

Towards a Science of Security Games

Thanh H. Nguyen¹, Debarun Kar¹, Matthew Brown¹, Arunesh Sinha¹, Milind Tambe¹, and Albert Xin Jiang²

¹ University of Southern California, Los Angeles, USA
[thanhnhng,dkar,mattheab,aruneshs,tambe]@usc.edu

² Trinity University, San Antonio, TX, USA
xjiang@trinity.edu

Abstract. Security is a critical concern around the world. In many domains from counter-terrorism to sustainability, limited security resources prevent complete security coverage at all times. Instead, these limited resources must be scheduled (or allocated or deployed), while simultaneously taking into account the importance of different targets, the responses of the adversaries to the security posture, and the potential uncertainties in adversary payoffs and observations, etc. Computational game theory can help generate such security schedules. Indeed, casting the problem as a Stackelberg game, we have developed new algorithms that are now deployed over multiple years in multiple applications for scheduling of security resources. These applications are leading to real-world use-inspired research in the emerging research area of “security games”. The research challenges posed by these applications include scaling up security games to real-world sized problems, handling multiple types of uncertainty, and dealing with bounded rationality of human adversaries.

1 Introduction

Security is a critical concern around the world that manifests in problems such as protecting our ports, airports, public transportation, and other critical national infrastructure from terrorists, in protecting our wildlife and forests from poachers and smugglers, and in curtailing the illegal flow of weapons, drugs, and money across international borders. In all of these problems, we have limited security resources which prevents security coverage on all the targets at all times; instead, security resources must be deployed intelligently taking into account differences in the importance of targets, the responses of the attackers to the security posture, and potential uncertainty over the types, capabilities, knowledge and priorities of attackers faced.

Game theory, which models interactions among multiple self-interested agents, is well-suited to the adversarial reasoning required for the security resource allocation and scheduling problem. Casting the problem as a Stackelberg game, we have developed new algorithms for efficiently solving such games that provide randomized patrolling or inspection strategies. These algorithms have led to successes and advances over previous human-designed approaches in security scheduling and allocation by addressing the key weakness of predictability in human-designed schedules. These algorithms are now deployed in multiple applications. The first application was ARMOR, which was

deployed at the Los Angeles International Airport (LAX) in 2007 to randomize checkpoints on the roadways entering the airport and canine patrol routes within the airport terminals [14]. Following that, came several other applications: IRIS, a game-theoretic scheduler for randomized deployment of the US Federal Air Marshals (FAMS), has been in use since 2009 [14]; PROTECT, which schedules the US Coast Guard’s randomized patrolling of ports, has been deployed in the port of Boston since April 2011 and is in use at the port of New York since February 2012 [35], and has spread to other ports such as Los Angeles/Long Beach, Houston, and others; another application for deploying escort boats to protect ferries has been deployed by the US Coast Guard since April 2013 [6]; and TRUSTS [43] which has been evaluated in field trials by the Los Angeles Sheriffs Department (LASD) in the LA Metro system. Most recently, PAWS— another game-theoretic application was tested by rangers in Uganda for protecting wildlife in Queen Elizabeth National Park in April 2014 [38]; MIDAS was tested by the US Coast Guard for protecting fisheries [10]. These initial successes point the way to major future applications in a wide range of security domains.

Given the many game-theoretic applications for solving real-world security problems, this book chapter provides an overview of the models and algorithms, key research challenges and a brief description of our successful deployments. Overall, the work in security games has produced numerous game-theoretic decision aids that are in daily use by security agencies to optimize their limited security resources. The implementation of these applications required addressing fundamental research challenges and has led to an emerging *science of security games* consisting of a general framework for modeling and solving security resource allocation problems. We categorize the research challenges associated with security games into four broad categories: (1) addressing scalability across a number of dimensions of the game, (2) tackling different forms of uncertainty that be present in the game, (3) addressing human adversaries’ bounded rationality and (4) evaluation of the framework in the field. Given the success in providing solutions for many security domains involving the protection of critical infrastructure, the science of security games has evolved and expanded to include new types of security domains for wildlife and environmental protection. These *green security games* address important global conservation problems and introduce additional research challenges that require incorporating new techniques such as planning and learning into security games.

The rest of the chapter is organized as follows: Section 2 introduces the general security games model, Section 3 describes the approaches used to tackle scalability issues, Section 4 describes the approaches to deal with uncertainty, Section 5 focuses on bounded rationality and Section 6 provides details of field evaluation of the science of security games.

2 Stackelberg Security Games

Stackelberg games were first introduced to model leadership and commitment [36]. The term Stackelberg Security Games (SSG) was first introduced by Kiekintveld et al [21] to describe specializations of a particular type of Stackelberg game for security as discussed below. This section provides details on this use of Stackelberg games for

modeling security domains. We first give a generic description of security domains followed by *security games*, the model by which security domains are formulated in the Stackelberg game framework.

2.1 Stackelberg Security Game

In Stackelberg Security Games, a defender must perpetually defend a set of targets T using a limited number of resources, whereas the attacker is able to surveil 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 on patrols or checkpoints, e.g., scheduling checkpoints at the LAX airport or assigning federal air marshals to protect flight tours. The pure strategy for an attacker represents an attack at a target, e.g., a flight. The *mixed strategy* of the defender is a probability distribution over the pure strategies. Additionally, with each target are also associated a set of payoff values that define the utilities for both the defender and the attacker in case of a successful or a failed attack.

A key assumption of Stackelberg Security Games (we will sometimes refer to them as simply security games) is that the payoff of an outcome depends only on the target attacked, and whether or not it is *covered* (protected) by the defender [21]. The payoffs do *not* depend on the remaining aspects of the defender allocation. For example, if an adversary succeeds in attacking target t_1 , the penalty for the defender is the same whether the defender was guarding target t_2 or not.

	Defender		Attacker	
Target	Covered	Uncovered	Covered	Uncovered
t_1	10	0	-1	1
t_2	0	-10	-1	1

Table 1. Example of a security game with two targets.

This allows us to compactly represent the payoffs of a security game. Specifically, a set of four payoffs is associated with each target. These four payoffs are the rewards and penalties to both the defender and the attacker in case of a successful or an unsuccessful attack, and are sufficient to define the utilities for both players for all possible outcomes in the security domain. More 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. Table 1 shows an example security game with two targets, t_1 and t_2 . In this example game, if the defender was covering target t_1 and the attacker attacked t_1 , the defender would get 10 units of reward whereas the attacker would receive -1 units. We make the assumption that in a security game it is always better for the defender to cover a 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 security games are not necessarily zero-sum.

2.2 Solution Concept: Strong Stackelberg Equilibrium

The solution to a security game 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 for himself. The defender's mixed strategy is a probability distribution over all pure strategies, where a pure strategy is an assignment of the defender's limited security resources to targets. This solution concept is known as a Stackelberg equilibrium [23].

The most commonly adopted version of this concept in related literature is called Strong Stackelberg Equilibrium (SSE) [2, 5, 30, 37]. In security games, the mixed strategy of the defender is equivalent to the probabilities that each target t is covered by the defender, denoted by $C = \{c_t\}$ [22]. Furthermore, it is enough to consider a pure strategy of the rational adversary [5], which is to attack a target t . The expected utility for defender for a strategy profile (C, t) is defined as $U_d(t, C) = c_t U_d^c(t) + (1 - c_t) U_d^u(t)$, and a similar form for the adversary. A SSE for the basic security games (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 defender plays a best-response: $U_d(t^*, C^*) \geq U_d(t(C), C)$ for all defender's strategy C where $t(C)$ is the attacker's response against the defender strategy C .*
2. *The attacker plays a best-response: $U_a(t^*, C^*) \geq U_a(t, C^*)$ for all target t .*
3. *The attacker breaks ties in favor of the defender: $U_d(t^*, C^*) \geq U_d(t', C^*)$ for all target t' such that $t' = \arg\max_t U_a(t, C^*)$*

The assumption that the follower will always break ties in favor of the leader 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 [37]. 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. Finally, although initial applications relied on the SSE solution concept, we have since proposed new solution concepts that are more robust against various uncertainties in the model [42, 1, 34] and have used these robust solution concepts in some of the later applications.

In the following sections, we present three key challenges in solving real-world security problems which are summarized in Figure 1: 1) scaling up to real-world sized security problems, 2) handling multiple uncertainties in security games, and 3) dealing with bounded rationality of human adversaries. While Figure 1 does not provide an exhaustive overview of all research in SSG, it provides a general overview of the areas of research, and a roadmap to the rest of the book chapter. In each case, we will use a domain example to motivate the specific challenge and then outline the key algorithmic innovation needed to address the challenge.

	Scalability				Uncertainty		Attacker Bounded Rationality
Challenge	Large defender strategy space	Large defender & attacker strategy spaces	Mobile resources & moving targets	Multiple boundedly rational attackers	Unifications of uncertainties	Dynamic execution uncertainty	Learning attacker behaviors
Domain Example	Federal Air Marshals Service	Road Network Security	Ferry Protection	Fishery Protection	Security in LAX Airport	Security in Transit System	Green Security Domains: wildlife/ fishery protection
Algorithmic Solution	ASPEN: strategy generation approach	RUGGED: double oracle approach	CASS: compact representation of strategy	MIDAS: cutting plane approach	URAC: multi-dimensional reduction & divide-and-conquer approach	Markov Decision Processes approach	Behavioral models & Human subject experiments

Fig. 1. Summary of Real-world Security Challenges

3 Addressing Scalability in Real-world Problems

For simple examples of security games, such as the one shown in the previous section, the Strong Stackelberg Equilibrium can be calculated by hand. However, as the size of the game increases, hand calculation is no longer feasible and an algorithmic approach for generating the SSE becomes necessary. Conitzer and Sandholm [5] provided the first complexity results and algorithms for computing optimal commitment strategies in Bayesian Stackelberg games, including both pure and mixed-strategy commitments. An improved algorithm for solving Bayesian Stackelberg games, DOBSS [30], is central to the fielded application ARMOR in use at the Los Angeles International Airport [14]. These early works required that the full set of pure strategies for both players be considered when modeling and solving Stackelberg security games. However, many real world problems feature billions of pure strategies for either the defender and / or the attacker. Such large problem instances cannot even be represented in modern computers, let alone solved using previous techniques.

In addition to large strategy spaces, there are other scalability challenges presented by different real world security domains. There are domains where, rather than being static, the targets are moving and thus the security resources need to be mobile and move in a continuous space to provide protection. There are also domains where the attacker may not conduct the careful surveillance and planning that is assumed for a Strong Stackelberg Equilibrium and thus it is important to model the bounded rationality of the attacker in order to predict their behavior. In the former case, both the defender and attacker’s strategy spaces are infinite. In the latter case, computing the optimal strategy for the defender given attacker behavioral (bounded rationality) model is computationally expensive. In this section, we thus highlight the critical scalability challenges faced to bring Stackelberg security games to the real world and the research contributions that served to address these challenges.

3.1 Scale Up with Large Defender Strategy Spaces

This section provides an example of a research challenge in security games where the number of defender strategies is too enormous to be enumerated in computer memory. In this section as in others that will follow, we will first provide a domain example motivating the challenge and then the algorithmic solution for the challenge.

Domain Example – IRIS for US Federal Air Marshals Service. The US Federal Air Marshals Service (FAMS) allocates air marshals to flights departing from and arriving in the United States to dissuade potential aggressors and prevent an attack should one occur. Flights are of different importance based on a variety of factors such as the numbers of passengers, the population of source and destination cities, and international flights from different countries. Security resource allocation in this domain is significantly more challenging than for ARMOR: a limited number of air marshals need to be scheduled to cover thousands of commercial flights each day. Furthermore, these air marshals must be scheduled on tours of flights that obey various constraints (e.g., the time required to board, fly, and disembark). Simply finding schedules for the marshals that meet all of these constraints is a computational challenge. For an example scenario with 1000 flights and 20 marshals, there are over 10^{41} possible schedules that could be considered. Yet there are currently tens of thousands of commercial flights flying each day, and public estimates state that there are thousands of air marshals that are scheduled daily by the FAMS [20]. Air marshals must be scheduled on tours of flights that obey logistical constraints (e.g., the time required to board, fly, and disembark). An example of a schedule is an air marshal assigned to a round trip from New York to London and back.

Against this background, the IRIS system (Intelligent Randomization In Scheduling) has been developed and deployed by FAMS since 2009 to randomize schedules of air marshals on international flights. In IRIS, the targets are the set of n flights and the attacker could potentially choose to attack one of these flights. The FAMS can assign $m < n$ air marshals that may be assigned to protect these flights.

Since the number of possible schedules exponentially increases with the number of flights and resources, DOBSS is no longer applicable to the FAMS domain. Instead, IRIS uses the much faster ASPEN algorithm [11] to generate the schedule for thousands of commercial flights per day.

Algorithmic Solution – Incremental Strategy Generation (ASPEN). In this section, we describe one particular algorithm ASPEN, that computes strong Stackelberg equilibria (SSE) in domains with a *very large* number of pure strategies (up to billions of actions) for the defender [11]. ASPEN builds on the insight that in many real-world security problems, there exist solutions with *small support sizes*, which are mixed strategies in which only a small set of pure strategies are played with positive probability [24]. ASPEN exploits this by using a *incremental strategy generation* approach for the defender, in which defender pure strategies are iteratively generated and added to the optimization formulation.

In ASPEN’s security game, the attacker can choose any of the flights to attack, and each air marshal can cover one schedule. Each schedule here is a feasible set of targets

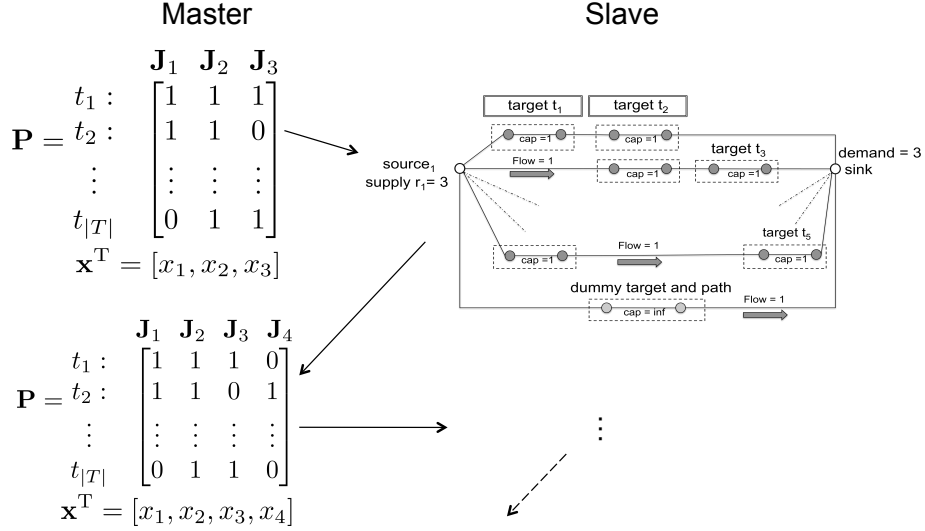


Fig. 2. Strategy generation employed in ASPEN: The schedules for a defender are generated iteratively. The *slave* problem is a novel minimum-cost integer flow formulation that computes the new pure strategy to be added to \mathbf{P} ; \mathbf{J}_4 is computed and added in this example.

that can be covered together; for the FAMS, each schedule would represent a flight tour which satisfies all the logistical constraints that an air marshal could fly. For example, $\{t_1, t_2\}$ would be a flight schedule, where t_1 is an outbound flight and t_2 is an inbound flight for one air marshal. A *joint schedule* then would assign every air marshal to a flight tour, and there could be exponentially many joint schedules in the domain. A pure strategy for the defender in this security game is a joint schedule. Thus for example, if there are two air marshals, one possible joint schedule would be $\{\{t_1, t_2\}, \{t_3, t_4\}\}$, where the first air marshal covers flights t_1 and t_2 , and the second covers flights t_3 and t_4 . As mentioned previously, ASPEN employs incremental strategy generation since all the defender pure strategies cannot be enumerated for such a massive problem. ASPEN decomposes the problem into a *master* problem and a *slave* problem, which are then solved iteratively. Given a number of pure strategies, the master solves for the defender and the attacker optimization constraints, while the slave is used to generate a new pure strategy for the defender in every iteration. *This incremental, iterative strategy generation process allows ASPEN to avoid generation of the entire set of pure strategies.* In other words, by exploiting the small support size mentioned above, only a few pure strategies get generated via the iterative process; and yet we are guaranteed to reach the optimal solution.

The iterative process is graphically depicted in Figure 2. The master operates on the pure strategies (joint schedules) generated thus far, which are represented using the matrix \mathbf{P} . Each column of \mathbf{P} , \mathbf{J}_j , is one pure strategy (or joint schedule). An entry P_{ij} in the matrix \mathbf{P} is 1 if a target t_i is covered by joint-schedule \mathbf{J}_j , and 0 otherwise. For example, in Figure 2, the joint schedule J_3 covers target t_1 but not target t_2 . The objec-

tive of the master problem is to compute x , the optimal mixed strategy of the defender over the pure strategies in \mathbf{P} . The objective of the slave problem is to generate the best joint schedule to add to \mathbf{P} . The best joint schedule is identified using the concept of *reduced costs*, which measures if a pure strategy can potentially increase the defender’s expected utility (the details of the approach are provided in [11]). While a naïve approach would be to iterate over all possible pure strategies to identify the pure strategy with the maximum potential, ASPEN uses a novel minimum-cost integer flow problem to efficiently identify the best pure strategy to add. ASPEN always converges on the optimal mixed strategy for the defender.

Employing incremental strategy generation for large optimization problems is not an “out-of-the-box” approach, the problem has to be formulated in a way that allows for domain properties to be exploited. The novel contribution of ASPEN is to provide a linear formulation for the master and a minimum-cost integer flow formulation for the slave, which enables the application of strategy generation techniques.

3.2 Scale Up with Large Defender & Attacker Strategy Spaces

Whereas the previous section focused on domains where only the defender’s strategy was difficult to enumerate, we now turn to domains where both defender and attacker strategies are difficult to enumerate. Once again we provide a domain example and then an algorithmic solution.

Domain Example – Road Network Security. One area of great importance is securing urban city networks, transportation networks, computer networks and other network centric security domains. For example, after the terrorist attacks in Mumbai of 2008 [4], the Mumbai police started setting up vehicular checkpoints on roads. We can model the problem faced by the Mumbai police as a security game between the Mumbai police and an attacker. In this urban security game, the pure strategies of the defender correspond to allocations of resources to edges in the network—for example, an allocation of police checkpoints to roads in the city. The pure strategies of the attacker correspond to paths from any *source* node to any *target* node—for example, a path from a landing spot on the coast to the airport.

The strategy space of the defender grows exponentially with the number of available resources, whereas the strategy space of the attacker grows exponentially with the size of the network. For example, in a fully connected graph with 20 nodes and 190 edges, the number of defender pure strategies for only 5 defender resources is $\binom{190}{5}$ or almost 2 billion, while the number of attacker pure strategies (i.e., paths without cycles) is on the order of 10^{18} . Real-world networks are significantly larger, e.g., the entire road network of the city of Mumbai has 9,503 nodes (intersections) and 20,416 edges (streets), and the security forces can deploy dozens (but not as many as number of edges) of resources. In addressing this computational challenge, novel algorithms based on incremental strategy generation have been able to generate randomized defender strategies that scale up to the entire road network of Mumbai [13].

Algorithmic Solution – Double Oracle Incremental Strategy Generation (RUGGED)

In domains such as the urban network security setting, the number of pure strategies of

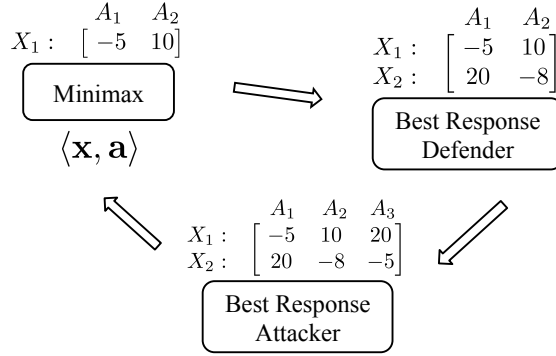


Fig. 3. Strategy Generation employed in RUGGED: The pure strategies for both the defender and the attacker are generated iteratively.

both the defender and the attacker are exponentially large. In this section, we describe the RUGGED algorithm [12], which generates pure strategies for both the defender and the attacker.

RUGGED models the domain as a zero-sum game, and computes the minimax equilibrium, since the minimax strategy is equivalent to the SSE in zero-sum games. Figure 3 shows the working of RUGGED: at each iteration, the Minimax module generates the optimal mixed strategies $\langle \mathbf{x}, \mathbf{a} \rangle$ for the two players for the current payoff matrix, the Best Response Defender module generates a new strategy for the defender that is a best response against the attacker's current strategy \mathbf{a} , and the Best Response Attacker module generates a new strategy for the attacker that is a best response against the defender's current strategy \mathbf{x} . The rows X_i in the figure are the pure strategies for the defender, they would correspond to an allocation of checkpoints in the urban road network domain. Similarly, the columns A_j are the pure strategies for the attacker, they represent the attack paths in the urban road network domain. The values in the matrix represent the payoffs to the defender. For example, in Figure 3, the row denoted by X_1 indicates that there was one checkpoint setup, and it provides a defender payoff of -5 against attacker strategy (path) A_1 , and a payoff of 10 against attacker strategy (path) A_2 .

In Figure 3, we show that RUGGED iterates over two oracles: the defender best response and the attacker best response oracles. In this case, the defender best response oracle has added a strategy X_2 , and the attacker best response oracle then adds a strategy A_3 . The algorithm stops when neither of the generated best responses improve on the current minimax strategies.

The contribution of RUGGED is to provide the mixed integer formulations for the best response modules which enable the application of such a strategy generation approach. The key once again is that RUGGED is able to converge to the optimal solution without enumerating the entire space of defender and attacker strategies. However, originally RUGGED could only compute the optimal solution for deploying up to 4 resources in real-city network with 250 nodes within a time frame of 10 hours (the complexity of this problem can be estimated by observing that both the best response problems

are NP-hard themselves [12]). More recent work [13] builds on RUGGED and proposes SNARES, which allows scale-up to the entire city of Mumbai, with 10–15 checkpoints.

3.3 Scale Up with Mobile Resources & Moving Targets

Whereas the previous two sections focused on incremental strategy generation as an approach for scale-up this section introduces another approach: use of compact marginal probability representations. This alternative approach is shown in use in the context of a new application of protecting ferries.

Domain Example – Ferry Protection for the US Coast Guard. The United States Coast Guard is responsible for protecting domestic ferries, including the Staten Island Ferry in New York, from potential terrorist attacks. here are a number of ferries carrying hundreds of passengers in many waterside cities. These ferries are attractive targets for an attacker who can approach the ferries with a small boat packed with explosives at any time; this attacker’s boat may only be detected when it comes close to the ferries. Small, fast, and well-armed patrol boats can provide protection to such ferries by detecting the attacker within a certain distance and stop him from attacking with the armed weapons. However, the numbers of patrol boats are often limited, thus the defender cannot protect the ferries at all times and locations. We thus developed a game-theoretic system for scheduling escort boat patrols to protect ferries, and this has been deployed at the Staten Island Ferry since 2013 [6].

The key research challenge is the fact that the ferries are continuously moving in a continuous domain, and the attacker could attack at any moment in time. This type of moving targets domain leads to game-theoretic models with continuous strategy spaces, which presents computational challenges. Our theoretical work showed that while it is “safe” to discretize the defender’s strategy space (in the sense that the solution quality provided by our work provides a lower bound), discretizing the attacker’s strategy space would result in loss of utility (in the sense that this would provide only an upper bound, and thus an unreliable guarantee of true solution quality). We developed a novel algorithm that uses a compact representation for the defender’s mixed strategy space while being able to exactly model the attacker’s continuous strategy space. The implemented algorithm, running on a laptop, is able to generate daily schedules for escort boats with guaranteed expected utility values.

Algorithmic Solution – Compact Strategy Representation (CASS). In this section, we describe the CASS (Solver for Continuous Attacker Strategy) algorithm [6] for solving security problems where the defender has mobile patrollers to protect a set of mobile targets against the attacker who can attack these moving targets at any time during their movement. In these security problems, the sets of pure strategies for both the defender and attacker are continuous w.r.t the continuous spatial and time components of the problem domain. The CASS algorithm attempts to compute the optimal mixed strategy for the defender without discretizing the attacker’s continuous strategy set; it exactly models this set using sub-interval analysis which exploits the piecewise-linear structure



Fig. 4. Escort boats protecting the Staten Island Ferry use strategies generated by our system.

of the attacker's expected utility function. The insight of CASS is to compactly represent the defender's mixed strategies as a *marginal* probability distribution, overcoming the short-coming of an exponential number of pure strategies for the defender.

CASS casts problems such as the ferry protection problem mentioned above as a *zero-sum* security game in which targets move along a *one-dimensional* domain, i.e., a straight line segment connecting two terminal points. This *one-dimensional* assumption is valid as in real-world domains such as ferry protection, ferries normally move back-and-forth in a straight line between two terminals (i.e., ports) around the world. Although the targets' locations vary w.r.t time changes, these targets have a fixed daily schedule, meaning that determining the locations of the targets at a certain time is straightforward. The defender has mobile patrollers (i.e., boats) that can move along between two terminals to protect the targets. While the defender is trying to protect the targets, the attacker will decide to attack a certain target at a certain time. The probability that the attacker successfully attacks depends on the positions of the patroller at that time. Specifically, each patroller possesses a protective circle of radius within which she can detect and try to intercept any attack, whereas she is incapable of detecting the attacker prior to that radius.

In CASS, the defender's strategy space is discretized and her mixed strategy is compactly represented using flow distributions. Figure 5 shows an example of a ferry transition graph in which each node of the graph indicates a particular pair of (location, time step) for the target. Here, there are three location points namely A, B, and C on a straight line where B lies between A and C. Initially, the target is at one of these location points at the 5-minute time step. Then the target moves to the next location point which is determined based on the connectivity between these points at the 10-minute time step and so on. For example, if the target is at the location point A at the 5-minute time step, denoted by (A, 5 min) in the transition graph, it can move to the location point B or stay at location point A at the 10-minute time step. The defender follows this transition graph to protect the target.




Fig. 5. An example of a ferry transition graph

A pure strategy for the defender is defined as a trajectory of this graph, e.g., the trajectory including (A, 5 min), (B, 10 min), and (C, 15 min) indicates a pure strategy for the defender. One key challenge of this representation for the defender’s pure strategies is that the transition graph consists of an exponential number of trajectories, i.e., $O(N^T)$ where N is the number of location points and T is the number of time steps. To address this challenge, CASS proposes a compact representation of the defender’s mixed strategy. Instead of directly computing a probability distribution over pure strategies for the defender, CASS attempts to compute the marginal probability that the defender will follow a certain edge of the transition graph, e.g., the probability of being at the node (A, 5 min) and moving to the node (B, 10 min). CASS shows that *any strategy in full representation can be mapped into a compact representation as well as compact representation does not lead to any loss in solution quality*. This compact representation allows CASS to reformulate the resource-allocation problem as computing the optimal *marginal* coverage of the defender over a number of $O(NT)$ the edges of the transition graph.

3.4 Scale Up with Boundedly Rational Attackers

One key challenge of real-world security problems is that the attacker is boundedly rational; the attacker’s target choice is non-optimal. In SSGs, attacker bounded rationality is often modeled via behavior models such as Quantal Response (QR) [25, 27]. In general, QR attempts to predict the probability the attacker will choose each target with the intuition is that the higher the expected utility at a target is, the more likely that the adversary will attack that target. Another behavioral model that was recently shown to provide higher prediction accuracy in predicting the attacker’s behavior than QR is Subjective Utility Quantal Response (SUQR) [29]. SUQR is motivated by the lens model which suggested that evaluation of adversaries over targets is based on a linear combination of multiple observable features [3]. However, handling multiple attackers with these behavioral models in the context of large defender’s strategy space is computational challenge. In this section, we mainly focus on handling the scalability problem given behavioral models of the attacker. The problem of handling the attacker’s bounded rationality (e.g., modeling and learning) is explained in detail in Section 5.

To handle the problem of large defender’s strategy space given behavioral models of attackers, we introduce yet another technique of scaling up, which is similar to the incremental strategy generation. Instead, here we use incremental marginal space refinement. We use the compact marginal representation, discussed earlier, but refine that space incrementally if the solution produces violates the necessary constraints.

Domain Example– Fishery Protection for US Coast Guard. Fisheries are a vital natural resource from both an ecological and economic standpoint. However, fish stocks around the world are threatened with collapse due to illegal, unreported, and unregulated (IUU) fishing. The United States Coast Guard (USCG) is tasked with the responsibility of protecting and maintaining the nation’s fisheries. To this end, the USCG deploys resources (both air and surface assets) to conduct patrols over fishery areas in order to deter and mitigate IUU fishing. Due to the large size of these patrol areas and the limited patrolling resources available, it is impossible to protect an entire fishery from IUU fishing at all times. Thus, an intelligent allocation of patrolling resources is critical for security agencies like the USCG.

Natural resource conservation domains such as fishery protection raise a number of new research challenges. In stark contrast to counter-terrorism settings, there is frequent interaction between the defender and attacker in these resource conservation domains. This distinction is important for three reasons. First, due to the comparatively low stakes of the interactions, rather than a handful of persons or groups, the defender must protect against numerous adversaries (potentially hundreds or even more), each of which may behave differently. Second, frequent interactions make it possible to collect data on the actions of the adversaries actions over time. Third, the adversaries are less strategic given the short planning windows between actions.

Algorithmic Solution – Incremental Constraint Generation (MIDAS). Generating effective strategies for domains such as fishery protection requires an algorithmic approach which is both *scalable* and *robust*. For scalability, the defender is responsible for protecting a large patrol area and therefore must consider a large strategy space. Even if the patrol area is discretized into a grid or graph structure, the defender must still reason over an exponential number of patrol strategies. For robustness, the defender must protect against *multiple* boundedly rational adversaries. Bounded rationality models, such as the quantal response (QR) model [27] and the subjective utility quantal response (SUQR) model [29], introduce stochastic actions, relaxing the strong assumption in classical game theory that all players are perfectly rational and utility maximizing. These models are able to better predict the actions of human adversaries and thus lead the defender to choose strategies that perform better in practice. However, both QR and SUQR are non-linear models resulting in a computationally difficult optimization problem for the defender. Combining these factors, MIDAS models a population of boundedly rational adversaries and utilizes available data to learn the behavior models of the adversaries using the subjective utility quantal response (SUQR) model in order to improve the way the defender allocates its patrolling resources.

Previous work on boundedly rational adversaries has considered the challenges of scalability and robustness separately, in [41, 39] and [38, 10], respectively. The MIDAS

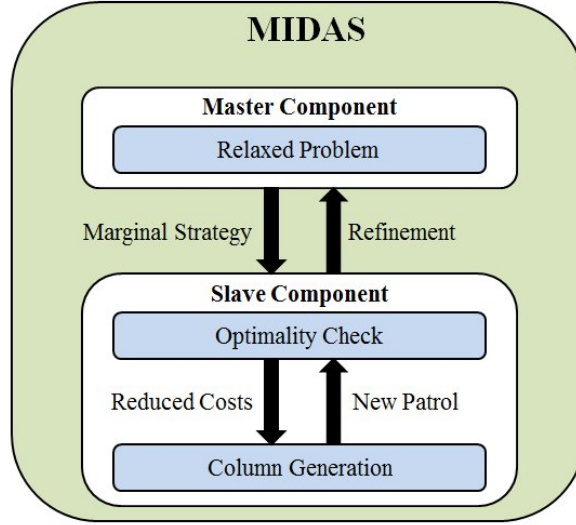


Fig. 6. Overview of the multiple iterative process within the MIDAS algorithm

algorithm was introduced to merge these two research threads for the first time by addressing scalability and robustness simultaneously. Figure 6 provides a visual overview of how MIDAS operates as an iterative process. Similar to the ASPEN algorithm described earlier, given the sheer complexity of the game being solved, the problem is decomposed using a master-slave formulation. The master utilizes multiple simplifications to create a relaxed version of the original problem which is more efficient to solve. First, a piecewise linear approximation of the security game is taken to make the optimization problem both linear and convex. Second, the complex spatio-temporal constraints associated with patrols are initially ignored and then incrementally added back using cut generation. In other words, we ignore the spatio-temporal constraint that a patroller cannot simply appear and disappear at different locations instantaneously; and that a patroller must pass through regions connecting two different regions if the patroller is to go from one region to another. This significantly simplifies the master problem.

Due to the relaxations, solving the master produces a marginal strategy x which is a probability distribution over targets. However, the defender ultimately needs a probability distribution over patrols. Additionally, since not all of the spatio-temporal constraints are considered in the master, the relaxed solution x may not be a feasible solution to the original problem. Therefore, the slave checks if the marginal strategy x can be expressed as a linear combination, i.e., probability distribution, of patrols. Otherwise, the marginal distribution is infeasible for the original problem. However, given the exponential number of patrol strategies, even performing this optimality check is intractable. Thus, column generation is used *within* the slave where only a small set of patrols is considered initially in the optimality check and the set is expanded over time. Much like previous examples of column generation in security games, e.g., [11], new patrols are




Fig. 7. Uncertainty space and algorithms

added by solving a minimum cost network flow problem using reduced cost information from the optimality check. If the optimality check fails, then the slave generates a cut which is returned to refine and constrain the master, incrementally bringing it closer to the original problem. The entire process is repeated until an optimal solution is found. Finally, MIDAS has been successfully deployed and evaluated by the USCG in the Gulf of Mexico.

4 Addressing Uncertainty in Real-world Problems

The standard security game model features a number of strong assumptions including that the defender has perfect information about the game payoff matrix as well as the attacker's behavioral model. Additionally, the defender is assumed to be capable of exactly executing the computed patrolling strategy. However, uncertainty is endemic in real-world security domains and thus it may be impossible or impractical for the defender to accurately estimate various aspects of the game. Also, there are any number of practicalities and unforeseen events that may force the defender to change their patrolling strategy. These types of uncertainty can significantly deteriorate the effectiveness of the defender's strategy and thus addressing uncertainty when generating strategies is a key challenge of solving real-world security problems. This section describes several approaches for dealing with various types of uncertainties in SSGs.

We first summarize the major types of uncertainties in SSGs as a 3-dimensional uncertainty space with the following three dimensions (Figure 7): 1) uncertainty in the adversary's payoffs; 2) uncertainty related to the defender's strategy (including uncertainty in the defender's execution and the attacker's observation); and 3) uncertainty in the adversary's rationality. These dimensions refer to three key attributes which affect both players' utilities. The origin of the uncertainty space corresponds to the case with no uncertainty. Figure 7 also shows existing algorithms for addressing uncertainty in SSGs which follow the two different approaches: 1) applying robust optimization

techniques using uncertainty intervals to represent uncertainty in SSGs. For example, BRASS [32] is a robust algorithm that only addresses attacker-payoff uncertainty, RECON [42] is another robust algorithm that focuses on addressing defender-strategy uncertainty, and Monotonic Maximin [15] is to handle the uncertainty in the attacker’s bounded rationality. Finally, URAC [28] is a unified robust algorithm that handles all types of uncertainty; and 2) following Bayesian Stackelberg game model with dynamic execution uncertainty in which the uncertainty is represented using Markov Decision Process (MDP) where the time factor is incorporated.

In the following, we present two algorithmic solutions which are the representatives of these two approaches: URAC – a unified robust algorithm to handle all types of uncertainty with uncertainty intervals and the MDP-based algorithm to handle execution uncertainty with an MDP representation of uncertainty.

4.1 Security Patrolling with Unified Uncertainty Space

Domain Example – Security in Los Angeles International Airport. Los Angeles International Airport (LAX) is the largest destination airport in the United States and serves 60-70 million passengers per year. The LAX police use diverse measures to protect the airport, which include vehicular checkpoints, police units patrolling the roads to the terminals, patrolling inside the terminals (with canines), and security screening and bag checks for passengers. The application of our game-theoretic approach is focused on two of these measures: (1) placing vehicle checkpoints on inbound roads that service the LAX terminals, including both location and timing, and (2) scheduling patrols for bomb-sniffing canine units at the different LAX terminals. The eight different terminals at LAX have very different characteristics, like physical size, passenger loads, international versus domestic flights, etc. These factors contribute to the differing risk assessments of these eight terminals. Furthermore, the numbers of available vehicle checkpoints and canine units are limited by resource constraints. Thus, it is challenging to optimally allocate these resources to improve their effectiveness while avoiding patterns in the scheduled deployments.

The ARMOR system (Assistant for Randomized Monitoring over Routes) focuses on two of the security measures at LAX (checkpoints and canine patrols) and optimizes security resource allocation using Bayesian Stackelberg games. Take the vehicle checkpoints model as an example. Assuming that there are n roads, the police’s strategy is placing $m < n$ checkpoints on these roads where m is the maximum number of checkpoints. ARMOR randomizes allocation of checkpoints to roads. The adversary may conduct surveillance of this mixed strategy and may potentially choose to attack through one of these roads. ARMOR models different types of attackers with different payoff functions, representing different capabilities and preferences for the attacker. ARMOR has been successfully deployed since August 2007 at LAX [14].

Although standard SSG-based solutions (i.e., DOBSS) have been demonstrated to improve the defender’s patrolling effectiveness significantly, there remains potential improvements that can be made to further enhance the quality of such solutions such as taking uncertainties in payoff values, in the attacker’s rationality, and in defender’s execution into account. Therefore, we propose the unified robust algorithm, URAC,



Fig. 8. LAX checkpoints are deployed using ARMOR.

to handle these types of uncertainties by maximizing the defender’s utility against the worst-case scenario resulting from these uncertainties.

Algorithmic Solution – Uncertainty Dimension Reduction (URAC). In this section, we present the robust URAC (Unified Robust Algorithmic framework for addressing unCertainties) algorithm for addressing a combination of all uncertainty types [28]. Consider an SSG where there is uncertainty in the attacker’s payoff, the defender’s strategy (including the defender’s execution and the attacker’s observation), and the attacker’s behavior, URAC represents all these uncertainty types (except for the attacker’s behaviors) using uncertainty intervals. Instead of knowing exactly values of these game attributes, the defender only has prior information w.r.t the upper bounds and lower bounds of these attributes. For example, the attacker’s reward if successfully attacking a target t is known to lie within the interval $[1, 3]$. Furthermore, URAC assumes the attacker monotonically responds to the defender’s strategy. In other words, the higher the expected utility of a target, the more likely that the attacker will attack that target; however, the precise attacking probability is unknown for the defender. This monotonicity assumption is motivated by the Quantal Response model — a well-known human behavioral model for capturing the attacker’s decision making [27].

Based on these uncertainty assumptions, URAC attempts to compute the optimal strategy for the defender by maximizing her utility against the worst-case scenario of uncertainty. The key challenge of this optimization problem is that it involves several types of uncertainty, resulting in multiple minimization steps for determining the worst-case scenario. Nevertheless, URAC introduces a unified representation of all these uncertainty types as a uncertainty set of attacker’s responses. Intuitively, despite of any type of uncertainty mentioned above, what finally affects the defender’s utility is the attacker’s response, which is unknown to the defender due to uncertainty. As a result, URAC can represent the robust optimization problem as a single maximin problem.

However, the infinite uncertainty set of the attacker’s responses depends on the planned mixed strategy for the defender, making this maximin problem difficult to solve if directly applying the traditional method (i.e., taking the dual maximization of the inner minimization of maximin and merging it with the outer maximization — maximin now can be represented a single maximization problem). Therefore, URAC proposes a divide-and-conquer method in which the defender’s strategy set is divided into subsets such that the uncertainty set of the attacker’s responses is the same for every defender strategy within each subset. This division leads to multiple sub-maximin problems which can be solved by using the traditional method. The optimal solution of the original maximin problem is now can be computed as a maximum over all the sub-maximin problems.

4.2 Security Patrolling with Dynamic Execution Uncertainty

Domain Example – TRUSTS for Security in Transit Systems. Urban transit systems face multiple security challenges, including deterring fare evasion, suppressing crime and counter-terrorism. In particular, in some urban transit systems, including the Los Angeles Metro Rail system, passengers are legally required to purchase tickets before entering but are not physically forced to do so (Figure 9). Instead, security personnel are dynamically deployed throughout the transit system, randomly inspecting passenger tickets. This proof-of-payment fare collection method is typically chosen as a more cost-effective alternative to direct fare collection, i.e., when the revenue lost to fare evasion is believed to be less than what it would cost to directly preclude it. In the case of Los Angeles Metro, with approximately 300,000 riders daily, this revenue loss can be significant; the annual cost has been estimated at \$5.6 million [9]. The Los Angeles Sheriffs Department (LASD) deploys uniformed patrols on board trains and at stations for fare-checking (and for other purposes such as crime prevention). The LASD’s current approach relies on humans for scheduling the patrols, which places a tremendous cognitive burden on the human schedulers who must take into account all of the scheduling complexities (e.g., train timings, switching time between trains, and schedule lengths).

The TRUSTS system (Tactical Randomization for Urban Security in Transit Systems) models the patrolling problem as a leader-follower Stackelberg game [43]. The leader (LASD) pre-commits to a mixed strategy patrol (a probability distribution over all pure strategies), and riders observe this mixed strategy before deciding whether to buy the ticket or not. Both ticket sales and fines issued for fare evasion translate into revenue for the government. Therefore the utility for the leader is the total revenue (total ticket sales plus penalties). The main computational challenge is the exponentially many possible patrol strategies, each subject to both the spatial and temporal constraints of travel within the transit network under consideration. To overcome this challenge, TRUSTS uses a compact representation of the strategy space which captures the spatiotemporal structure of the domain.

The LASD conducted field tests of this TRUSTS system in the LA Metro in 2012, and one of the feedback comments from the officers was that patrols are often interrupted due to execution uncertainty such as emergencies and arrests.



Fig. 9. TRUSTS for transit systems

Algorithmic Solution – Marginal MDP Strategy Representation. Utilizing techniques from planning under uncertainty (in particular Markov Decision Processes), we proposed a general approach to dynamic patrolling games in uncertain environments, which provides patrol strategies with contingency plans[16]. This led to schedules now being loaded onto smartphones and given to officers. If interruptions occur, the schedules are then automatically updated on the smartphone app. The LASD has conducted successful field evaluations using the smartphone app, and the TSA is currently evaluating it toward nationwide deployment. We now describe the solution approach in more detail. Note that the targets, e.g., trains normally follow predetermined schedules, thus timing is an important aspect which determines the effectiveness of the defender’s patrolling schedules (the defender needs to be at the right location at a specific time in order to protect these moving targets). However, as a result of execution uncertainty (e.g., emergencies or errors), the defender could not carry out her planned patrolling schedule in later time steps. For example, in real-world trials for TRUSTS carried out by Los Angeles Sheriff’s Department (LASD), there is interruption (due to writing citations, felony arrests, and handling emergencies) in a significant fraction of the executions, causing the officers to miss the train they are supposed to catch as following the pre-generated patrolling schedule.

In this section, we present the Bayesian Stackelberg game model for security patrolling with dynamic execution uncertainty introduced by [16] in which the uncertainty is represented using Markov Decision Processes (MDP). The key advantage of this game-theoretic model is that patrol schedules which are computed based on Stackelberg equilibrium have contingency plans to deal with interruptions and are robust against execution uncertainty. Specifically, the security problem with execution uncertainty is represented as a two-player Bayesian Stackelberg game between the defender and the attacker. The defender has multiple patrol units while there are also multiple types of attackers which are unknown to the defender. A (naive) patrol schedule consists of a set of sequenced commands in the following form: at time t , the patrol unit should be at location l and execute patrol action a . This patrol action a will take the

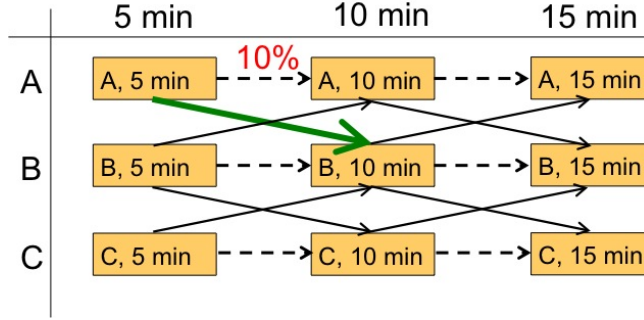


Fig. 10. An example of execution uncertainty in a transition graph

unit to the next location and time if successfully executed. However, due to execution uncertainty, the patrol unit may end up at a different location and time. Figure 10 shows an example of execution uncertainty in a transition graph where if the patrol unit is currently at location A at the 5-minute time step, she is supposed to take the on-train action to move to location B in the next time step. However, unlike CASS for ferry protection in which the defender's action is deterministic, there is a 10% chance that she will still stay at location A due to execution uncertainty. This interactions of the defender with the environment when executing patrol can be represented as an MDP.

In essence, the transition graph as represented above is augmented to indicate the possibility that there are multiple uncertain outcomes possible from a given state. Solving this transition graph results in marginals over MDP policies. When a sample MDP policy is obtained and loaded on to a smart phone, it provides a patroller not only the current action, but contingency actions should the current action fail or succeed. So the MDP policy provides options for the patroller, allowing the system to handle execution uncertainty. A key challenge of computing the SSE for this type of security problem is that the dimension of the space of mixed strategies for the defender is exponential in the number of states in terms of the defender's times and locations. Therefore, instead of directly computing the mixed strategy, the defender attempts to compute the marginal probabilities of each patrolling unit reaching a state $s = (t, l)$, and taking action a which have dimensions polynomial in the sizes of the MDPs (the details of this approach are provided in [16]).

5 Addressing Bounded Rationality in Real-world Problems

Game theory models the strategic interactions between multiple players who are assumed to be perfectly rational, i.e., they will always select the optimal strategy available to them. This assumption may be applicable for high-stakes security domains such as infrastructure protection where presumably the adversary will conduct careful surveillance and planning before attacking. However, there are other security domains where the adversary may not be perfectly rational due to short planning windows or because the adversary is less strategic due to lower stakes associated with attacking. Security




Fig. 11. Examples of illegal activities in green security domains

strategies generated under the assumption of a perfectly rational adversary are not necessarily as effective as would be feasible against a less-than-optimal response. Therefore, addressing the boundedly rationality exhibited by human adversaries is a fundamental challenge for applying security games to wide variety of domains.

Domain Example – Green Security Domains. A number of our newer applications are focused on resource conservation, through suppression of environmental crime. One area is protecting forests [17], where we must protect a continuous forest area from extractors by patrols through the forest that seek to deter such extraction activity. With limited resources for performing such patrols, a patrol strategy will seek to distribute the patrols throughout the forest, in space and time, in order to minimize the resulting amount of extraction that occurs or maximize the degree of forest protection. This problem can be formulated as a Stackelberg game and the focus is on computing optimal allocations of patrol density [17].

Endangered species poaching is reaching critical levels as the populations of these species plummet to unsustainable numbers. The global tiger population, for example, has dropped over 95% from the start of the 1900s and has resulted in three out of nine species extinctions. Depending on the area and animals poached, motivations for poaching range from profit to sustenance, with the former being more common when profitable species such as tigers, elephants, and rhinos are the targets. To counter poaching efforts and to rebuild the species' populations, countries have set up protected wildlife reserves and conservation agencies tasked with defending these large reserves. Because of the size of the reserves and the common lack of law enforcement resources, conservation agencies are at a significant disadvantage when it comes to deterring and capturing poachers. Agencies use patrolling as a primary method of securing the park. Due to their limited resources, however, patrol managers must carefully create patrols that account for many different variables (e.g., limited patrol units to send out, multiple locations that poachers can attack at varying distances to the outpost).

Behavioral Modeling and Learning. Recently, we have conducted some research on applying ideas from behavioral game theory (e.g., prospect theory [18] and quantal response [26]) within security game algorithms. One line of approaches is based on

the quantal response model to predict the behaviors of the human adversary, and then to compute optimal defender strategies against such behavior of the adversary. These include BRQR [40] which follows the logit quantal response (QR) [26] model, and subsequent work on subjective-utility quantal response (SUQR) models [29]. The parameters of these models are estimated by experimental tuning. Data from a large set of participants on the Amazon Mechanical Turk (AMT) were collected and used to learn the parameters of the behavioral models to predict future attacks. In real-world domains like fisheries protection, or wildlife crime, there are repeated interactions between the defender and the adversary, where the game progresses in "rounds". We call this a Repeated Stackelberg Security Game (RSSG) where in each round the defender would play a particular strategy and the adversary would observe that strategy and act accordingly. In order to simulate this scenario and conduct experiments to identify adversary behavior in such repeated settings, an online RSSG game was developed (shown in Fig. 12) and deployed.



Fig. 12. Interface of the Wildlife Poaching game to simulate an RSSG

In our game, human subjects play the role of poachers looking to place a snare to hunt a hippopotamus in a protected wildlife park. The portion of the park shown in the map is actually a Google Maps view of a portion of the Queen Elizabeth National Park (QENP) in Uganda. The region shown is divided into a 5*5 grid, i.e. 25 distinct cells. Overlaid on the Google Maps view of the park is a heat-map, which represents the rangers' mixed strategy x — a cell i with higher coverage probability x_i is shown more in red, while a cell with lower coverage probability is shown more in green. As the subjects play the game and click on a particular region on the map, they were given

detailed information about the poacher’s reward, penalty and coverage probability at that region. However, the participants are unaware of the exact location of the rangers while playing the game, i.e. they do not know the pure strategy that will be played by the rangers, which is drawn randomly from mixed strategy x shown on the game interface. In our game, there were 9 rangers protecting this park, with each ranger protecting one grid cell. Therefore, at any point in time, only 9 out of the 25 distinct regions in the park are protected. A player succeeds if he places a snare in a region which is not protected by a ranger, else he is unsuccessful. Similar to the Guards and Treasures game, here also we recruited human subjects on AMT and asked them to play this game repeatedly for a set of rounds with the defender strategy changing per round based on the behavioral model being used to learn the adversary’s behavior.

While behavioral models like (QR) [26] and SUQR [29] assume that there is a homogeneous population of adversaries, in the real-world we face heterogeneous populations of adversaries. Therefore Bayesian SUQR was proposed to learn the behavioral model for each attack [38]. Protection Assistant for Wildlife Security (PAWS) is an application which was originally created using Bayesian SUQR. However, in real-world security domains, we may have very limited data, or may only have some limited information on the biases displayed by adversaries. An alternative approach is based on robust optimization: instead of assuming a particular model of human decision making, try to achieve good defender expected utility against a range of possible models. One instance of this approach is MATCH [34], which guarantees a bound for the loss of the defender to be within a constant factor of the adversary loss if the adversary responds non-optimally. Another robust solution concept is monotonic maximin [15], which tries to optimize defender utility against the worst-case monotonic adversary behavior, where monotonicity is the property that actions with higher expected utility is played with higher probability. Recently, there has been attempts to combine such robust-optimization approaches with available behavior data [10] for RSSGs, resulting in a new human behavior model called Robust SUQR. However, one question of research is how these proposed models and algorithms will fare against human subjects in RSSGs. This has been explored in recent research [19] in the ‘first-of-its-kind’ human subjects experiments in RSSGs over a period of 46 weeks with the ‘Wildlife Poaching’ game, a brief summary of which is presented below.

In our human subjects experiments in RSSGs, we observe that: (i) Existing approaches (QR, SUQR, Bayesian SUQR) [29, 38, 10] perform poorly in initial rounds, while Bayesian SUQR which is the basis for PAWS [38], perform poorly through-out all rounds; (ii) Surprisingly, simpler models like SUQR which were originally proposed for single-shot games performed better than recent advances like Bayesian SUQR and Robust SUQR which are geared specifically towards addressing repeated SSGs. Therefore, we proposed a new model called SHARP [19] which addresses the limitations of the existing models (Stochastic Human behavior model with AttRactiveness and Probability weighting) in the following way: (i) Modeling the adversarys adaptive decision making process in repeated SSGs, SHARP reasons based on success or failure of the adversarys past actions on exposed portions of the attack surface, where attack surface is defined as the n -dimensional space of the features used to model adversary behavior; (ii) Addressing limited exposure to significant portions of the attack surface in initial




Fig. 13. ARMOR evaluation results.




Fig. 14. PROTECT evaluation results: pre deployment (left) and post deployment patrols (right).

rounds, SHARP reasons about similarity between exposed and unexposed areas of the attack surface, and also incorporates a discounting parameter to mitigate adversary's lack of exposure to enough of the attack surface; (iii) Addressing the limitation that existing models do not account for the adversary's weighting of probabilities, we incorporate a two parameter probability weighting function. Based on our human subjects experiments highlighted in [19], we observe that SHARP completely outperforms existing approaches consistently over all rounds, most notably in initial rounds.

6 Addressing Field Evaluation in Real-world Problems

Evidence showing the benefits of the algorithms discussed in the previous sections is definitely an important issue that is necessary for us to answer. Unlike conceptual ideas, where we can run thousands of careful simulations under controlled conditions, we cannot conduct such experiments in the real world with our deployed applications. Nor can we provide a proof of 100% security – there is no such thing.

Instead, we focus on the specific question of: are our game-theoretic algorithm better at security resource optimization or security allocation than how they were allocated

previously, which was typically relying on human schedulers or a simple dice roll for security scheduling (simple dice roll is often the other automation that is used or offered as an alternative to our methods). We have used the following methods to illustrate these ideas. These methods range from simulations to actual field tests.

1. **Simulations (including using a machine learning attacker):** We provide simulations of security schedules, e.g., randomized patrols, assignments, comparing our approach to earlier approaches based on techniques used by human schedulers. We have a machine learning based attacker who learns any patterns and then chooses to attack the facility being protected. Game-theoretic schedulers are seen to perform significantly better in providing higher levels of protections [33, 14]. This is also shown in Figure 13.
2. **Human adversaries in the lab:** We have worked with a large number of human subjects and security experts (security officials) to have them get through randomized security schedules, where some are schedules generated by our algorithms, and some are baseline approaches for comparison. Human subjects are paid money based on the reward they collect by successfully intruding through our security schedules; again our game-theoretic schedulers perform significantly better ([31]).
3. **Actual security schedules before and after:** For some security applications, we have data on how scheduling was done by humans (before our algorithms were deployed) and how schedules are generated after deployment of our algorithms. For measures of interest to security agencies, e.g., predictability in schedules, we can compare the actual human-generated schedules vs our algorithmic schedules. Again, game-theoretic schedulers are seen to perform significantly better by avoiding predictability and yet ensuring that more important targets are covered with higher frequency of patrols. Some of this data is published [35] and is also shown in Figure 14.
4. **“Adversary” teams simulate attack:** In some cases, security agencies have deployed adversary perspective teams or mock attacker teams that will attempt to conduct surveillance to plan attacks; this is done before and after our algorithms have been deployed to check which security deployments worked better. This was done by the US Coast Guard indicating that the game-theoretic scheduler provided higher levels of deterrence [35].
5. **Real-time comparison: human vs algorithm:** This is a test we ran on the metro trains in Los Angeles. For a day of patrol scheduling, we provided head-to-head comparison of human schedulers trying to schedule 90 officers on patrols vs an automated game-theoretic scheduler. External evaluators then provided an evaluation of these patrols; the evaluators did not know who had generated each of the schedules. The results show that while human schedulers required significant effort even for generating one schedule (almost a day), and the game-theoretic scheduler ran quickly, the external evaluators rated the game theoretic schedulers higher (with statistical significance) [7].
6. **Actual data from deployment:** This is another test run on the metro trains in LA. We had a comparison of game-theoretic scheduler vs an alternative (in this case a uniform random scheduler augmented with real time human intelligence) to check fare evaders. In 21 days of patrols, the game-theoretic scheduler led to significantly higher numbers of fare evaders captured than the alternative [7, 8].

7. **Domain expert evaluation (internal and external):** There have been of course significant numbers of evaluations done by domain experts comparing their own scheduling method with game theoretic schedulers and repeatedly the game theoretic schedulers have come out ahead. The fact that our software is now in use for several years at several different important airports, ports, air-traffic, and so on, is an indicator to us that the domain experts must consider this software of some value.

7 Conclusion

Security is recognized as a world-wide challenge and game theory is an increasingly important paradigm for reasoning about complex security resource allocation. We have shown that the general model of security games is applicable (with appropriate variations) to varied security scenarios. There are applications deployed in the real world that have led to a measurable improvement in security. We presented approaches to address four significant challenges: scalability, uncertainty, bounded rationality and field evaluation in security games.

In short, we introduced specific techniques to handle each of these challenges. For scalability, we introduced three approaches: (i) incremental strategy generation for addressing the problem of large defender strategy spaces; (ii) double oracle incremental strategy generation w.r.t large defender & attacker strategy spaces; (iii) compact representation of strategies for the case of mobile resources and moving targets; and (iv) cutting plane (incremental constraint generation) for handling multiple boundedly rational attacker. For handling uncertainty we introduced two approaches: (i) dimensionality reduction in uncertainty space for addressing a unification of uncertainties; and (ii) Markov Decision Process with marginal strategy representation w.r.t dynamic execution uncertainty. In terms of handling attacker bounded rationality, we propose different behavioral models to capture the attackers' behaviors and introduce human subject experiments with game simulation to learn such behavioral models. Finally, for addressing field evaluation in real-world problems, we discussed two approaches: (i) data from deployment; and (ii) mock attacker team.

While the deployed game theoretic applications have provided a promising start, significant amount of research remains to be done. These are large-scale interdisciplinary research challenges that call upon multiagent researchers to work with researchers in other disciplines, be "on the ground" with domain experts and examine real-world constraints and challenges that cannot be abstracted away.

References

1. An, B., Tambe, M., Ordóñez, F., Shieh, E., Kiekintveld, C.: Refinement of Strong Stackelberg Equilibria in Security Games. In: Proc. of the 25th Conference on Artificial Intelligence. pp. 587–593 (2011)
2. Breton, M., Alg, A., Haurie, A.: Sequential Stackelberg Equilibria in Two-Person Games. Optimization Theory and Applications 59(1), 71–97 (1988)
3. Brunswik, E.: The conceptual framework of psychology, vol. 1. Univ of Chicago Pr (1952)

4. Chandran, R., Beitchman, G.: Battle for Mumbai Ends, Death Toll Rises to 195. Times of India (29 November 2008), http://articles.timesofindia.indiatimes.com/2008-11-29/india/27930171_1_taj-hotel-three-terrorists-nariman-house
5. Conitzer, V., Sandholm, T.: Computing the Optimal Strategy to Commit to. In: Proc. of the ACM Conference on Electronic Commerce (ACM-EC). pp. 82–90 (2006)
6. Fang, F., Jiang, A., Tambe, M.: Optimal patrol strategy for protecting moving targets with multiple mobile resources. In: AAMAS (2013)
7. Fave, F.M.D., Brown, M., Zhang, C., Shieh, E., Jiang, A.X., Rosoff, H., Tambe, M., Sullivan, J.: Security games in the field: an initial study on a transit system(extended abstract). In: International Conference on Autonomous Agents and Multiagent Systems (AAMAS) [Short paper] (2014)
8. Fave, F.M.D., Jiang, A.X., Yin, Z., Zhang, C., Tambe, M., Kraus, S., Sullivan, J.: Game-theoretic security patrolling with dynamic execution uncertainty and a case study on a real transit system. Journal of Artificial Intelligence Research, 50:321–367 (2014)
9. Hamilton, B.A.: Faregating Analysis. Report Commissioned by the LA Metro (2007), http://boardarchives.metro.net/Items/2007/11_November/20071115EMACItem27.pdf
10. Haskell, W.B., Kar, D., Fang, F., Tambe, M., Cheung, S., Denicola, L.E.: Robust protection of fisheries with compass (2014)
11. Jain, M., Kardes, E., Kiekintveld, C., Ordonez, F., Tambe, M.: Security Games with Arbitrary Schedules: A Branch and Price Approach. In: Proc. of The 24th AAAI Conference on Artificial Intelligence. pp. 792–797 (2010)
12. Jain, M., Korzhyk, D., Vanek, O., Pechoucek, M., Conitzer, V., Tambe, M.: A Double Oracle Algorithm for Zero-Sum Security games on Graphs. In: Proc. of The 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS) (2011)
13. Jain, M., Tambe, M., Conitzer, V.: Security scheduling for real-world networks. In: AAMAS (2013)
14. Jain, M., Tsai, J., Pita, J., Kiekintveld, C., Rathi, S., Tambe, M., Ordonez, F.: Software Assistants for Randomized Patrol Planning for the LAX Airport Police and the Federal Air Marshal Service. Interfaces 40, 267–290 (2010)
15. Jiang, A.X., Nguyen, T.H., Tambe, M., Procaccia, A.D.: Monotonic maximin: A robust stackelberg solution against boundedly rational followers. In: Conference on Decision and Game Theory for Security (GameSec) (2013)
16. Jiang, A., Yin, Z., Kraus, S., Zhang, C., Tambe, M.: Game-theoretic randomization for security patrolling with dynamic execution uncertainty. In: AAMAS (2013)
17. Johnson, M., Fang, F., Yang, R., Tambe, M., Albers, H.: Patrolling to Maximize Pristine Forest Area. In: Proc. of the AAAI Spring Symposium on Game Theory for Security, Sustainability and Health (2012)
18. Kahneman, D., Tversky, A.: Prospect Theory: An Analysis of Decision Under Risk. Econometrica 47(2), 263–291 (1979)
19. Kar, D., Fang, F., Fave, F.D., Sintov, N., Tambe, M.: a game of thrones: When human behavior models compete in repeated stackelberg security games. In: International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015) (2015)
20. Keteyian, A.: TSA: Federal Air Marshals (2010), <http://www.cbsnews.com/stories/2010/02/01/earlyshow/main6162291.shtml>, *retrieved Feb 1, 2011*
21. Kiekintveld, C., Jain, M., Tsai, J., Pita, J., Tambe, M., Ordonez, F.: Computing Optimal Randomized Resource Allocations for Massive Security Games. In: Proc. of The 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS). pp. 689–696 (2009)

22. Korzhyk, D., Conitzer, V., Parr, R.: Complexity of computing optimal stackelberg strategies in security resource allocation games. In: Proc. of The 24th AAAI Conference on Artificial Intelligence. pp. 805–810 (2010)
23. Leitmann, G.: On Generalized Stackelberg Strategies. *Optimization Theory and Applications* 26(4), 637–643 (1978)
24. Lipton, R., Markakis, E., Mehta, A.: Playing large games using simple strategies. In: EC: Proceedings of the ACM Conference on Electronic Commerce. pp. 36–41. ACM New York, NY, USA (2003)
25. McFadden, D.: Conditional logit analysis of qualitative choice behavior. Tech. rep. (1972)
26. McFadden, D.: Quantal choice analysis: A survey. *Annals of Economic and Social Measurement* 5(4), 363–390 (1976)
27. McKelvey, R.D., Palfrey, T.R.: Quantal Response Equilibria for Normal Form Games. *Games and Economic Behavior* 10(1), 6–38 (1995)
28. Nguyen, T., Jiang, A., Tambe, M.: Stop the compartmentalization: Unied robust algorithms for handling uncertainties in security games. In: International Conference on Autonomous Agents and Multiagent Systems (AAMAS) (2014)
29. Nguyen, T.H., Yang, R., Azaria, A., Kraus, S., Tambe, M.: Analyzing the effectiveness of adversary modeling in security games. In: Conference on Artificial Intelligence (AAAI) (2013)
30. Paruchuri, P., Pearce, J.P., Marecki, J., Tambe, M., Ordonez, F., Kraus, S.: Playing Games with Security: An Efficient Exact Algorithm for Bayesian Stackelberg Games. In: Proc. of The 7th International Conference on Autonomous Agents and Multiagent Systems (AAMAS). pp. 895–902 (2008)
31. Pita, J., Bellamane, H., Jain, M., Kiekintveld, C., Tsai, J., Ordez, F., Tambe, M.: Security applications: lessons of real-world deployment. *ACM SIGecom Exchanges* 8(2) (2009)
32. Pita, J., Jain, M., Ordez, F., Tambe, M., Kraus, S., Magori-Cohen, R.: Effective solutions for real-world stackelberg games: When agents must deal with human uncertainties. In: The Eighth International Conference on Autonomous Agents and Multiagent Systems (2009)
33. Pita, J., Jain, M., Western, C., Portway, C., Tambe, M., Ordonez, F., Kraus, S., Parachuri, P.: Deployed ARMOR protection: The Application of a Game-Theoretic Model for Security at the Los Angeles International Airport. In: Proc. of The 7th International Conference on Autonomous Agents and Multiagent Systems (AAMAS). pp. 125–132 (2008)
34. Pita, J., John, R., Maheswaran, R., Tambe, M., Kraus, S.: A robust approach to addressing human adversaries in security games. In: European Conference on Artificial Intelligence (ECAI) (2012)
35. Shieh, E., An, B., Yang, R., Tambe, M., Baldwin, C., DiRenzo, J., Maule, B., Meyer, G.: PROTECT: A Deployed Game Theoretic System to Protect the Ports of the United States. In: Proc. of The 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS) (2012)
36. von Stackelberg, H.: *Marktform und Gleichgewicht*. Springer, Vienna (1934)
37. von Stengel, B., Zamir, S.: Leadership with Commitment to Mixed Strategies. Tech. Rep. LSE-CDAM-2004-01, CDAM Research Report (2004)
38. Yang, R., Ford, B., Tambe, M., Lemieux, A.: Adaptive resource allocation for wildlife protection against illegal poachers. In: International Conference on Autonomous Agents and Multiagent Systems (AAMAS) (2014)
39. Yang, R., Jiang, A.X., Tambe, M., Ordóñez, F.: Scaling-up security games with boundedly rational adversaries: a cutting-plane approach. In: Proceedings of the Twenty-Third international joint conference on Artificial Intelligence. pp. 404–410. AAAI Press (2013)
40. Yang, R., Kiekintveld, C., Ordonez, F., Tambe, M., John, R.: Improving Resource Allocation Strategy Against Human Adversaries in Security Games. In: IJCAI (2011)

41. Yang, R., Ordonez, F., Tambe, M.: Computing optimal strategy against quantal response in security games. In: Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 2. pp. 847–854. International Foundation for Autonomous Agents and Multiagent Systems (2012)
42. Yin, Z., Jain, M., Tambe, M., Ordonez, F.: Risk-Averse Strategies for Security Games with Execution and Observational Uncertainty. In: Proc. of The 25th AAAI Conference on Artificial Intelligence (AAAI). pp. 758–763 (2011)
43. Yin, Z., Jiang, A., Johnson, M., Tambe, M., Kiekintveld, C., Leyton-Brown, K., Sandholm, T., Sullivan, J.: TRUSTS: Scheduling Randomized Patrols for Fare Inspection in Transit Systems. In: Proc. of The 24th Conference on Innovative Applications of Artificial Intelligence (IAAI) (2012)