

Deployed Security Games for Patrol Planning

Fernando Ordóñez, Milind Tambe, Juan F. Jara, Manish Jain, Christopher Kiekintveld, Jason Tsai

Abstract Nations and organizations need to secure locations of economic, military, or political importance from groups or individuals that can cause harm. The fact that there are limited security resources prevents complete security coverage, which allows adversaries to observe and exploit patterns in patrolling or monitoring, and enables them to plan attacks that avoid existing patrols. The use of randomized security policies that are more difficult for adversaries to predict and exploit can counter their surveillance capabilities and improve security. In this chapter we describe the recent development of models to assist security forces in randomizing their patrols and their deployment in real applications.

The systems deployed are based on fast algorithms for solving large instances of Bayesian Stackelberg games that capture the interaction between security forces and adversaries. Here we describe a generic mathematical formulation of these models, present some of the results that have allowed these systems to be deployed in practice, and outline remaining future challenges. We discuss the deployment of these systems in two real-world security applications: 1) The police at the Los Angeles International Airport uses these models to randomize the placement of checkpoints on roads entering the airport and the routes of canine unit patrols within the airport terminals. 2) The Federal Air Marshal Service uses these models to randomize the schedules of air marshals on international flights.

Fernando Ordóñez and Juan F. Jara
Industrial Engineering Department, University of Chile, Republica 701, Santiago, Chile, e-mail: fordon@dii.uchile.cl

Milind Tambe, Manish Jain and Jason Tsai
Computer Science Department, University of Southern California, Los Angeles, 90089 e-mail: milind,manishja,jason@usc.edu

Christopher Kiekintveld
Computer Science Department, University of Texas, El Paso, 79968 e-mail: cdkiekintveld@utep.edu

1 Introduction

Nations and organizations need to secure locations of economic, military, or political importance from groups or individuals that can cause harm. Protecting such critical sites and targets, such as airports, historical landmarks, power generation facilities, and political figures, is a challenging task for police and security agencies worldwide. The growing threat of international terrorism has exacerbated this challenge in recent years. For instance, transportation networks such as buses, trains, and airplanes carry millions of people per day to their destinations, making them a prime target for terrorists and extremely difficult to protect for law enforcement agencies. The September 11, 2001 attack on the World Trade Center in New York City via commercial airliners resulted in \$27.2 billion of direct short term costs [30] as well as a government-reported 2,974 lives lost. The 2004 Madrid commuter train bombings resulted in 191 lives lost, 1,755 wounded, and an estimated cost of 212 million Euros [11]. Finally, in the 2005 London subway and bus bombings, 52 lives were lost, 700 were wounded, and there was an estimated economic cost of 2 billion pounds [43].

Measures for protecting potential target areas include monitoring entrances or in-bound roads and patrolling the network at transfer points and aboard transportation vehicles. However, limited resources imply that it is typically impossible to provide full security coverage at all times. Furthermore, adversaries can observe security arrangements over time and exploit any predictable patterns to their advantage. One way to mitigate the ability of adversaries to exploit patterns is the judicious use of randomization in scheduling the actions of security forces. For example, police patrols, baggage screenings, vehicle checkpoints, and other security procedures are often randomized. However, security forces face many difficulties in effectively randomizing their operations. One of these difficulties is how to weigh the different actions the defender could take. A strategy in which all targets are equally likely to be defended fails to take into account that some targets are more attractive or vulnerable than others. A defense strategy that weighs the protection of each target against the value of that target still fails to account for the possibility that the attacker is intelligent and will update their strategy based on the actions of the defender. Asking a human to generate a random security policy has additional difficulties as humans are not good at generating truly random behavior [48, 44], and can easily fall into predictable patterns. Furthermore, in transportation networks and many other security domains, the problem of scheduling security forces is prohibitively large, even without considering randomization. Creating a schedule by hand is a costly and labor-intensive process.

Our work on randomized patrol planning has led to a number of deployed software assistants that address many of these key difficulties of randomization and provide an easy-to-use solution for security forces. These assistants use game-theoretic models and solution algorithms to determine good randomization strategies that take into account target values and assume intelligent adversary responses to security measures. Game theory is a well-established paradigm for reasoning about situations with multiple self-interested decision makers [20]. We model security games

as Stackelberg games [46] between the defender (i.e., the security forces), and the attacker (i.e., a terrorist adversary). Stackelberg games are a bilevel model [8] that account for the ability of an attacker to gather information about the defense strategy before planning an attack. These games specify different payoff values for both players in the event of an attack on every potential target. Extending these games to Bayesian Stackelberg games [16] allows us to capture uncertainty about these payoffs in the game model. Solutions to these games provide a randomized policy for the defense strategy, which can be used to generate specific schedules for security patrols.

In this article we describe how we applied this game-theoretic approach in two different software solutions that provide assistance in scheduling real security operations. The ARMOR program [37] was developed for the Los Angeles airport (LAX) police, and randomizes checkpoints on the roadways entering the airport and canine patrol routes within the airport terminals. The IRIS program [45] was developed for the Federal Air Marshal Service (FAMS) to assist with randomly scheduling air marshals on flights. These software assistants are interactive and domain experts can change domain parameters when necessary. Underlying each of these tools is a model of the domain as a Bayesian Stackelberg game, along with fast solution algorithms for computing an optimal solution to the game model. These algorithms use various techniques for exploiting structure in the security domains to speed up the computation and enable large real-world problem instances to be solved in reasonable amounts of time [36, 35, 26]. As we highlight later in this article, developing these assistants requires a substantial amount of work in calibrating the Stackelberg game model to capture expert's knowledge of the security domain. This is critical so that the defender strategies proposed are reasonable and useful. Clustering and data mining methods can help in formulating a representative security game in situations where there is sufficient information of events.

The rest of the article is organized as follows. Related work is discussed in Section 2. The Bayesian Stackelberg security game models, technical formulation and solution algorithms are discussed in Section 3. The LAX and FAMS domains and the software assistants developed are described in Section 4. In Section 5, we illustrate how to use clustering methods to automatically build a Stackelberg security game for a network patrolling problem. We present our conclusions in Section 6. This book chapter is based on our previous work [35, 26, 23], and extends it by describing a data driven process to build a Stackelberg security game.

2 Related Work

There are three main areas of related work that we review here: Optimization techniques for patrol planning that do not take the strategic behavior of adversaries into account, Stackelberg game models used in diverse security problems, and other game-theoretic models used for security.

The first area of related work applies optimization techniques to model a security domain, but do not address the strategic aspects of the problem. These methods provide a randomization strategy for the defender, but they do not take into account the fact that the adversaries can observe the defender's actions and then adjust their behavior. Examples of such authors or approaches include [39, 36] which are based on learning, Markov Decision Processes (MDPs) and Partially Observable Markov Decision Processes (POMDPs). As part of this work, the authors model the patrolling problem with locations and varying incident rates in each of the locations and solve for optimal routes using a MDP framework. Another example is the "Hypercube Queueing Model" [28] which is based on queueing theory and depicts the detailed spatial operation of urban police departments and emergency medical services. It has found application in police beat design, in allocation of patrolling time, etc. Such frameworks can address many of the problems we raise, including different target values and increasing uncertainty by using many possible patrol routes. However, they fail to account for the possibility that an intelligent attacker will observe and exploit patterns in the security policy. If a policy is based on the historical frequency of attacks, it is essentially a reactive policy and an intelligent attacker will always be one step ahead.

A second area of work uses Stackelberg games to model a variety of security domains. Bier et al. [10] give a strong endorsement of this type of modeling for security problems. Game-theoretic models have been applied in a variety of homeland security settings, such as protecting critical infrastructure [15, 37, 32]. Lawrence et al. [50] apply Stackelberg games in the context of screening visitors entering the United States. In their work, they model the U.S. Government as the leader who specifies the biometric identification strategy to maximize the detection probability using finger print matches, and the follower is the terrorist who can manipulate the image quality of the finger print. They have also been used for studying missile defense systems [13] and for studying the development of an adversary's weapon systems [14]. A family of Stackelberg games known as inspection games is closely related to the security games we are interested in and includes models of arms inspections and border patrols [6]. Other recent work uses Stackelberg games to obtain randomized patrolling in a generic "police and robbers" scenario [21] and perimeter patrols [4].

Our work belongs to this line of research, focusing on Stackelberg games for patrol planning. Our work differs from the previous work mainly in the solution approach used and the domain constraints considered, which have arisen from our work deploying these systems in the real world. In addition to the ARMOR and IRIS systems that will be discussed in detail in this chapter, we are currently working on designing new game-theoretic scheduling assistants for other security agencies. For instance, GUARDS [38] is a system for scheduling activities being developed for the Transportation Security Administration. GUARDS is being evaluated at an undisclosed airport for potential nationwide deployment. Finally, PROTECT [5] is in use for scheduling the patrols of the United States Coast Guard in the port of Boston; it is currently being deployed in the port of New York and may be deployed at multiple other ports in the United States.

The third area of related work is the application of game-theoretic techniques that are not based on Stackelberg games to security applications. Security problems are increasingly studied using game-theoretic analysis, ranging from computer network security [49, 41] to terrorism [40]. Babu et al. [7] have worked on modeling passenger security system at US airports using linear programming approaches, however, their objective is to classify the passengers in various groups and then screen them based on the group they belong to.

3 Methodology

A generic Stackelberg game has two players, a *leader* and a *follower*. These players need not represent individuals, but could also be groups that cooperate to execute a joint strategy, such as a police force or terrorist organization. For the modeling of security, the leadership role is assumed by the police, and the role of follower by criminals. Decisions made by each player are where to protect and where to attack respectively. Thus, having police act first reflects the fact that patrols conducted by police officers are observable by criminals and, in the long run, the latter are able to estimate the probability of encountering police in a given sector. Thus, the decision of offenders is carried out once the likelihood of facing the police is observed.

The actions for the security forces represent the action of scheduling a patrol or security procedure to protect a set of targets, e.g. a checkpoint at the LAX airport or assigning federal air marshals to a flight. The actions for an adversary represent possible attacks at one of the targets being protected, e.g. a terminal at LAX or a certain flight.

3.1 Stackelberg Equilibrium

In a Stackelberg game each player has a set of possible *pure strategies*, denoted $\sigma_d \in \Sigma_d$ and $\sigma_a \in \Sigma_a$. A *mixed strategy* allows a player to play a probability distribution over pure strategies, denoted $\delta_d \in \Delta_d$ and $\delta_a \in \Delta_a$. Payoffs for each player are defined over all possible joint pure-strategy outcomes: $\Omega_d : \Sigma_a \times \Sigma_d \rightarrow \mathcal{R}$ for the defender and similarly for each attacker. The payoff functions are extended to mixed strategies in the standard way by taking the expectation over pure-strategy outcomes. The follower can observe the leader's strategy, and then act in a way to optimize its own payoffs. Formally, the attacker's strategy in a Stackelberg security game becomes a function that selects a strategy for each possible leader strategy: $F_a : \Delta_d \rightarrow \Delta_a$.

The most common solution concept in game theory is a *Nash equilibrium*, which is a profile of strategies for each player in which no player can gain by unilaterally changing to another strategy [33]. Stackelberg equilibrium is a refinement of Nash equilibrium specific to Stackelberg games. It is a form of sub-game perfect equi-

librium in that it requires that each player select the best-response in any subgame of the original game (where subgames correspond to partial sequences of actions). The effect is to eliminate equilibrium profiles that are supported by non-credible threats off the equilibrium path. Subgame perfection is a natural requirement, but it does not guarantee a unique solution in cases where the follower is indifferent among a set of strategies. The literature contains two forms of Stackelberg equilibria that identify unique outcomes, first proposed by Leitmann [29], and typically called “strong” and “weak” after Breton et. al. [12]. The strong form assumes that the follower will always choose the optimal strategy for the leader in cases of indifference, while the weak form assumes that the follower will choose the worst strategy for the leader. Unlike the weak form, strong Stackelberg equilibria are known to exist in all Stackelberg games [9]. A standard argument suggests that the leader is often able to induce the favorable strong form by selecting a strategy arbitrarily close to the equilibrium which causes the follower to strictly prefer the desired strategy [47]. We adopt strong Stackelberg equilibrium as our solution concept in part for these reasons, but also because it is the most commonly used in related literature [33, 16, 34].

Definition 1 *A set of strategies (δ_d, F_a) form a Strong Stackelberg Equilibrium (SSE) if they satisfy the following:*

1. *The leader plays a best-response:*

$$\Omega_d(\delta_d, F_a(\delta_d)) \geq \Omega_d(\delta'_d, F_a(\delta'_d)) \forall \delta'_d \in \Delta_d.$$
2. *The follower plays a best-response:*

$$\Omega_a(\delta_d, F_a(\delta_d)) \geq \Omega_a(\delta_d, \delta_a) \forall \delta_d \in \Delta_d, \delta_a \in \Delta_a.$$
3. *The follower breaks ties optimally for the leader:*

$$\Omega_d(\delta_d, F_d(\delta_d)) \geq \Omega_d(\delta_d, \delta_a) \forall \delta_d \in \Delta_d, \delta_a \in \Delta_a^*(\delta_d),$$
where $\Delta_a^(\delta_d)$ is the set of follower best-responses, as above.*

Whether or not the Stackelberg leader benefits from the ability to commit depends on whether commitment to mixed strategies is allowed. Committing to a pure strategy can be either good or bad for the leader; for example, in the “Rock, Paper, and Scissors” game, forcing commitment to a pure strategy would guarantee a loss. However, it has been shown that the ability to commit to a mixed strategy always weakly increases the leader’s payoffs in equilibrium profiles of the game [47]. In the context of a Stackelberg security game, a deterministic policy is a liability for the defender (the leader), but a credible randomized security policy is an advantage. Our model allows commitment to mixed strategies by the defender.

The Bayesian extension to the Stackelberg game allows for multiple types of players, with each type associated with its own payoff values. For the security games of interest in this chapter, we assume that there is only one leader type (e.g. only one police force), although there are multiple follower types (e.g. multiple adversary types trying to infiltrate security). The set of follower types is denoted by Γ . Each type γ is represented by a different payoff matrix. The leader does not know the follower’s type. The goal is to *find the optimal mixed strategy* for the leader to commit to, given that each follower type will know the mixed strategy of the leader when choosing its own strategy. Payoffs for each type are defined over all possible

joint pure-strategy outcomes: $\Omega_d : \Sigma_a^I \times \Sigma_d \rightarrow \mathcal{R}$ for the defender and similarly for each attacker type. The leader’s best response is now a weighted best response to the followers’ responses, where the weights are based on the probability of occurrence of each type. The strategy of each attacker type γ becomes: $F_a^\gamma : \Delta_d \rightarrow \Delta_a^\gamma$, which still satisfies constraints 2 and 3 in Definition 1.

3.2 Security Game Representation

There are two major problems with using conventional methods to represent security games in normal form. First, many solution methods require the use of a Harsanyi transformation when dealing with Bayesian games [22]. The Harsanyi transformation converts a Bayesian game into a normal-form game, but the new game may be exponentially larger than the original Bayesian game. Our compact representation avoids this Harsanyi transformation, and instead we directly operate on the Bayesian game. Operating directly on the Bayesian representation is possible in our model because the evaluation of the leader strategy against a Harsanyi-transformed game matrix is equivalent to its evaluation against each of the game matrices for the individual follower types. (For more details, see the Appendix; a further detailed explanation appears in [34]). The second problem arises because the defender has many possible resources to schedule in the security policy. This can also lead to a combinatorial explosion in a standard normal-form representation. For example, if the leader has m resources to defend n entities, then normal-form representations model this problem as a single leader with $\binom{n}{m}$ rows, each row corresponding to a leader action of covering m targets with security resources. However, in our compact representation, the game representation would only include n rows, each row corresponding to whether the corresponding target was covered or not. Such a representation is equivalent to the normal form representation for the class of problems we address in this work (see [26] for additional details). This compactness in our representation is possible because the payoffs for the leader in these games simply depend on whether the attacked target was covered or not, and not on what other targets were covered (or not covered). The representation we use here avoids both of these potential problems, using methods similar to other compact representations for games [27, 24].

We now introduce our compact representation for security games. Let $T = \{t_1, \dots, t_n\}$ be a set of *targets* that may be attacked, corresponding to pure strategies for the attacker. The defender has a set of resources available to *cover* these targets, $R = \{r_1, \dots, r_m\}$ (for example, in the FAMS domain, targets could be flights and resources could be federal air marshals). Associated with each target are four payoffs defining the possible outcomes for an attack on the target, as shown in Table 1. There are two cases, depending on whether or not the target is covered by the defender. The defender’s payoff for an uncovered attack when facing an adversary of type γ is denoted $U_d^{\gamma,u}(t)$, and for a covered attack $U_d^{\gamma,c}(t)$. Similarly, $U_a^{\gamma,u}(t)$ and $U_a^{\gamma,c}(t)$ are the payoffs of the attacker.

| | Covered | Uncovered |
|----------|---------|-----------|
| Defender | 5 | -20 |
| Attacker | -10 | 30 |

Table 1 Example payoffs for an attack on a target.

A crucial feature of the model is that payoffs depend only on the target attacked, and whether or not it is covered by the defender. The payoffs do *not* depend on the remaining aspects of the schedule, such as whether any unattacked target is covered or which specific defense resource provides coverage. For example, if an adversary succeeds in attacking Terminal 1, the penalty for the defender is the same whether the defender was guarding Terminal 2 or 3. Therefore, from a payoff perspective, many resource allocations by the defender are identical. We exploit this by summarizing the payoff-relevant aspects of the defender's strategy in a *coverage* vector, C , that gives the probability that each target is covered, c_t . The analogous attack vector A^γ gives the probability of attacking a target by a follower of type γ . We restrict the attack vector for each follower type to attack a single target with probability 1. This is without loss of generality because a strong Stackelberg equilibrium (SSE) solution still exists under this restriction [34]. Thus, the follower of type γ can choose any pure strategy $\sigma_a^\gamma \in \Sigma_a^\gamma$, that is, attack any one target from the set of targets.

The payoff for a defender when a specific target t is attacked by an adversary of type γ is given by $U_d^\gamma(t, C)$ and is defined in Equation 1. Thus, the expectation of $U_d^\gamma(t, C)$ over t gives U_d^γ , which is the defender's expected payoff given coverage vector C when facing an adversary of type γ whose attack vector is A^γ . U_d^γ is defined in Equation 2. The same notation applies for each follower type, replacing d with a . Thus, $U_a^\gamma(t, C)$ gives the payoff to the attacker when a target t is attacked by an adversary of type γ . We will see $U_a^\gamma(t, C)$ and $U_d^\gamma(t, C)$ used in the MILP discussed later. We also define the useful notion of the *attack set* in Equation 3, $\Lambda^\gamma(C)$, which contains all targets that yield the maximum expected payoff for the attacker type γ given coverage C . This *attack set* is used by the adversary to break ties when calculating a strong Stackelberg equilibrium. Moreover, in these security games, exactly one adversary is attacking in one instance of the game; however, the adversary could be of any type and the defender does not know the type of the adversary faced.

$$U_d^\gamma(t, C) = c_t U_d^{\gamma,c}(t) + (1 - c_t) U_d^{\gamma,u}(t) \quad (1)$$

$$U_d^\gamma(C, A^\gamma) = \sum_{t \in T} a_t^\gamma \cdot (c_t \cdot U_d^{\gamma,c}(t) + (1 - c_t) U_d^{\gamma,u}(t)) \quad (2)$$

$$\Lambda^\gamma(C) = \{t : U_d^\gamma(t, C) \geq U_d^\gamma(t', C) \forall t' \in T\}. \quad (3)$$

In a strong Stackelberg equilibrium, the attacker selects the target in the attack set with maximum payoff for the defender. Let t^* denote this optimal target. Then the expected SSE payoff for the defender when facing this adversary of type γ with probability p^γ is $\hat{U}_d^\gamma(C) = U_d^\gamma(t^*, C) \times p^\gamma$, and for the attacker $\hat{U}_a^\gamma(C) = U_a^\gamma(t^*, C)$.

3.3 Solution Method

We introduce the ERASER-C algorithm (Efficient Randomized Allocation of Security Resources with Constraints), which takes as input a security game in the compact form described in Section 3.2 and solves for an optimal coverage vector corresponding to a SSE strategy for the defender. We allow resources to be assigned to *schedules* covering multiple targets. The set of legal schedules $S = \{s_1 \dots s_l\}$ is a subset of the power set of the targets, with restrictions on this set representing scheduling constraints. We define the relationship between targets and schedules with the function $H : S \times T \rightarrow \{0, 1\}$, which evaluates to 1 if and only if t is covered in s . The defender’s strategy is now an assignment of resources to schedules, rather than targets. Another important notion is the presence of *resource types*, $\Omega = \{\omega_1, \dots, \omega_v\}$, each with the capability to cover a different subset of S . The number of available resources of each type is given by the function $\mathcal{R}(\omega)$. Coverage capabilities for each type are given by the function $Ca : S \times \Omega \rightarrow \{0, 1\}$, which is 1 if the type is able to cover the given schedule and 0 otherwise.¹

The combination of schedules and resource types captures key elements of the security domains. For example, in FAMS, federal air marshals are resources, and flights are potential targets, with payoff values defined by risk analysis of the flight. Due to location and timing constraints, however, a marshal cannot be on all possible flights. For example, a marshal in New York cannot board flights flying out of Los Angeles. Legal schedules can be used to define the set of possible flights that a federal air marshal could fly, given these constraints. Resource types are used to define the initial state (notably, location) of a marshal, which defines a subset of legal schedules that any given marshal could fly.

Adding scheduling and resource coverage constraints reduces the space of feasible coverage vectors. Consider an example with a single federal air marshal defending three flights. Suppose that there are two legal schedules, covering targets $\{1, 2\}$ and $\{2, 3\}$. Given only these schedules, it is not possible to implement a coverage vector that places 50% probability on both targets 1 and 3, with no coverage of target 2.

The algorithm is a mixed-integer linear program (MILP) described in (4)–(11), with notation presented in Table 2. Constraints (5) and (12) force each adversary to select a pure strategy attacking a single target. The coverage vector C is constrained by the number of available resources through (8) and the coverage in each target to be in the range $[0, 1]$ by (13). The coverage of each schedule must sum to the contributions of the individual resource types, specified by constraint (6). The mapping between the coverage of schedules and coverage of targets is enforced in Equation (7). Constraint (8) restricts the schedule so that only the available number of resources of each type are used. Constraint (9) enforces that no probability may be assigned infeasible schedules for each resource type. The defender’s expected payoff is defined with constraint (10) when follower γ attacks target A^γ . Since the

¹ Our current implementation uses complete matrices to represent H and Ca , but sparse representations could offer additional performance improvements.

objective maximizes d^γ , for any optimal solution $d^\gamma = U_d^\gamma(C, A^\gamma)$. This also implies that C is maximal, given A^γ for any optimal solution, since d^γ is maximized. In a similar way, constraint (11) forces the attacker to select a strategy in the attack set of C . If the attack vector specifies a target that is not maximal, this constraint is violated. Therefore, taken together, the objective and constraints (10)–(11) imply that C and A^γ are mutual best-responses for the defender and the adversary in any solution. Thus, the defender mixed strategy C and the adversary attack vector A^γ for each adversary type γ form a strong Stackelberg equilibrium of the security Stackelberg game.

| Symbol | Meaning |
|-----------------------|--|
| d^γ | Reward of defender against adversary of type γ |
| k^γ | Reward of adversary type γ |
| p^γ | Probability of occurrence of adversary of type γ |
| Γ | Set of adversary types |
| T | Set of targets |
| A^γ | Attack vector for the adversary of type γ |
| a_t^γ | Probability of adversary of type γ attacking target t |
| C | Coverage vector of the defender |
| c_t | Probability of defender covering target t |
| $h(s, \omega)$ | Probability of coverage of schedule s by defender type ω |
| x_s | Total coverage probability over schedule s |
| S | Set of valid schedules |
| Ω | Set of resource types |
| $Ca(s, \omega)$ | Capability: 1 if type ω can cover schedule s ; 0 otherwise |
| $\mathcal{R}(\omega)$ | Number of available resources of type ω |
| $H(s, t)$ | Mapping: 1 if schedule s covers target t ; 0 otherwise |
| M | Huge positive constant |
| $U_d^\gamma(t, C)$ | Utility of the defender when facing adversary type γ who attacks target t when defender coverage is C |
| $U_a^\gamma(t, C)$ | Utility of the adversary of type γ when target t is attacked and defender coverage is C |

Table 2 Notation Table

$$\max_{a,c,q,h,d,k} \sum_{\gamma \in \Gamma} d^\gamma p^\gamma \quad (4)$$

$$\text{s.t.} \quad \sum_{t \in T} a_t^\gamma = 1 \quad \gamma \in \Gamma \quad (5)$$

$$\sum_{\omega \in \Omega} h_{s,\omega} = x_s \quad s \in S \quad (6)$$

$$\sum_{s \in S} x_s H(s, t) = c_t \quad t \in T \quad (7)$$

$$\sum_{s \in \mathcal{S}} h_{s,\omega} Ca(s, \omega) \leq \mathcal{R}(\omega) \quad \omega \in \Omega \quad (8)$$

$$h_{s,\omega} \leq Ca(s, \omega) \quad s, \omega \in \mathcal{S} \times \Omega \quad (9)$$

$$d^\gamma - U_d^\gamma(t, C) \leq (1 - a_t^\gamma) \cdot M \quad t \in T, \gamma \in \Gamma \quad (10)$$

$$0 \leq k^\gamma - U_a^\gamma(t, C) \leq (1 - a_t^\gamma) \cdot M \quad t \in T, \gamma \in \Gamma \quad (11)$$

$$a_t^\gamma \in \{0, 1\} \quad t \in T, \gamma \in \Gamma \quad (12)$$

$$c_t \in [0, 1] \quad t \in T \quad (13)$$

$$x_s \in [0, 1] \quad s \in \mathcal{S} \quad (14)$$

$$h_{s,\omega} \in [0, 1] \quad s, \omega \in \mathcal{S} \times \Omega \quad (15)$$

The payoff values $U_d^\gamma(t, C)$ and $U_a^\gamma(t, C)$ are calculated based on Equations 1 and 2. The values of $U_d^{\gamma,c}$ and $U_d^{\gamma,u}$ used in these equations are the payoff values to the defender when a target is covered and uncovered respectively. These values are provided by the domain experts, as described in Section 5. Similarly, the payoff values for the adversaries are also provided by the domain experts.

The values of other model parameters are calculated based on the user input and the game specification. Police officers and canines are the resources for ARMOR for checkpoint and ARMOR for canine respectively. ARMOR does not differentiate between different resources (for example, all canines are assumed to be equally capable), and hence there is exactly one resource type Ω . The number of resources \mathcal{R} , i.e. checkpoints or canines, is directly input by the user in the system. In the case of ARMOR, the set of legal schedules is an assignment of a checkpoint to an inbound road, and is automatically generated by the system since ARMOR is aware of the road map of the airport. The capability matrix Ca in ARMOR consists of all ones since any resource could be assigned to any target. For example, any canine could be scheduled to any terminal.

Similarly, all the model parameters are defined based on user input and domain constraints in IRIS. The federal air marshals are the resources for IRIS. In IRIS, the different FAM Offices form the different resource types. This information has already been supplied to IRIS by the domain experts. The numbers of resources of each type \mathcal{R} , that is the number of federal air marshals in each office, is directly input in IRIS by the end users. The set of legal schedules \mathcal{S} is provided as an input to the system by the FAMS in IRIS. Each schedule in IRIS is a sequence of flights that a federal air marshal can take to complete a tour. In IRIS, the capability matrix Ca is defined based on resource types; for example, federal air marshals at the FAM office based in Los Angeles can only cover schedules flying out of Los Angeles, and hence only those schedules would have their capabilities set to 1. The mapping M is also calculated by the systems based on the domain specifications. For example, in IRIS, if schedule s is to take flight f1 followed by flight f2, then the row in M corresponding to s would have ones only for columns corresponding to f1 and f2.

Kiekintveld et al. [26] have shown that the ERASER-C MILP corresponds to a SSE of the security game. The intuition behind the proof are two claims: (1) the coverage probability of the leader and the attack set of the follower are mutual best-

responses by the construction of the MILP, and (2) the coverage probability of the leader gives the leader the optimal utility.

4 Software Systems Deployed at the LAX and FAMS Domains

Both LAX and FAMS are security scenarios in which there is a leader/follower dynamic between the security forces and terrorist adversaries. In both domains there are limited resources available to protect a very large space of possible targets, so it is not possible to provide complete coverage. Finally, the targets have diverse values and vulnerabilities in each domain. The domains, however, differ primarily due to size. In the LAX security domain there are eight terminals that must be protected, while the air marshals are responsible for protecting tens of thousands of commercial flights each day. This difference in size requires, in addition to scalable solution algorithms, different types of interfaces to have domain experts specify each game. Finally, while in the LAX domain all security resources can reach all targets, in the FAMS domain, the security resources must satisfy more complicated constraints (for example, a given marshal cannot be assigned to two flights with overlapping time schedules).

In this section, we describe both security domains (LAX and FAMS) and discuss the architecture of the software systems developed for these domains. We begin with a description of the generic software architecture and then describe each domain and their specific software assistant. We finish this section with a list of lessons learned in doing these deployments.

4.1 Software Assistants

We now describe in detail the system architecture for each of the two software assistants, focusing primarily on the ARMOR system but providing some discussion of IRIS as a point of comparison. We paid particular attention to organization acceptance during the development process. The end users of both ARMOR and IRIS are security officers, and the system must be simple enough for them to be comfortable using it on a regular basis. In particular, the systems are designed to hide as much of the complexity of the game-theoretic models as possible, while still allowing enough flexibility for the users to input important parameters that change regularly. This required considerable effort in both user interface design and identifying ways to simplify and reduce the inputs required by the system to specify a game model. In the case of IRIS, it was also very important to build in functionality to import data from other systems to ease the burden of data entry (e.g., importing flight information from existing databases). Finally, the schedules that the system produces must be presented in a format that is easy to understand, with tools that allow final modifications if necessary.

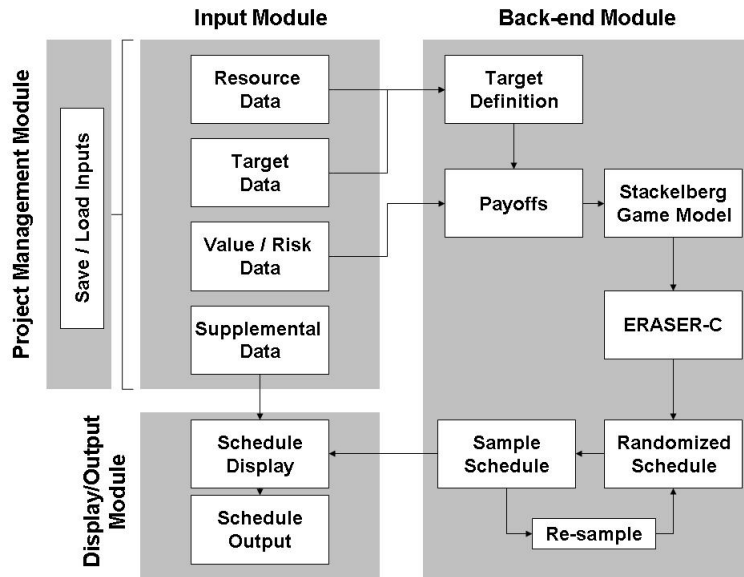


Fig. 1 General Structure for the Security Assistants

Both ARMOR and IRIS are stand-alone desktop applications. ARMOR was developed in the Microsoft .NET framework, while IRIS is a standalone Java application. Due to security concerns, both systems are run on machines that are not connected to any network. The underlying solution methods use the open source glpk² toolkit to solve the necessary mixed-integer programs. The general structure of the two applications is shown in Figure 1. The core architecture can be divided into three modules, which we describe in detail in the subsequent sections:

1. **Input:** Interface for the user to enter parameters and domain knowledge.
2. **Back-end:** Inputs are translated into a game model, which is passed to the Bayesian Stackelberg game solver and then to a final process that generates a specific sample schedule based on the computed probabilities.
3. **Display Module:** The final schedule is presented to the user, with options to modify the output if necessary.

We rely on the users and domain experts to provide the knowledge required to specify the game model. While some elements of the model do not change over time, others change frequently. For these, we must provide the users a convenient way to enter the necessary values. The basic inputs that both ARMOR and IRIS require fall into four categories: (1) the number of available resources and their capabilities, (2) the set of targets, (3) payoff values for each target, and (4) supplemental data to improve the user experience (e.g., names and labels). Both applications allow users to save and re-use this information across multiple executions.

² <http://www.gnu.org/software/glpk/>

The balance of how much information is hard-coded and how much is entered by the user is quite different for ARMOR and IRIS. For example, in ARMOR the set of targets is hard-coded because the number of terminals at LAX changes very rarely. However, in IRIS the flight information may change every time the system is run, so this is part of the user input. Determining which parameters were necessary to expose to the user was a significant task, and required several iterations with the domain experts and end users to strike the right balance between the complexity of the inputs and the flexibility of the system to capture the necessary information.

The Back-end module is fairly common to the two applications. This model builds a specific instance of a Bayesian Stackelberg game, based on all of the data provided by domain experts and entered through the GUI by end users. Some of the necessary information is hard-coded in each system, while other inputs can be modified by the user during the scheduling process.

Once an explicit game model has been generated it is passed as input to the ERASER-C mixed-integer program. This model is solved using the standard open source solver GLPK in these applications. ERASER-C returns an optimal mixed strategy for the defender — a probability distribution over the defender’s actions — which represents a randomized policy for allocating the security resources of either LAX or FAMS. We sample the randomized schedule found to generate a specific schedule for the security forces. This sample schedule specifies exactly where and when each resource should be assigned to each target. If necessary, it is also possible to “re-sample” from the randomized schedule to get another specific schedule, though this capability is used rarely. Any specific constraints that the schedules must satisfy, are taken into consideration when final schedules are sampled. These sampled schedules are then displayed for the user through the Display Module.

The output module presents the generated sampled schedule to the user. The user can then review the schedule and accept it as is, or add additional constraints and run the scheduling process again. Since the specifics of Input and Display Modules are domain dependent we describe both of them, first for LAX and then for FAMS.

4.2 LAX Domain: ARMOR

LAX is the fifth busiest airport in the United States, the largest destination airport in the United States, and serves 60-70 million passengers per year [1, 42]. LAX is known to be a prime terrorist target on the west coast of the United States, with multiple arrests of plotters attempting to attack LAX [42]. To protect LAX, the airport police have designed a security system that utilizes multiple rings of protection. As is evident to anyone traveling through the airport, these rings include such things as vehicular checkpoints, police units patrolling the roads to the terminals, patrolling inside the terminals (with canines), and security screening and bag checks for passengers. Airport police use intelligent randomization within two of these rings: (1) placing vehicle checkpoints on inbound roads that service the LAX terminals, including both location and timing (a checkpoint is shown in Figure 2) (2) scheduling



Fig. 2 Security Checkpoints and Canine Patrols at LAX

patrols for bomb-sniffing canine units at the different LAX terminals (as shown in Figure 2). The numbers of available vehicle checkpoints and canine units are limited by resource constraints, so randomization is used as a method to increase the effectiveness of these resources while avoiding creating patterns in deployment.

The eight different terminals at LAX have very different characteristics, leading to different assessments of the value/risk for each terminal. For example, international flights are concentrated at a few terminals, while terminals have varying physical size and passenger loads. Because uncertainty about the adversary was identified by airport police as a key problem, the model should take into account the different types of adversaries that may be encountered. For example, there may be both hard-line, well-funded international terrorists planning attacks as well as amateur individuals. The payoff values for different attack scenarios should depend on the type of attacker and their capabilities.

The interface for the ARMOR checkpoints program is shown in Figure 3 and provides options for the number of available resources, the number of scheduled days, the time slots to schedule, and the monthly calendar. A spreadsheet is used to display the proposed schedule and provide additional opportunities for the end users to modify the schedules in an iterative process. Three options are provided to change the possible scheduling actions: (i) number of checkpoints allowed during a particular time slot; (ii) the time interval of each time slot; (iii) the number of days to schedule over. Furthermore, three options are given to the user to enforce constraints onto the schedule: (i) forced checkpoint; (ii) forbidden checkpoint; (iii) at least one checkpoint. These constraints are intended to be used sparingly to accommodate situations where a user, faced with exceptional circumstances and extra knowledge, wishes to influence the output of the game. The user can impose these specific actions in the schedule using the spreadsheet interface. Each restriction is represented by a different color in the spreadsheet. The interface for the ARMOR Canine Patrols at LAX has similar features.

ARMOR generates a different game for each time slot on each day. The number of defender resources in the model is the number of canine units/checkpoints spec-

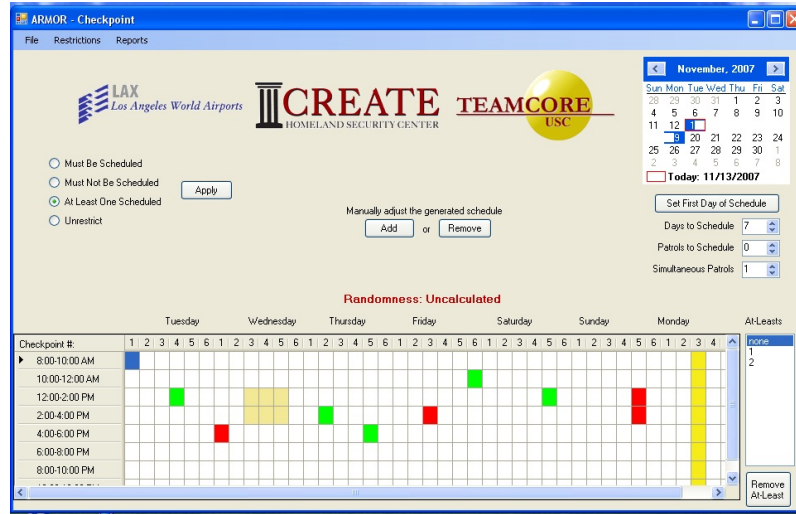


Fig. 3 ARMOR Interface

ified by the user. The number of targets is the number of terminals for the canines system, and the number of inbound roads for the checkpoints system. Generating the game matrix also requires values for the payoffs associated with each possible target. These payoff values depend on a variety of conditions, such as passenger loads, cost of the infrastructure, publicity to the adversary, etc. Domain experts provided us with formulae to automatically generate payoff values for all possible combinations of such conditions, which we encode in ARMOR. The system is also provided with estimates of the passenger load and other elements (the details of these formulae and tools cannot be discussed due to security concerns). For any given day, ARMOR is able to take the conditions for this day and select appropriate payoff values for the targets. As a result, it is not necessary for LAX police officers to enter these values by hand to generate each schedule, which is both time-consuming and error-prone. The system still retains a high degree of flexibility because values are precomputed and stored for a wide range of possible conditions.

The generated schedule of checkpoints and canines is presented to the user via a spreadsheet. Each row in the output spreadsheet corresponds to one hour. Each column in the sheet corresponds to a terminal. Each entry in the sheet represents a schedule generated by ARMOR. The familiarity of the police officers with spreadsheets helped in the acceptance of the ARMOR schedules.

When ARMOR identifies that user constraints are causing unreasonably low likelihood of scheduling a checkpoint, it presents the schedule to the user with alerts. The user may then alter the schedule by modifying the forbidden/required checkpoints, or possibly by directly altering the schedule. Both possibilities are accommodated in ARMOR. If the user simply adds or removes constraints, ARMOR can create a new schedule. Once the schedule is finalized, it can be saved for actual use,

thus completing the system cycle. This full process was designed to specifically meet the requirements at LAX for checkpoint and canine allocation.

4.3 FAMS Domain: IRIS

The Federal Air Marshals Service (FAMS) places undercover law enforcement personnel aboard flights originating in and departing from the United States to dissuade potential aggressors and prevent an attack should one occur [3]. The exact methods used to evaluate the risks posed by individual flights is not made public by the service, but we can identify many factors that might influence such an evaluation. For example, flights have different numbers of passengers, and some fly over densely populated areas while others do not. International flights also serve different countries, which may pose different risks. Special events can also change the risks for particular flights at certain times [2].

The scale of the domain is massive. There are currently tens of thousands of commercial flights scheduled each day, and public estimates state that there are thousands of air marshals. 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. Our task is made more difficult by the need to find a randomized policy that meets these scheduling constraints, while also accounting for the different values of each flight.

The FAMS domain is considerably larger, and the information required to build a game model in this domain changes much more frequently. For these reasons the application is considerably more complex than ARMOR in terms of the user interface and the mechanisms required to input all of the necessary information. This additional complexity is necessary in this domain to accurately capture the situation, and provide all of the functionality requested by the end users. However, it does place a greater burden on the users to learn the system, and scheduling is a more time-consuming process than in ARMOR. Again, finding the right level of complexity to expose to the users was an iterative process that involved many discussions with the users and domain experts.

In the FAMS domain, we require information about the available air marshals, their scheduling constraints, the possible flights, and information about the risks/values to associate with each flight. The data about resources includes information about the number and location of air marshals, as well as the conditions that define legal flight schedules. Flight information includes various data about each flight, including flight number, carrier, origin, destination, aircraft type, etc. Finally, some information is collected to improve usability, even though it is not strictly necessary for the game-theoretic analysis. This includes naming schemes for airports and airlines and other information that allows the system to output schedules in a more usable format or to interface easily with other systems. IRIS also includes

functionality to import data from existing databases with flight data and other information. This greatly reduces the amount of data entry necessary to create a schedule.

Specifying the payoff values for every possible flight was a particular challenge in this domain, since there are thousands of flights to consider. We use an attributed-based system to elicit these values, based on the Threat, Vulnerability and Consequence (TVC) model for estimating terrorism risk [51]. By eliciting values for attributes of flights rather than specific flights, we are able to dramatically reduce the number of entries required by the user. Each flight is then given an aggregate value based on these components; the specific calculations used to determine flight risk are sensitive information, and cannot be revealed. The values of the attributes for each flight can be populated automatically from existing databases. To allow for specific intelligence or exceptional circumstances, the individual payoff values for any flight can also be directly edited by the end user. However, this is only rarely necessary and the majority of the analysis can be effectively automated.

This preference elicitation system of IRIS has substantially reduced the number of values that must be entered by the user. During a restricted test run on real data, the attribute-based approach called for a total of 114 values to input regardless of the number of flights. By contrast, there were 2,571 valid flights over a week, each requiring 4 payoff values, summing to 10,284 user-entered values without the attribute-based preference elicitation system. The attribute-based approach clearly requires far fewer inputs and remains constant as the number of flights increases, allowing for excellent scalability as we deal with larger and larger sets of flights. Equally importantly, attribute-based risk assessment is an intuitive and highly-scalable method that can be used in any problem where people must distill numerous attributes of a situation into a single value for a large number of situations that share the same attributes.

The generated schedules are presented to the user via the application window. The schedule created is shown in the interface, and allows the users to view more detailed information about each target. The user is also able to output the schedule to a file which can then be used to analyze the schedule in more detail. The sample assignment of federal air marshals to flight schedules is exactly a schedule that could be used by the FAMS. At this point, the scheduling assistant allows the expert using the system to create numerous sample schedules based on the same optimal mixed strategy or to change the assignment of federal air marshals to flight schedules by hand to create a final schedule that meets the needs of the FAMS. Of course, the user can also adjust any of the parameters entered and re-solve the game completely. The output of IRIS is in the same format as the other systems used by the FAMS officers. It has not been presented here for simplicity and because of security concerns.

4.4 Lessons Learned

The design and deployment of ARMOR and IRIS has posed numerous challenges. We outline some key lessons learned during the design and deployment of these

tools. Firstly, there is a critical need for randomization in security operations. Security officials are aware that requiring humans to generate randomized schedules is unsatisfactory because as psychological studies have often shown [48, 44], humans have difficulty in randomizing, and they can also fall into predictable patterns. Instead, game-theoretic randomization that appropriately weighs the costs and benefits of different actions, and randomizes with appropriate weights leads to improved results. Security officials were therefore extremely enthusiastic in their reception of our research, and eager to apply it to their practices.

Secondly, organizational acceptance is a key issue. In creating solutions for people, we must be cognizant of how difficult it will be for a user to adopt our solution. Each deviation from existing methodology is a step away from the familiar that we must convince the user to accept. Instead of asking people to make numerous and sometimes unnecessary changes, minimizing these differences and complexities can help pave the way towards a successful implementation. For example, tweaking the GUI to achieve a look and feel that the user is familiar and comfortable with can help the user understand the system faster and better. Similarly, because infrastructural changes are often costly and/or time-consuming, ease of incorporating our work into their daily routine is essential. For example, using inputs and creating outputs that were in the same format as existing protocols minimized the additional work that our assistant would create for the security officers and lead to easier acceptance of the system.

Thirdly, it is important to provide the users with operational flexibility. When initially generating schedules for canine patrols, we created a very detailed schedule, micro-managing the patrols. This did not get as positive a reception from the officers. Instead, an abstract schedule that afforded the officers some flexibility to respond to situations on the ground was better received.

5 A Generic Network Security Problem

As noted above, implementing a Stackelberg security game model to plan patrols is a difficult process that to date has been undertaken with substantial effort in collaboration with the security providers. In many situations, however, there is enough information about the security process that a data-driven process could be used to assist security providers in defining the actions and payoffs of the security game. In this section we illustrate recent work that aims to automatically build a Stackelberg security game for the problem of patrolling a street network to prevent crime. The proposed approach uses data mining tools on a database of past reported crime and events to identify the locations to be patrolled, the times at which the game changes, and the types of adversaries faced. The idea is to exploit temporal and spatial patterns of crime on the area to be patrolled to determine the priorities on how to use the limited security resources.

We consider the street network depicted in Figure 4 which corresponds to a centric commercial, turistic, and economic district in Santiago, Chile. This is a busy part

of the city usually with large crowds on the street and that historically concentrates a high number of crimes, for the most part theft or minor agressions. This type of crime in particular can be deterred or reduced with appropriate patrolling by police. To represent the problem of deciding where to patrol as a Stackelberg security game, security providers need to identify the specific points on this street network that concentrate crime and determine the payoffs defenders and attackers would receive if crimes at these locations are committed or are prevented. In this security game, police patrols on foot would go to the points selected following the random optimal mixed strategy that maximizes the defender's utility. Different types of criminals would, knowing the optimal mixed strategy of the police patrols, then decide where to attack on the network, if at all. We assume that both police and criminals appear at the point selected, without interacting in other parts of the network. In addition to the description of the street network, we obtained from the Chilean national police force information about reported crimes in the area and the police reports for a two year period. Each reported crime has a location, a date and time, and a description of the crime (classification of crime [robbery, theft, etc.], amount stolen, level of violence, etc.) The police reports include information about the available resources in each shift, which helps estimate the police resources used for preventive patrolling.

5.1 Building a Data-Driven Security Game

This information is then processed in an automated data-driven procedure to build a security game in five steps: 1) Define the amount of data that will be relevant to calibrate the security game, 2) Determine locations to patrol, 3) Identify attacker types from data, 4) Determine times to patrol, 5) Determine payoffs for leader and followers.

- Step 1: In determining which data to use to build a representative security game we must strike a balance between selecting too much data and too little. Sufficient past data should be included so that significant but perhaps rare patterns of crime are taken into consideration. However, if too much data is taken into account we run the risk of representing crime patterns that no longer exist. You must rely on expert opinion to estimate how representative past data is of the current security situation leading to an estimate of how much of the past information to use in identifying the locations of the patrols, the types of adversaries, and the utilities for each. In the results we show below we used a time window of 2 years of data (from December 15, 2002 until December 14, 2004) to build a week long game (for the week of December 15-22, 2004).
- Step 2: We used an off-the-shelf clustering software to identify the locations to be patrolled from the density plot of reported crime displayed in Figure 5. These locations are selected anywhere on the road network in a way that summarizes the geographical distribution of crimes without requiring a massive number of locations. We used the software DBSCAN (Density-Based Spatial Clustering of Applications with Noise)[18]. This is a density segmentation tool which also re-

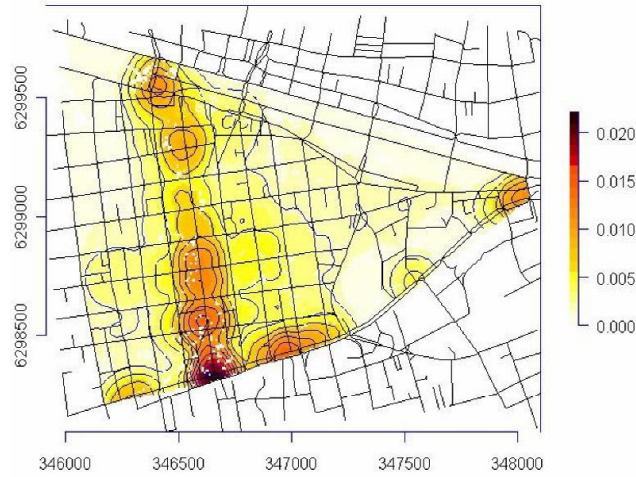


Fig. 4 Patrolling area with a density plot of reported crimes in the period 12/15/2002 - 12/14/2004.

moves the noise in the data and automatically selects the number of segments to consider. In the results we obtained, DBSCAN identifies 119 locations to protect in which there are at least 10 crimes within a radius of 20 meters. These points represent 89.23% of the reported crimes. We note that a number of good clustering algorithms can help in identifying a set of locations that are representative of the spatial crime distribution.

- Step 3: We follow the Knowledge Discovery in Databases (KDD) scheme [19] to process the database of reported crimes and identify different types of attackers. The KDD approach is a generic scheme that outlines a series of procedures to, among other things, create a target data set, remove data noise and outliers, handle missing data, identify useful features in the data, etc. Each of the processes can be implemented with any of a number of existing tools. For the selection of attributes, we chose a wrapper technique that automatically selects the attributes that help segmentation [17]. To identify the clusters of crimes we use a k-means clustering model. We found that this model was superior to alternative clustering models we tried (X-means, Expectation Maximization) for this problem, both in runtime and the quality of solutions found, which are more easily interpretable. The number of reported crimes in each cluster informs us of the frequency of different types of crime and thus the likelihood of facing each. The crimes in the 2 year database were classified into 9 significant clusters that were characterized by 24 significant attributes.
- Step 4: Since the security conditions change during the day and the Stackelberg security game describes static conditions, we separate the day into different time intervals (or blocks) in which the security conditions remain almost constant. The different types of crime identified in Step 3 include 3 different time blocks which are found to be significant. Intersecting these times with the police patrolling shifts gives

us a total of 7 time intervals, or blocks, where the likelihood and composition of different types of crime and patrolling resources are kept about constant.

| Block | From | To | Block | From | To | Block | From | To |
|-------|-------|-------|-------|-------|-------|-------|-------|-------|
| S1 | 0:00 | 6:59 | S2 | 7:00 | 9:59 | S3 | 10:00 | 14:59 |
| S4 | 15:00 | 17:59 | S5 | 18:00 | 19:59 | S6 | 20:00 | 21:59 |
| S7 | 22:00 | 23:59 | | | | | | |

In blocks S2 and S3 there are 23 patrolling units available, in blocks S4, S5 and S6 there are 24 patrolling units, and in blocks S1 and S7 there are 9 patrolling units. Here, one patrolling unit corresponds to a pair of policemen on foot.

We determine the probability of facing each type of adversary by the frequency with which each of the 9 types of crimes occur. To make this frequency more dependent on recent events, the past event data is scaled with an exponential decay function. Table 3 below shows these frequencies for each of the 9 types of crimes over the 7 time blocks found.

| Cluster | S1 | S2 | S3 | S4 | S5 | S6 | S7 | Total |
|---------|-------|-------|-------|-------|-------|-------|-------|-------|
| 0 | 0.234 | 0.516 | 0.624 | 0 | 0.603 | 0.562 | 0.395 | 1815 |
| 1 | 0.078 | 0.057 | 0.048 | 0.142 | 0.049 | 0.079 | 0.097 | 679 |
| 2 | 0 | 0 | 0 | 0.470 | 0 | 0 | 0 | 545 |
| 3 | 0.032 | 0.018 | 0.018 | 0 | 0.012 | 0.027 | 0.050 | 369 |
| 4 | 0 | 0 | 0 | 0.260 | 0 | 0 | 0 | 405 |
| 5 | 0.253 | 0.091 | 0.063 | 0.079 | 0.066 | 0.093 | 0.150 | 808 |
| 6 | 0.023 | 0.027 | 0.022 | 0.048 | 0.033 | 0.016 | 0.024 | 419 |
| 7 | 0 | 0 | 0 | 0 | 0 | 0.223 | 0.285 | 575 |
| 8 | 0.381 | 0.291 | 0.225 | 0 | 0.238 | 0 | 0 | 1110 |
| Total | 727 | 457 | 1892 | 1217 | 939 | 881 | 612 | |

Table 3 Probability of facing each follower in the different time blocks

Step 5: In this work we determine the payoffs for the attacker as a valuation of the monetary payoff of being successful or getting caught for each type of crime. In the case of the police, we estimate that the payoff for catching a criminal is zero (for all types) while the penalty for a successful crime equals the expected amount earned by that type of criminal. We first determine from the information on the database the average expected reward for the criminal in a successful attack. To determine the penalty of an unsuccessful attack we estimated the expected number of days in jail for that type of crime and evaluated the amount of forgone earnings for the criminal for not being able to commit crimes during that period. We note that there are a number of alternative models that can be incorporated here, in particular models of risk aversion that better represent human behavior in adversarial environments, such as prospect theory [25] or quantal response [31]. Table 4 below presents the values of payoffs for the Stackelberg game for each of the 9 types of adversaries.

| Cluster | Average Utility | Prison Time | Average Cost |
|---------|-----------------|-------------|--------------|
| 0 | 182 | 61 | 638 |
| 1 | 209 | 1752 | 731 |
| 2 | 136 | 63 | 476 |
| 3 | 451 | 1746 | 1579 |
| 4 | 175 | 1747 | 614 |
| 5 | 217 | 1686 | 761 |
| 6 | 139 | 74 | 486 |
| 7 | 138 | 1757 | 485 |
| 8 | 218 | 1739 | 764 |

Table 4 Expected payoff each type of criminal, in US dollars, if attack is successful (Average Utility) and if attack is unsuccessful (Average Cost, using a 40% discount rate while in prison).

5.2 Additional Considerations in a Data-Driven Security Game

The procedure above helps security providers build a Stackelberg security game to determine efficient patrols in an urban street network. This game can then be formulated as the mixed integer programs described in Section 3 and solved to optimality. A solution for this problem is depicted in Figure 5 below. The color at each node corresponds to the amount of coverage in the optimal mixed strategy for a certain time block. To implement this solution, the police should sample from this distribution to decide which locations to patrol each day in every time block.




Fig. 5 Optimal mixed strategy on locations selected. Node color corresponds to probability of coverage of the node, with a darker color indicating a higher probability of coverage.

The game developed can also be used to evaluate the current practice and the proposed patrol plan. Currently police direct their preventive patrols to the locations where the highest concentration of crime is expected to occur, based on recent past activity (2 weeks). We assume that the highest concentration of crime are the locations where the game predicts the highest payoff for the adversary, therefore directing the patrols to the maximum payoff locations leading to a minimax strategy. Table 5 presents the defender’s expected profits in each time block under each of 4 different strategies: the optimal mixed and pure strategies of the Stackelberg game and Minimax. We note that the utility for the leader is always better in the Stackelberg game (mixed).

| Block of Time | S1 | S2 | S3 | S4 | S5 | S6 | S7 |
|---------------------|-------|-------|-------|-------|-------|-------|-------|
| Stackelberg (mixed) | -1.87 | -0.31 | -0.29 | -0.24 | -0.21 | -0.22 | -1.63 |
| Stackelberg (pure) | -8.87 | -3.90 | -3.62 | -3.27 | -3.53 | -3.48 | -8.30 |
| Maximin (mixed) | -5.40 | -2.42 | -2.21 | -1.94 | -2.18 | -2.18 | -5.10 |
| Maximin (pure) | -8.87 | -3.95 | -3.67 | -3.27 | -3.53 | -3.48 | -8.30 |

Table 5 Defender’s expected utility in different time blocks

The set of tools described here hope to complement the experience and intuition of law enforcement. There is much information that is difficult to include in decisions on how to patrol. This is the case in part because of the amount of data and in part because the data is not being collected, or is biased. We note that a better description of the security problem can be obtained, and thus a better security game formulated, by incorporating additional sources of information, such as surveys of victimization, physical description of places, etc. We believe this is an interesting avenue of future research to create robust systems that would be more easily deployable in diverse settings.

6 Conclusions

Monitoring and patrolling are key components of law enforcement in security domains. In generating schedules for these patrols, it is important to account for varying weights of the targets being protected as well as the fact that potential attackers can often observe the procedures being used. This article describes scheduling assistants for the LAX police, ARMOR, and the FAMS, IRIS, which provide game-theoretic solutions to this problem. The two systems assist the security forces in generating randomized patrols while ensuring that differences in importance of different targets are preserved. A critical observation in the deployment of these scheduling assistants is the difficulty faced in reducing a complex security domain to a Stackelberg game model. To address this difficulty we present a data-mining based model to assist security personnel in defining the Stackelberg security game from historic data.

ARMOR and IRIS make use of algorithmic advances in multi-agent systems research to solve the class of massive security games with complex constraints that were not previously solvable in realistic time-frames. Thus, although our applications were designed to be deployed at LAX and FAMS, they provide a general framework for solving patrolling scheduling problems in other domains as well.

Our approach of using Stackelberg games to model real-world security problems is applicable in a wide range of domains that share the following attributes: (i) there are intelligent players, (ii) one player's strategy is observable by the other player, (iii) player's have varying preferences among targets, and (iv) it is not possible to provide full coverage of all targets. Some examples of similar security situations include security in computer networks, checkpoints at subway stations, security inspections at ports and monitoring of other mediums of public transport.

Ultimately the security providers (Police, Air Marshals) are the judge of the usefulness of these Stackelberg security game models. As in any model it is critical to allow for expert knowledge to inform the system and provide feedback on the quality of solutions. With this in mind the development of the interface of these deployed systems has been an important aspect of this work. This research and these applications have been effective in helping in the security officers with scheduling and patrolling concerns. Thus, ARMOR and IRIS represent successful transitions of game-theoretic advances to applications that have been in use and effective in the real world. There are a number of additional improvements to these systems that could be done in the future to facilitate deployment to different domains. Some lines of future research include methods to incorporate qualitative information (estimates of unreported crime, fear of crime, etc.) to construct the Stackelberg games; coordination of different security resources and considering attackers who deviate from rational behavior (due to differences in information or human bias).

References

1. General Description: Just the Facts. 2007. <http://www.lawa.org/lax/justTheFact.cfm>.
2. Federal Air Marshal Service. 2008. http://en.wikipedia.org/wiki/Federal_Air_Marshal_Service.
3. TSA: Federal Air Marshals. 2008. <http://www.tsa.gov/lawenforcement/programs/fams.shtm>.
4. N. Agmon, V. Sadov, G. A. Kaminka, and S. Kraus. The Impact of Adversarial Knowledge on Adversarial Planning in Perimeter Patrol. In *AAMAS*, volume 1, pages 55–62, 2008.
5. B. An, J. Pita, E. Shieh, M. Tambe, C. Kiekintveld, and J. Marecki. GUARDS and PROTECT: Next generation applications of security games. *ACM SIGecom Exchanges*, 10(1).
6. R. Avenhaus, B. von Stengel, and S. Zamir. Inspection Games. In R. J. Aumann and S. Hart, editors, *Handbook of Game Theory*, volume 3, chapter 51, pages 1947–1987. North-Holland, Amsterdam, 2002.
7. L. Babu, L. Lin, and R. Batta. Passenger Grouping Under Constant Threat Probability in an Airport Security System. In *European Journal of Operational Research*, volume 168, pages 633 – 644, 2006.
8. J. F. Bard. *Practical Bilevel Optimization: Algorithms and Applications*, volume 30 of *Non-convex Optimization and Its Applications*. Springer, 1999.
9. T. Basar and G. J. Olsder. *Dynamic Noncooperative Game Theory*. Academic Press, San Diego, CA, 2nd edition, 1995.

10. V. M. Bier. Choosing What to Protect. *Risk Analysis*, 27(3):607–620, 2007.
11. M. Blanco, A. Valino, J. Heijs, T. Baumert, and J. G. Gomez. The Economic Cost of March 11: Measuring the direct economic cost of the terrorist attack on March 11, 2004 in Madrid. *Terrorism and Political Violence*, 19(4):489–509, 2007.
12. M. Breton, A. Alg, and A. Haurie. Sequential Stackelberg Equilibria in Two-Person Games. *Optimization Theory and Applications*, 59(1):71–97, 1988.
13. G. Brown, M. Carlyle, J. Kline, and K. Wood. A Two-Sided Optimization for Theater Ballistic Missile Defense. In *Operations Research*, volume 53, pages 263–275, 2005.
14. G. Brown, M. Carlyle, J. Roysset, and K. Wood. On The Complexity of Delaying an Adversary’s Project. In B. Golden, S. Raghavan, and E. Wasil, editors, *The Next Wave in Computing, Optimization and Decision Technologies*, pages 3–17. Springer, 2005.
15. G. Brown, M. Carlyle, J. Salmeron, and K. Wood. Defending Critical Infrastructure. In *Interfaces*, volume 36, pages 530 – 544, 2006.
16. V. Conitzer and T. Sandholm. Computing the Optimal Strategy to Commit to. In *ACM EC-06*, pages 82–90, 2006.
17. J. G. Dy and C. E. Brodley. Feature selection for unsupervised learning. *Journal of Machine Learning Research*, 5:845–889, 2004.
18. M. Ester, H.-P. Kriegel, J. Sander, and X. Xu. A density-based algorithm for discovering clusters in large spatial database with noise. Technical report, Institute for Computer Science, University of Munich, 1996.
19. U. Fayyad, G. Piatetsky-Shapiro, and P. Smyth. From data mining to knowledge discovery: An overview. *AI Magazine*, 17(3):37–54, 1996.
20. D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, 1991.
21. N. Gatti. Game Theoretical Insights in Strategic Patrolling: Model and Algorithm in Normal-form. In *ECAI-08*, pages 403–407, 2008.
22. J. C. Harsanyi and R. Selten. A Generalized Nash Solution for Two-person Bargaining Games With Incomplete Information. *Management Science*, 18(5):80–106, 1972.
23. M. Jain, J. Tsai, J. Pita, C. Kiekintveld, S. Rathi, F. Ordóñez, and M. Tambe. Software assistants for patrol planning at lax and federal air marshals service. *Interfaces*, 40(4):267–290, 2010.
24. A. Jiang and K. Leyton-Brown. A Polynomial-time Algorithm for Action-Graph Games. In *Artificial Intelligence*, pages 679–684, 2006.
25. D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292, 1979.
26. C. Kiekintveld, M. Jain, J. Tsai, J. Pita, M. Tambe, and F. Ordóñez. Computing Optimal Randomized Resource Allocations for Massive Security Games. In *AAMAS-09*, 2009.
27. D. Koller and B. Milch. Multi-agent Influence Diagrams for Representing and Solving Games. *Games and Economic Behavior*, 45(1):181–221, 2003.
28. R. Larson. A Hypercube Queueing Modeling for Facility Location and Redistricting in Urban Emergency Services. In *Journal of Computers and Operations Research*, volume 1, pages 67–95, 1974.
29. G. Leitmann. On Generalized Stackelberg Strategies. *Journal of Optimization Theory and Applications*, 26(4):637–643, 1978.
30. R. Looney. Economic Costs to the United States Stemming From the 9/11 Attacks. *Strategic Insights*, 1(6), August 2002.
31. R. D. McKelvey and T. R. Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 2:638, 1995.
32. Nie, R. Batta, Drury, and Lin. Optimal Placement of Suicide Bomber Detectors. In *Military Operations Research*, volume 12, pages 65–78, 2007.
33. M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
34. P. Paruchuri, J. P. Pearce, J. Marecki, M. Tambe, F. Ordóñez, and S. Kraus. Playing Games with Security: An Efficient Exact Algorithm for Bayesian Stackelberg games. In *AAMAS-08*, pages 895–902, 2008.
35. P. Paruchuri, J. P. Pearce, M. Tambe, F. Ordóñez, and S. Kraus. An Efficient Heuristic Approach for Security Against Multiple Adversaries. In *AAMAS*, 2007.

36. P. Paruchuri, M. Tambe, F. Ordonez, and S. Kraus. Security in Multiagent Systems by Policy Randomization. In *AAMAS*, 2006.
37. J. Pita, M. Jain, C. Western, C. Portway, M. Tambe, F. Ordóñez, S. Kraus, and P. Parachuri. Deployed ARMOR protection: The application of a game-theoretic model for security at the Los Angeles International Airport. In *AAMAS-08 (Industry Track)*, 2008.
38. J. Pita, M. Tambe, C. Kiekintveld, S. Cullen, and E. Steigerwald. GUARDS - game theoretic security allocation on a national scale. In *International Conference on Autonomous Agents and Multiagent Systems*, 2011.
39. S. Ruan, C. Meirina, F. Yu, K. R. Pattipati, and R. L. Popp. Patrolling in a Stochastic Environment. In *10th Intl. Command and Control Research and Tech. Symp.*, 2005.
40. T. Sandler and D. G. A. M. Terrorism and Game Theory. *Simulation and Gaming*, 34(3):319–337, 2003.
41. V. Srivastava, J. Neel, A. B. MacKenzie, R. Menon, L. A. Dasilva, J. E. Hicks, J. H. Reed, and R. P. Gilles. Using Game Theory to Analyze Wireless Ad Hoc Networks. *IEEE Communications Surveys and Tutorials*, 7(4), 2005.
42. D. Stevens and et. al. Implementing Security Improvement Options at Los Angeles International Airport. 2006. http://www.rand.org/pubs/documented_briefings/2006/RAND_DB499-1.pdf.
43. P. Thornton. London Bombings: Economic Cost of Attacks Estimated at 2bn, July 2005. <http://www.independent.co.uk/news/business/news/economic-cost-of-attacks-estimated-at-1632bn-499281.html>.
44. M. Treisman and A. Faulkner. Generation of Random Sequences by Human Subjects: Cognitive Operations or Psychological Process? *Journal of Experimental Psychology*, 116(4):337–355, 1987.
45. J. Tsai, S. Rathi, C. Kiekintveld, F. Ordóñez, and M. Tambe. IRIS - A Tool for Strategic Security Application in Transportation Networks. In *AAMAS Industry Track*, 2009.
46. H. von Stackelberg. *Marktform und Gleichgewicht*. Springer, Vienna, 1934.
47. B. von Stengel and S. Zamir. Leadership with Commitment to Mixed Strategies. Technical Report LSE-CDAM-2004-01, CDAM Research Report, 2004.
48. W. A. Wagenaar. Generation of Random Sequences by Human Subjects: A Critical Survey of Literature. *Psychological Bulletin*, 77(1):65–72, 1972.
49. K. wei Lye and J. M. Wing. Game Strategies in Network Security. *International Journal of Information Security*, 4(1–2):71–86, 2005.
50. L. M. Wein. Homeland Security: From Mathematical Models to Policy Implementation. In *Operations Research*, 2008.
51. H. Willis, A. Morral, T. Kelly, and J. Medby. *Estimating Terrorism Risk*. RAND Corporation, 2005.