over-fishing.

Opportunistic crime [Zhang *et al.*, 2015] refers to the problem of urban crime where criminals are not committed to detailed plans and are flexible in the execution of their plans, as opportunities arise. Protecting against such urban crime has been studied as a Stackelberg game and evaluated for deterring fare evasion within the Los Angeles Metro System (TRUSTS) and for crime prevention at the University of Southern California.

While many deployed applications have originated from research work (and spin-offs) conducted at the University of Southern California, there are a growing number of independent applications including traffic monitoring in Israel and drug interdiction in Chile and Argentina (see proceedings of GameSec'17). These demonstrate the generality of the SSG framework in tackling real-world security problems.

4 Key Technical Innovations

We survey the technical advances in SSGs and variants in two parts, focusing on technical approaches in deployed applications and then presenting various modifications to the basic SSG model proposed in the literature. In order to easily find the references to these approaches and models, we cite key relevant papers for the approaches in Table 2 and the modified models in Table 3.

4.1 Deployed Applications and Extensions

The increasing intricacy of SSG models for new applications has also resulted in a number of algorithmic advances for SSE. We identify three main dimensions that affect the SSE computation: (1) large models, (2) bounded rationality models, and (3) uncertainty. First, the growing size of applications has presented challenges in computing the SSE. For example, the defender strategy space in recent applications like TSG is larger than 10^{33}. In the road networks security problem, the adversary space is of the order 10^{18} [Nguyen *et al.*, 2016].

Next, in many domains such as green security, it is unrealistic to assume that the adversary is fully rational. This motivates work in SSGs that focuses on bounded rationality models for the adversary. These have used behavior models such as quantal response models, lens-QR models and prospect theory inspired models [Kar *et al.*, 2015]. These parametrized models not only inherently present a problem of learning the parameters, they also greatly exacerbate the SSE computation problem due to the non-linearity introduced by the behavior models.

Finally, the parameters of any game model are never known with certainty; incorporating such uncertainty in the SSE computation can mitigate the effect of worst-case realizations of parameters. In the literature, models of uncertainty include both Bayesian and interval models of uncertainty for parameter values. There are various other kinds of possible uncertainty such as action outcome uncertainty (modeled as MDPs) as well as uncertainty about underlying state of the world (partial observability) [Nguyen *et al.*, 2014]. Uncertainty considerations can greatly (though not always) increase

Approach	Key Paper(s)	General-sum	Bounded rationality	Multi-step
1a	[Kiekintveld *et al.*, 2009]	Yes	No	No
1b	[Jain *et al.*, 2010] [Jain *et al.*, 2011]	Yes	No	No
1c	[Brown *et al.*, 2016]	No	No	No
1d	[Basak *et al.*, 2016]	Yes	No	No
2a	[Nguyen *et al.*, 2013] [Kar *et al.*, 2015]	N/A	Yes	Yes
2b	[Yang *et al.*, 2013]	N/A	Yes	Yes
3a	[Nguyen *et al.*, 2014]	Yes	No	No
3b	[Blum *et al.*, 2014]	Yes	No	Yes
3c	[Kar *et al.*, 2017] [Zhang *et al.*, 2015]	N/A	Yes	Yes

Table 2: Summary of key papers for technical approaches in SSGs. The papers cited are the first paper that introduced a new approach in a series of papers for the corresponding approach. N/A means not applicable.

the computational burden in computing the SSE, and motivate learning game parameters as a research topic in itself.

In response to these challenges, various innovative approaches have been proposed in the literature. These approaches can be broadly classified according to the three challenges they address (the three dimensions stated above). We present these approaches in the following list with approaches 1a,1b,1c,1d addressing challenge 1 (large models), approaches 2a,2b addressing challenge 2 (bounded rationality), and approaches 3a,3b,3c addressing challenge 3 (uncertainty).

1.a) **Marginal strategy representation.** An approach based on a relaxed representation of the SSE optimization problem using unconstrained (no scheduling constraints) coverage (marginal) probabilities. The resulting marginal solution may not be a feasible mixed strategy in the worst case, but can sometimes be repaired. The relaxed problem is solved based on a bucket-filling approach revealing and exploiting features of SSGs. This approach is much faster than solving a linear optimization and enables solving very large SSGs directly, and is also useful as a heuristic in more complex algorithms

1.b) **Incremental strategy generation.** An approach based on solving the SSE problem with restricted action spaces for the defender or the adversary, and then iteratively expanding the action space with relevant actions. This includes methods based on column generation or the double oracle approach or branch and price techniques.

1.c) **Polytope decomposition approach.** An approach building on the marginal approach by identifying feasible sub-polytopes of the defender mixed strategy polytope around the marginal solution and taking their convex hull to yield marginal solutions that are feasible mixed strategies. This new approach to linear optimization is computationally efficient in some scenarios and applicable to problems beyond SSGs.

1.d) **Abstraction.** An approach often used in AI is abstraction, that is, formulating a smaller problem from a given large problem, solving the smaller problem and mapping the solution back to the large problem. An approach to abstraction is to combine similar actions, and while it may lead to a loss in solution quality, it can significantly improve the computational efficiency. Abstraction has been extensively used in Poker games research and has been applied to improve scalability in SSGs as well.

2.a) **Modeling adversary bounded rational behavior.** Inspired by psychology and behavioral economics models, parametrized models of bounded rationality such as *quantal response*, *subjective utility quantal response* and prospect theoretical models have been studied for SSG. These models provide general techniques for modeling bounded rationality in games, which can be broadly applied to other types of game interaction. Unfortunately, the (typical) non-convexity of these model greatly increases the computational burden.

2.b) **Incremental strategy generation with non-linear adversary behavior.** An approach that is an efficient realization of the well-known optimization technique of cutting-planes with non-linear bounded rationality models in SSGs. The approach is a general technique for solving non-linear optimization using cutting plane methods that could be used for general nonlinear optimization problems.

3.a) **Robust optimization.** These techniques have addressed uncertainty in game parameters using concepts of maximin robustness and minimax regret robustness among others forms of optimization. The approaches here also provide general models of adversary behavior called *monotone maximin* that generalize those in 2(a) above, and techniques of solving games with such models.

3.b) **Learning defender strategies against an unknown rational adversary.** An approach that assumes the rational adversary's utility is unknown and learns the optimal defender strategy by repeatedly playing intelligently chosen strategies against the adversary. This approach guarantees to learn the optimal defender strategy within a polynomial number of defender strategies.

3.c) **Learning bounded rational adversary model.** While early work in learning bounded rational adversary models used basic maximum likelihood estimation, recent approaches have integrated defender strategy in the learning itself to propose various techniques based on graphical models or decision trees.

Model	Key Paper(s)	General-sum	Bounded rationality	Multi-step
4a	[Letchford and Vorobeychik, 2013] [Panda and Vorobeychik, 2017]	Yes	No	Yes
4b	[Guo *et al.*, 2016][Gholami *et al.*, 2016]	Yes	No and Yes	No
4c	[Basilico *et al.*, 2012][Basilico *et al.*, 2016]	Yes	No	Yes
4d	[Blocki *et al.*, 2013][Blocki *et al.*, 2015]	Yes	No	No

Table 3: Summary of key papers for SSG inspired models.

Besides the above techniques, there are broad theoretical characterizations of classes of SSGs: e.g., the complexity of SSE computation in the rational adversary setting [Xu, 2016], a theoretical characterization of the equilibrium [Korzhyk *et al.*, 2011] in SSGs, and sample complexity of learning [Sinha *et al.*, 2016].

4.2 Novel Potential Applications

There is an exciting set of emerging SSG inspired models for entirely new kinds of problems in the literature with corresponding techniques of solving them (citations in Table 3).

4.a) **Plan interdiction games.** These games present a Stackelberg game model of security combined with planning representations in which the defender chooses a mitigation strategy that interdicts potential attack actions, and the attacker responds by computing an optimal attack plan that circumvents the deployed mitigations. The work has focused on interdiction of very elaborate Markov decision process based plans of the adversary, which applies to adversaries in cybersecurity.

4.b) **Coalitional security games.** This deals with scenarios where attackers can form coalitions and addresses the problem of optimally inhibiting the formation of attacker coalitions. The motivating domain is breaking up terrorist links or cells or breaking up collusion between multiple attackers. This topic marries aspects of coalitional game theory with SSGs.

4.c) **Patrolling games.** An extensive-form infinite-horizon game with commitment by the defender, where decision nodes are potentially infinite. The attacker can undertake actions during the execution of the defender's strategy. The most prominent application is patrolling environments against intruders, which is inspired by the well-known game model of pursuit-evasion, but has been extended in various ways including the introduction of alarm systems.

4.d) **Audit games.** These games study economic considerations in the design of audit mechanisms, focusing on effective resource allocation and appropriate punishment schemes. The audit game model is a generalization of a security game model with an additional punishment parameter. The models are applied to audits to ensure privacy policy compliance in US hospitals.

Unlike the basic SSG model, some of the models above, such as 4b and 4c, have multiple attackers. In literature, various theoretical extensions of the SSG model such as multiple attackers [Korzhyk *et al.*, 2017], multiple defenders [Lou and Vorobeychik, 2015], and Bayesian generalization of SSGs [Li *et al.*, 2016] have also been considered. Influenced by research on SSGs, Stackelberg games have been used broadly in literature for other applications such as drug design against viruses, software code testing, and adversarial learning.

4.3 Open Problems

We present some immediate open problems in SSGs. Scalability remains an issue despite many existing approaches, especially in handling uncertainty. As a result, most deployed applications do not handle uncertainty so there is a pressing need to address scalability when dealing with uncertainty. Moreover, scalability has resulted in a number of approaches that perform well in practice but are not known to be polynomial time such as the polytope decomposition method. In problems using bounded rationality models, the question of how to further improve the prediction of adversary behavior using recent advances in learning is an interesting research direction. Two other across-the-board issues are discussed below.

Deception: A fundamental issue in any security situation is deception. The mythical Trojan Horse is a classic example of deception. Classes of computer malware are called trojan horses, symbolizing the deceptive behavior inherent in malware. Deception by the defender has been studied in SSG literature, albeit in simple one-time interaction settings using signaling enabled by the advantage of extra information available to the defender [Xu *et al.*, 2015]. More broadly, deception in game theory arises from asymmetric information. Some proposals in SSG based approaches to cybersecurity have looked at expending resources to provide additional information advantage to the defender. This is in the form of honeypots [Gutierrez and Kiekintveld, 2017], which are dummy systems meant to lure the adversary to attack them. Honeypots, in addition to minimizing loss from the current attack, fool adversaries to reveal their secrets.

Deception can be complicated, especially when used by both the defender and adversary [Guo *et al.*, 2017]. Research on deception by signaling or by designing the game to provide an information advantage in a sequential game could enable applications of SSGs for highly complex defender-adversary interaction.

Active and secure defense: The SSG models have mostly specified the defender's actions as defending targets, with some exceptions (e.g., plan interdiction and coalition games). More broadly, defense can be a combination of defensive actions, offensive actions, and information seeking actions. In a highly tactical environment, a defender must use all options

and decide which actions to take or not. It is important to note that actions taken by defender may not always improve the security situation, as an observant adversary may be able to figure out weaknesses in defense such as the lack of information that a defender is seeking to obtain. Thus, defender actions may need to be covert and any defense strategy must be itself secure from adversarial attacks.

5 Other Approaches to Security

While SSGs have been quite successful in tackling security problems, a number of other game-theoretic approaches to security have been proposed. Some of these approaches focus on different modeling aspects of security, while others focus on technical generalizations of SSGs. A prominent model of security interactions is known as interdependent security games [Laszka *et al.*, 2014]. This simultaneous move model focuses on the interdependence between multiple defenders within the underlying system that is being protected. For technical generalization, there are theoretical papers on commitment in extensive form game [Letchford *et al.*, 2014], stochastic games [Letchford *et al.*, 2014] and games with imperfect information [Cermak *et al.*, 2016]. While theoretical in nature, these techniques could be useful in future as the applications become increasingly complex.

6 Future Applications

We discuss potential application domains where SSGs can be applied. Some initial applications in these domains are in a nascent stage but there is scope for much more research and new applications. These require innovations in models and algorithms that extend beyond those discussed previously.

Changing terrorism threats: The terrorist threat scenario has been evolving over the last decade. In response to efficient anti-terrorism measures, the nature of threats has moved from scenarios of well-planned attacks to lone wolf attacks. This presents a whole new modeling problem in itself where terrorists act like opportunistic criminals but are determined to carry out an attack. Models need to incorporate the action of actively seeking information about the potential lone wolf, which needs to conducted with limited resources.

Green security challenges: The current green security applications have addressed problems in wildlife and natural resources protection for specific scenarios. However, every nature park presents unique challenges requiring novel techniques to address them. The scale of these parks, their diversity, the different needs of human patrollers, and different types of crimes against the environment imply that a single model or software and a fixed level of autonomy will not work across the world. However, an entirely different model and approach for each park is also not desirable. Thus, the challenge is to develop a flexible and tunable model of green security problems.

Elaborating on a few points above, incorporating real time information in patrolling strategy is important; this presents an opportunity to use drones for passive patrolling of large areas [Bondi *et al.*, 2018] and also use of other static sensors. Of course, drones are costly (especially in resource constrained countries) and thus it is critical to use such resources optimally. Another aspect of wildlife crime is the varied types of crimes that include poaching, illegal logging, encroachment of park lands, and others. This presents a multi-objective problem that requires innovations in multi-objective game theory.

Cyber-security applications: Cybersecurity is a quintessential example of a complex security challenge. There are a variety of sub-problems in cyber-security that involve interaction between defenders, adversaries, and users. In contrast to physical security, a number of characteristics of cybersecurity make the problem more complex, including:

- Changing state of the world. The underlying state in cyber-security includes dynamically changing components such as operating systems, software applications, communication networks, etc. This makes it hard to technically specify the underling system even without considering defender-adversary interaction. This is generally known as a dynamic *attack surface*.
- Enormous problem size with unobserved actions. The possible actions over multiple time steps for both the defender and adversary in cybersecurity domain are enormous. Further, imperfect and incomplete information is the norm in cybersecurity. Attacks are often stealthy and go unnoticed for days or months. Scalability in presence of such complications is an even greater challenge in cyber-security than other domains.
- Unknown actions. A complete specification of all possible actions is often not feasible, resulting in what are known as *zero-day* attacks. Handling zero-day attack actions is a critical modeling problem in cyber-security.
- Multiple players. The attack surface in cyber-security includes human actor such as the users of the cyber-system. Attacks on the cyber-system often succeed by deceiving users to reveal critical credential and these are referred to as social engineering attacks in the literature. The presence of these user agents adds another player(s) to the game apart from the defender and adversary, making the problem significantly harder.

One subproblem that researchers have looked at in cyber-security using the SSG framework is the problem of allocating limited human resources to inspecting a large number of alerts from any intrusion prevention system [Schlenker *et al.*, 2017]. Looking forward, the problem of protection against social engineering attacks, balancing between performance and protection when using anti-virus software and ensuring compliance with security policies are topics amenable to SSG solutions. Other applications include fighting against spear phishing [Zhao *et al.*, 2016].

Privacy applications: While privacy is sometimes conflated with security, privacy presents distinct challenges from traditional security issues. A potential application that we discussed was SSGs for privacy audits. Broadly, privacy issues with any software system always involve a trade-off between privacy and usefulness of the system such as location privacy in social networks. This balance is meaningful only when there is a model of how the privacy-compromising adversary

will act. Game theory is an apt model for such interactions. However, privacy problems possess the same three issues as cybersecurity so there are significant research challenges in applying SSG for privacy problems.

Combating fraud in e-commerce: The main function of many e-commerce platforms is to guide buyer impressions to sellers, where a buyer impression means one buyer click on a product. Buyer impressions are usually allocated through a ranking system that displays the sellers' products according to the conversion rate, which is the probability that a buyer buys the product if he clicks on it. This is done to increase the total number of transactions. Since the sellers usually cannot control the conversion rate of their products, they usually spend much effort on getting more buyer impressions. A legal approach to obtain more buyer impressions is advertising, which is costly. Many sellers resort to illegal ways such as artificially raising the conversion rate through fake transactions, wherein sellers control a number of buyer accounts and use them to buy their own products. Such fraudulent behaviors severely decrease the effectiveness of impression allocation and jeopardize the business environment.

Currently, e-commerce platforms mainly rely on fraud detection techniques to combat fraudulent behaviors. However, fraudulent behaviors and techniques are always evolving, resulting in a low detection rate. A promising approach is to design an optimal mechanism for the platform (the leader) to deter fraudulent behaviors by the followers who learn the platform's policy through their interactions with the platform. The key research challenges include: 1) learning heterogeneous sellers' behavior models from trading data, 2) computing the optimal policy with millions of sellers and continuous strategy space, 3) designing policies robust to market evolution and uncertainty, and 4) balancing other optimization objectives of the platform while deterring fraudulent behaviors.

7 Discussion

The SSG research area has been application driven and presents large-scale interdisciplinary research challenges that call upon multi-agent researchers to work with researchers in other disciplines, be *on the ground* with the domain experts, and examine real-world constraints and challenges that cannot be abstracted away. Our goal with this survey is to point out some fundamental publications in this area that shed light on how these practical challenges were addressed. Clearly, there are many more publications beyond what we are able to cite in this survey. As such, there are a number of resources (mostly online) to explore SSGs further and also for getting starting with research in this area. The conferences that SSG papers have appeared in include AAAI, AAMAS, IJCAI and GameSec. A recent tutorial named "Advances in Game Theory for Security and Privacy" in ACM-EC 2017 (slides available online) provides more material for perusal.

References

[Basak *et al.*, 2016] Anjon Basak, Fei Fang, Thanh Nguyen, and Christopher Kiekintveld. Abstraction methods for solving graph-based security games. In *AAMAS*, 2016.

[Basilico *et al.*, 2012] Nicola Basilico, Nicola Gatti, and Francesco Amigoni. Patrolling security games: Definition and algorithms for solving large instances with single patroller and single intruder. *Artificial Intelligence*, 184:78–123, 2012.

[Basilico *et al.*, 2016] Nicola Basilico, Giuseppe De Nittis, and Nicola Gatti. A security game combining patrolling and alarm-triggered responses under spatial and detection uncertainties. In *AAAI*, pages 404–410, 2016.

[Blocki *et al.*, 2013] Jeremiah Blocki, Nicolas Christin, Anupam Datta, Ariel D. Procaccia, and Arunesh Sinha. Audit games. In *IJCAI*, 2013.

[Blocki *et al.*, 2015] Jeremiah Blocki, Nicolas Christin, Anupam Datta, Ariel D. Procaccia, and Arunesh Sinha. Audit games with multiple defender resources. In *AAAI*, 2015.

[Blum *et al.*, 2014] Avrim Blum, Nika Haghtalab, and Ariel D Procaccia. Learning optimal commitment to overcome insecurity. In *NIPS*, 2014.

[Bondi *et al.*, 2018] Elizabeth Bondi, Fei Fang, Mark Hamilton, Donnabell Dmello Debarun Kar, Robert Hannaford Jongmoo Choi, Arvind Iyer, Lucas Joppa, Milind Tambe, and Ram Nevatia. Spot poachers in action: Augmenting conservation drones with automatic detection in near real time. In *IAAI*, 2018.

[Brown *et al.*, 2016] Matthew Brown, Arunesh Sinha, Aaron Schlenker, and Milind Tambe. One size does not fit all: A game-theoretic approach for dynamically and effectively screening for threats. In *AAAI*, 2016.

[Cermak *et al.*, 2016] Jiri Cermak, Branislav Bosansky, Karel Durkota, Viliam Lisy, and Christopher Kiekintveld. Using correlated strategies for computing stackelberg equilibria in extensive-form games. In *AAAI*, 2016.

[Conitzer and Sandholm, 2006] Vincent Conitzer and Tuomas Sandholm. Computing the optimal strategy to commit to. In *ACM conference on Electronic commerce*, 2006.

[Fang and Nguyen, 2016] Fei Fang and Thanh Nguyen. Green security games: Apply game theory to addressing green security challenges. *ACM SIGecom Exchanges*, 15(1), 2016.

[Fang *et al.*, 2015] Fei Fang, Peter Stone, and Milind Tambe. When security games go green: Designing defender strategies to prevent poaching and illegal fishing. In *IJCAI*, 2015.

[Fang *et al.*, 2016] Fei Fang, Thanh Hong Nguyen, Rob Pickles, Wai Y Lam, Gopalasamy R Clements, Bo An, Amandeep Singh, Milind Tambe, Andrew Lemieux, et al. Deploying paws: Field optimization of the protection assistant for wildlife security. In *IAAI*, 2016.

[Gan *et al.*, 2015] Jiarui Gan, Bo An, and Yevgeniy Vorobeychik. Security games with protection externalities. In *AAAI*, 2015.

[Gholami *et al.*, 2016] Shahrzad Gholami, Bryan Wilder, Matthew Brown, Dana Thomas, Nicole Sintov, and Milind

Tambe. Divide to defend: Collusive security games. In *GameSec*, 2016.

[Guo *et al.*, 2016] Qingyu Guo, Bo An, Yevgeniy Vorobeychik, Long Tran-Thanh, Jiarui Gan, and Chunyan Miao. Coalitional security games. In *AAMAS*, 2016.

[Guo *et al.*, 2017] Qingyu Guo, Bo An, Branislav Bošanskỳ, and Christopher Kiekintveld. Comparing strategic secrecy and stackelberg commitment in security games. In *IJCAI*, 2017.

[Gutierrez and Kiekintveld, 2017] Marcus Gutierrez and Christopher Kiekintveld. Adapting with honeypot configurations to detect evolving exploits. In *AAMAS*, 2017.

[Jain *et al.*, 2010] Manish Jain, Erim Kardes, Christopher Kiekintveld, Fernando Ordonez, and Milind Tambe. Security Games with Arbitrary Schedules: A Branch and Price Approach. In *AAAI*, 2010.

[Jain *et al.*, 2011] Manish Jain, Dmytro Korzhyk, Ondřej Vaněk, Vincent Conitzer, Michal Pěchouček, and Milind Tambe. A double oracle algorithm for zero-sum security games on graphs. In *AAMAS*, 2011.

[Kar *et al.*, 2015] Debarun Kar, Fei Fang, Francesco Delle Fave, Nicole Sintov, and Milind Tambe. A game of thrones: when human behavior models compete in repeated stackelberg security games. In *AAMAS*, 2015.

[Kar *et al.*, 2017] Debarun Kar, Benjamin Ford, Shahrzad Gholami, Fei Fang, Andrew Plumptre, Milind Tambe, et al. Cloudy with a chance of poaching: adversary behavior modeling and forecasting with real-world poaching data. In *AAMAS*, 2017.

[Kiekintveld *et al.*, 2009] Christopher Kiekintveld, Manish Jain, Jason Tsai, James Pita, Fernando Ordóñez, and Milind Tambe. Computing optimal randomized resource allocations for massive security games. In *AAMAS*, 2009.

[Korzhyk *et al.*, 2011] Dmytro Korzhyk, Zhengyu Yin, Christopher Kiekintveld, Vincent Conitzer, and Milind Tambe. Stackelberg vs. nash in security games: An extended investigation of interchangeability, equivalence, and uniqueness. *Journal of Artificial Intelligence Research*, 41, 2011.

[Korzhyk *et al.*, 2017] Dmytro Korzhyk, Vincent Conitzer, and Ronald Parr. Security games with multiple attacker resources. In *IJCAI*, 2017.

[Laszka *et al.*, 2014] Aron Laszka, Mark Felegyhazi, and Levente Buttyán. A survey of interdependent security games. *ACM Comput. Surv.*, 2014.

[Letchford and Vorobeychik, 2013] Joshua Letchford and Yevgeniy Vorobeychik. Optimal interdiction of attack plans. In *AAMAS*, 2013.

[Letchford *et al.*, 2014] Joshua Letchford, Dmytro Korzhyk, and Vincent Conitzer. On the value of commitment. *JAAMAS*, 28(6), 2014.

[Li *et al.*, 2016] Yuqian Li, Vincent Conitzer, and Dmytro Korzhyk. Catcher-evader games. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, pages 329–337. AAAI Press, 2016.

[Lou and Vorobeychik, 2015] Jian Lou and Yevgeniy Vorobeychik. Equilibrium analysis of multi-defender security games. In *IJCAI*, 2015.

[Nguyen *et al.*, 2013] Thanh Nguyen, Rong Yang, Amos Azaria, Sarit Kraus, and Milind Tambe. Analyzing the effectiveness of adversary modeling in security games. In *AAAI*, 2013.

[Nguyen *et al.*, 2014] Thanh Nguyen, Albert Jiang, and Milind Tambe. Stop the compartmentalization: Unified robust algorithms for handling uncertainties in security games. In *AAMAS*, 2014.

[Nguyen *et al.*, 2016] Thanh Nguyen, Debarun Kar, Matthew Brown, Arunesh Sinha, Albert Jiang, and Milind Tambe. Towards a science of security games. In *Math. Sciences with Multidisciplinary Appl.* Springer, 2016.

[Panda and Vorobeychik, 2017] Swetasudha Panda and Yevgeniy Vorobeychik. Near-optimal interdiction of factored mdps. In *UAI*, 2017.

[Schlenker *et al.*, 2017] Aaron Schlenker, Haifeng Xu, Mina Guirguis, Christopher Kiekintveld, Arunesh Sinha, Milind Tambe, Solomon Sonya, Darryl Balderas, and Noah Dunstatter. Don't bury your head in warnings: A game-theoretic approach for intelligent allocation of cyber-security alerts. In *IJCAI*, 2017.

[Sinha *et al.*, 2016] Arunesh Sinha, Debarun Kar, and Milind Tambe. Learning adversary behavior in security games: A PAC model perspective. In *AAMAS*, 2016.

[von Stengel and Zamir, 2004] Bernhard von Stengel and Shmuel Zamir. Leadership with Commitment to Mixed Strategies. Technical Report LSE-CDAM-2004-01, CDAM Research Report, 2004.

[Xu *et al.*, 2015] Haifeng Xu, Zinovi Rabinovich, Shaddin Dughmi, and Milind Tambe. Exploring information asymmetry in two-stage security games. In *AAAI*, 2015.

[Xu, 2016] Haifeng Xu. The mysteries of security games: Equilibrium computation becomes combinatorial algorithm design. In *ACM Economics and Computation*, 2016.

[Yang *et al.*, 2013] Rong Yang, Albert Jiang, Milind Tambe, and Fernando Ordóñez. Scaling-up security games with boundedly rational adversaries: a cutting-plane approach. In *IJCAI*, pages 404–410. AAAI Press, 2013.

[Zhang *et al.*, 2015] Chao Zhang, Arunesh Sinha, and Milind Tambe. Keeping pace with criminals: Designing patrol allocation against adaptive opportunistic criminals. In *AAMAS*, 2015.

[Zhao *et al.*, 2016] Mengchen Zhao, Bo An, and Christopher Kiekintveld. Optimizing personalized email filtering thresholds to mitigate sequential spear phishing attacks. In *AAAI*, 2016.