

PROTECT – A Deployed Game Theoretic System for Strategic Security Allocation for the United States Coast Guard

Bo An, Eric Shieh, Rong Yang, Milind Tambe

University of Southern California
{boa, eshieh, yangrong, tambe}@usc.edu

Craig Baldwin, Joseph DiRenzo, Ben Maule, Garrett Meyer

United States Coast Guard
{Craig.W.Baldwin, Joseph.DiRenzo, Ben.J.Maule, Garrett.R.Meyer}@uscg.mil

Abstract

While three deployed applications of game theory for security have recently been reported, we as a community of agents and AI researchers remain in the early stages of these deployments; there is a continuing need to understand the core principles for innovative security applications of game theory. Towards that end, this paper presents PROTECT, a game-theoretic system deployed by the United States Coast Guard (USCG) in the Port of Boston for scheduling their patrols. USCG has termed the deployment of PROTECT in Boston a success; PROTECT is currently being tested in the Port of New York, with the potential for nationwide deployment.

PROTECT is premised on an attacker-defender Stackelberg game model and offers five key innovations. First, this system is a departure from the assumption of perfect adversary rationality noted in previous work, relying instead on a quantal response (QR) model of the adversary's behavior — to the best of our knowledge, this is the first real-world deployment of the QR model. Second, to improve PROTECT's efficiency, we generate a compact representation of the defender's strategy space, exploiting equivalence and dominance. Third, we show how to practically model a real maritime patrolling problem as a Stackelberg game. Fourth, our experimental results illustrate that PROTECT's QR model more robustly handles real-world uncertainties than a perfect rationality model. Finally, in evaluating PROTECT, this paper for the first time provides real-world data: (i) comparison of human-generated vs PROTECT security schedules, and (ii) results from an Adversarial Perspective Team's (human mock attackers) analysis.

Introduction

The last several years have witnessed the successful application of game theory in allocating limited resources to protect critical infrastructures. The real-world deployed systems include ARMOR, used by the Los Angeles International Airport (Jain et al. 2010) to randomize checkpoints of roadways and canine patrols; IRIS, which helps the US Federal Air Marshal Service (Jain et al. 2010) in scheduling air marshals on international flights; and GUARDS (Pita et al. 2011), which is under evaluation by the US Transportation Security Administration to allocate the resources available for airport protection. Yet we as a community of agents

and AI researchers remain in the early stages of these deployments, and must continue to develop core principles of applying game theory for security.


To this end, this paper presents a new game-theoretic security application to aid the United States Coast Guard (USCG), called *Port Resilience Operational/Tactical Enforcement to Combat Terrorism* (PROTECT). The USCG's mission includes maritime security of the US coasts, ports, and inland waterways; a security domain that faces increased risks in the context of threats such as terrorism and drug trafficking. Given a particular port and the variety of critical infrastructure that an adversary may attack within the port, USCG conducts patrols to protect this infrastructure; however, while the adversary has the opportunity to observe patrol patterns, limited security resources imply that USCG patrols cannot be at every location 24/7. To assist the USCG in allocating its patrolling resources, similar to previous applications (Jain et al. 2010; Pita et al. 2011), PROTECT uses an attacker-defender Stackelberg game framework, with USCG as the defender against terrorist adversaries that conduct surveillance before potentially launching an attack. PROTECT's solution is to provide a mixed strategy, i.e., randomized patrol patterns taking into account the importance of different targets, and the adversary's surveillance and anticipated reaction to USCG patrols.

While PROTECT builds on previous work, this paper highlights five key innovations. The first and most important is PROTECT's departure from the assumption of perfect rationality (which was embedded in previous applications) on the part of the human adversaries. The assumption of perfect rationality is well-recognized as a limitation of classical game theory, and a number of approaches have been proposed to model bounded rationality in behavioral game-theoretic approaches (Camerer 2003). Within this behavioral framework, quantal response equilibrium has emerged as a promising approach to model human bounded rationality (Camerer 2003; McKelvey and Palfrey 1995; Wright and Leyton-Brown 2010) including recent results illustrating the benefits of the quantal response (QR) model in security games contexts (Yang et al. 2011). Therefore, PROTECT uses a novel algorithm called PASAQ (Yang, Tambe, and Ordóñez 2012) based on the QR model of a human adversary. To the best of our knowledge, this is the first time

that the QR model has been used in a real-world security application.

Second, PROTECT improves PASAQ’s efficiency via a compact representation of defender strategies exploiting dominance and equivalence analysis. Experimental results show the significant benefits of this compact representation. Third, PROTECT addresses practical concerns (such as grouping patrol points into areas) of modeling real-world maritime patrolling application in a Stackelberg framework. Fourth, this paper presents a detailed simulation analysis of PROTECT’s robustness to uncertainty that may arise in the real-world. For various cases of added uncertainty, the paper shows that PROTECT’s quantal-response-based approach leads to significantly improved robustness when compared to an approach that assumes full attacker rationality.

PROTECT has been in use at the Port of Boston since April 2011 and been evaluated by the USCG. This evaluation brings forth our final key contribution: for the first time, this paper provides real-world data comparing human-generated and game-theoretic schedules. We also provide results from an Adversarial Perspective Team’s (APT) analysis and comparison of patrols before and after the use of the PROTECT system from a viewpoint of an attacker. Given the success of PROTECT in Boston, PROTECT is currently being tested in the Port of New York, and based on the outcome there, it may potentially be extended to other ports in the US.



(a) PROTECT is being used in Boston (b) Extending PROTECT to NY

Figure 1: USCG boats patrolling the ports of Boston and NY

USCG and PROTECT’s Goals

The USCG continues to face challenges from potential terrorists within the maritime environment, which includes both the Maritime Global Commons and the ports and waterways that make up the United States Maritime Transportation System. The former Director of National Intelligence, Dennis Blair noted in 2010 a persistent threat “from al-Qa’ida and potentially others who share its anti-Western ideology. A major terrorist attack may emanate from either outside or inside the United States” (Blair 2010). This threat was reinforced in May of 2011 following the raid on Osama Bin Laden’s home, where a large trove of material was uncovered, including plans to attack an oil tanker. “There is an indication of intent, with operatives seeking the size and construction of tankers, and concluding it’s best to

blow them up from the inside because of the strength of their hulls” (Dozier 2011). These oil tankers transit the U.S. Maritime Transportation System. The USCG plays a key role in the security of this system and the protection of seaports to support the economy, environment, and way of life in the US (Young and Orchard 2011).

Coupled with challenging economic times, USCG must operate as effectively as possible, achieving maximum benefit from every hour spent on patrol. To that end, the goal of PROTECT is to use game theory to assist the USCG in maximizing its effectiveness in the Ports, Waterways, and Coastal Security (PWCS) Mission. The PROTECT system, focused on the PWCS patrols, addresses how the USCG should optimally patrol critical infrastructure in a port to maximize protection, knowing that the adversary may conduct surveillance and then launch an attack.

PWCS patrols are focused on protecting critical infrastructure; without the resources to provide one hundred percent on scene presence at any, let alone all of the critical infrastructure, optimized use of security resources is critical. Towards that end, unpredictability creates situations of uncertainty for an enemy and can be enough to deem a target less appealing. While randomizing patrol patterns is key, PROTECT also addresses the fact that the targets are of unequal value, understanding that the adversary will adapt to whatever patrol patterns USCG conducts. The output of PROTECT is a schedule of patrols which includes when the patrols are to begin, what critical infrastructure to visit for each patrol, and what activities to perform at each critical infrastructure. While initially pilot tested in the Port of Boston, the solution technique is intended to be generalizable and applicable to other ports.

Key Innovations in PROTECT

Stackelberg games have been well established in the literature (Conitzer and Sandholm 2006; Korzhyk, Conitzer, and Parr 2011; Fudenberg and Tirole 1991) and PROTECT models the PWCS patrol problem as a Stackelberg game with USCG as the leader (defender) and the terrorist adversaries in the role of the follower. The choice of this framework was also supported by prior successful applications of Stackelberg games (Jain et al. 2010; Paruchuri et al. 2008; Pita et al. 2011). (For those unfamiliar with Stackelberg games, the sidebar article provides an introduction.)

In this Stackelberg game framework, the defender commits to a mixed (randomized) strategy of patrols, which is known to the attacker. This is a reasonable approximation of the practice since the attacker conducts surveillance to learn the mixed strategies that the defender carries out, and responds with a pure strategy of an attack on a target. The optimization objective is to find the optimal mixed strategy for the defender.

In the rest of this section, we begin by discussing how to practically cast this real-world maritime patrolling problem of PWCS patrols as a Stackelberg game. We also show how to reduce the number of defender strategies before addressing the most important innovation in PROTECT: its use of the quantal response model.

Game Modeling

To model the USCG patrolling domain as a Stackelberg game, we need to define (i) the set of attacker strategies, (ii) the set of defender strategies, and (iii) the payoff function. These strategies and payoffs center on the targets in a port — ports, such as the Port of Boston, have a significant number of potential targets (critical infrastructure). As discussed above, our Stackelberg game formulation assumes that the attacker learns the defender’s strategy through conducting surveillance, and can then launch an attack. Thus, the attacks an attacker can launch on different possible targets are considered as his pure strategies.

However, the definition of defender strategies is not as straightforward. Patrols last for some fixed duration during the day as specified by USCG, e.g., 4 hours. Our first attempt was to model each target as a node in a graph and allow patrol paths to go from each individual target to (almost all) other targets in the port, generating an almost complete graph on the targets. This method yields the most flexible set of patrol routes that would fit within the maximum duration, covering any permutation of targets within a single patrol. This method unfortunately faced significant challenges: (i) it required determining the travel time for a patrol boat for each pair of targets, a daunting knowledge acquisition task given the hundreds of pairs of targets; (ii) it did not maximize the use of port geography whereby boat crews could observe multiple targets at once and; (iii) it was perceived as micromanaging the activities of the USCG boat crews, which was undesirable.

Our improved approach to generating defender strategies therefore grouped nearby targets into patrol areas (in real world scenarios such as the Port of Boston, some targets are very close to each other and it is thus natural to group targets together according to their geographic locations). The presence of patrol areas led the USCG to redefine the set of defensive activities to be performed on patrol areas to provide a more accurate and expressive model of the patrols. Activities that take a longer time provide the defender a higher payoff compared to activities that take a shorter time to complete. This impacts the final patrol schedule as one patrol may visit fewer areas but conduct longer duration defensive activities at the areas, while another patrol may have more areas with shorter duration activities.

To generate all the permutations of patrol schedules, a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is created with the patrol areas as vertices \mathcal{V} and adjacent patrol areas as edges \mathcal{E} . Using the graph of patrol areas, PROTECT generates all possible patrol schedules, each of which is a closed walk of \mathcal{G} that starts and ends at the patrol area $b \in \mathcal{V}$, the base patrol area for the USCG. Each patrol schedule is a sequence of patrol areas and associated defensive activities at each patrol area in the sequence, and are constrained by a maximum patrol time τ . (Note that even when the defender just passes by a patrol area, this is treated as an activity.) The defender may visit a patrol area multiple times in a schedule due to geographic constraints and the fact that each patrol is a closed walk. For instance, the defender in each patrol should visit the base patrol area at least twice since she needs to start the patrol from the base and finally come back to the base to finish the

patrol.

The graph \mathcal{G} along with the constraints b and τ are used to generate the defender strategies (patrol schedules). Given each patrol schedule, the total patrol schedule time is calculated (this also includes traversal time between areas, but we ignore it in the following for expository purposes); we then verify that the total time is less than or equal to the maximum patrol time τ . After generating all possible patrol schedules, a game is formed where the set of defender strategies is composed of patrol schedules and the set of attacker strategies is the set of targets. The attacker’s strategy was based on targets instead of patrol areas because an attacker will choose to attack a single target.

Table 1 gives an example, where the rows correspond to the defender’s strategies and the columns correspond to the attacker’s strategies. In this example, there are two possible defensive activities, activity k_1 and k_2 , where k_2 provides more effective protection (also takes more time) for the defender than k_1 . Suppose that the time bound disallows more than two k_2 activities (given the time required for k_2) within a patrol. Patrol area 1 has two targets (target 1 and 2) while patrol areas 2 and 3 each have one target (target 3 and 4 respectively). In the table, a patrol schedule is composed of a sequence of patrol areas and a defensive activity in each area. The patrol schedules are ordered so that the first patrol area in the schedule denotes which patrol area the defender needs to visit first. In this example, patrol area 1 is the base patrol area, and all of the patrol schedules begin and end at patrol area 1. For example, the patrol schedule in row 2 first visits patrol area 1 with activity k_2 , then travels to patrol area 2 with activity k_1 , and returns back to patrol area 1 with activity k_1 . For the payoffs, if a target i is the attacker’s choice and the attack fails, then the defender would gain a reward R_i^d while the attacker would receive a penalty P_i^a , else the defender would receive a penalty P_i^d and the attacker would gain a reward R_i^a . Furthermore, let G_{ij}^d be the payoff for the defender if the defender chooses patrol j and the attacker chooses to attack target i . G_{ij}^d can be represented as a linear combination of the defender reward/penalty on target i and A_{ij} , the effectiveness probability of the defensive activity performed on target i for patrol j , as described by Equation 1. A_{ij} depends on the most effective activity on target i in patrol j . The value of A_{ij} is 0 if target i is not in patrol j . If patrol j only includes one activity in a patrol area that covers target i , then we determine its payoff using the following equation (any additional activity may provide an additional incremental benefit in that area and we discuss this in the following section).

$$G_{ij}^d = A_{ij}R_i^d + (1 - A_{ij})P_i^d \quad (1)$$

For instance, suppose target 1 is covered using k_1 in strategy 5, and the value of A_{15} is 0.5. If $R_1^d = 150$ and $P_1^d = -50$, then $G_{15}^d = 0.5(150) + (1 - 0.5)(-50) = 50$. (G_{ij}^a would be computed in a similar fashion.)

In the USCG problem, rewards and penalties are based on an analysis completed by a contracted company of risk analysts that looked at the targets in the Port of Boston and assigned corresponding values for each one. The types of fac-

tors taken into consideration for generating these values include economic damage and injury/loss of life. Meanwhile, the effectiveness probability, A_{ij} , for different defensive activities are decided based on the duration of the activities. Longer activities lead to a higher possibility of capturing the attackers.

While loss of life and property helps in assessing damage in case of a successful attack, assessing payoffs requires that we determine whether the loss is viewed symmetrically by the defender and attacker. Similarly, whether the payoffs are viewed symmetrically for the attacker and defender also holds for the scenario when there is a failed attack. These questions go to the heart of determining whether security games should be modeled as zero-sum games (Tambe 2011). Past work in security games (e.g., ARMOR (Jain et al. 2010), IRIS (Jain et al. 2010), GUARDS (Pita et al. 2011)) has used non-zero-sum game models, e.g., one assumption made is that the attacker might view publicity of a failed attack as a positive outcome. However, non-zero-sum games require further knowledge acquisition efforts to model the asymmetry in payoffs. For simplicity, as the first step PROTECT starts with the assumption of a zero-sum game. However, the algorithm used in PROTECT is not restricted to zero-sum games as in the future, USCG proposes to relax this assumption. It is also important to note that while Table 1 shows point estimates of payoffs, we recognize that estimates may not be accurate. To that end, in the experimental results section, we evaluated the robustness of our approach when there is payoff noise, observation noise, and execution error.

Patrol Schedule	Target 1	Target 2	Target 3	Target 4
$(1:k_1), (2:k_1), (1:k_1)$	50,-50	...	15,-15	...
$(1:k_2), (2:k_1), (1:k_1)$...	60,-60
$(1:k_1), (2:k_1), (1:k_2)$	-20,20
$(1:k_1), (3:k_1), (2:k_1), (1:k_1)$	50,-50
$(1:k_1), (2:k_1), (3:k_1), (1:k_1)$	15,-15	...

Table 1: Portion of a simplified example of a game matrix

Compact Representation

In our game, the number of defender strategies, i.e., patrol schedules, grows combinatorially, generating a scale-up challenge. To achieve scale-up, PROTECT uses a compact representation of the patrol schedules using two ideas: (i) combining equivalent patrol schedules and; (ii) removal of dominated patrol schedules.

With respect to equivalence, different permutations of patrol schedules provide identical payoff results. Furthermore, if an area is visited multiple times with different activities in a schedule, we only consider the activity that provides the defender the highest payoff, not the incremental benefit due to additional activities. This decision is made in consideration of the tradeoff between modeling accuracy and efficiency. On the one hand, the additional value of more activities is small. Currently, the patrol time of each schedule is relatively short (e.g., 1 hour) and the defender may visit a patrol area more than once within the short period and will conduct

an activity each time. For instance, the defender may pass by a patrol area 10 minutes after conducting a more effective activity at the same patrol area. The additional value of the pass by activity given the more effective activity is therefore very small. On the other hand, it leads to significant computational benefits which are described in this section if we just consider the most effective activity in each patrol.

Therefore, many patrol schedules are equivalent if the set of patrol areas visited and the most effective defensive activities in each patrol area in the schedules are the same even if their order differs. Such equivalent patrol schedules are combined into a single compact defender strategy, represented as a set of patrol areas and defensive activities (and minus any ordering information). The idea of combining equivalent actions is similar to action abstraction for solving large scale dynamic games (Gilpin 2009). Table 2 presents a compact version of Table 1, which shows how the game matrix is simplified by using equivalence to form compact defender strategies, e.g., the patrol schedules in the rows 2-3 from Table 1 are represented as a compact strategy $\Gamma_2 = \{(1,k_2), (2,k_1)\}$ in Table 2.

Compact Strategy	Target 1	Target 2	Target 3	Target 4
$\Gamma_1 = \{(1:k_1), (2:k_1)\}$	50,-50	30,-30	15,-15	-20,20
$\Gamma_2 = \{(1:k_2), (2:k_1)\}$	100,-100	60,-60	15,-15	-20,20
$\Gamma_3 = \{(1:k_1), (2:k_1), (3:k_1)\}$	50,-50	30,-30	15,-15	10,-10

Table 2: Example compact strategies and game matrix

Next, the idea of dominance is illustrated using Table 2 and noting the difference between Γ_1 and Γ_2 is the defensive activity on patrol area 1. Since activity k_2 gives the defender a higher payoff than k_1 , Γ_1 can be removed from the set of defender strategies because Γ_2 covers the same patrol areas while giving a higher payoff for patrol area 1. To generate the set of compact defender strategies, a naive approach would be to first generate the full set of patrol schedules and then prune the dominated and equivalent schedules. Instead, PROTECT generates compact strategies in the following manner: we start with patrols that visit the most patrol areas with the least effective activities within the patrol time limit; these activities take a shorter amount of time but we can cover more areas within the given time limit. Then we gradually consider patrols visiting less patrol areas but with increasingly effective activities. This process will stop when we have considered all patrols in which all patrol areas are covered with the most effective activities and cannot include any additional patrol area.

Figure 2 shows a high level view of the steps of the algorithm using the compact representation. The compact strategies are used instead of full patrol schedules to generate the game matrix. Once the optimal probability distribution is calculated for the compact strategies, the strategies with a probability greater than 0 are expanded to a complete set of patrol schedules.

In this expansion from a compact strategy to a full set of patrol schedules, we need to determine the probability of choosing each patrol schedule, since a compact strategy may correspond to multiple patrol schedules. The focus here is




Figure 2: Flow chart of the PROTECT system

to increase the difficulty for the attacker to conduct surveillance by increasing unpredictability¹, which we achieve by randomizing uniformly over all expansions of the compact defender strategies. The uniform distribution provides the maximum entropy (greatest unpredictability). Thus, all the patrol schedules generated from a single compact strategy are assigned a probability of v_i/w_i where v_i is the probability of choosing a compact strategy Γ_i and w_i is the total number of expanded patrol schedules for Γ_i . The complete set of patrol schedules and the associated probabilities are then sampled and provided to the USCG, along with the start time of the patrol generated via uniform random sampling.

Human Adversary Modeling

While previous game-theoretic security applications have assumed a perfectly rational attacker, PROTECT takes a step forward by addressing this limitation of classical game theory. Instead, PROTECT uses a model of a boundedly rational adversary by using a quantal response (QR) model of an adversary, which has shown to be a promising model of human decision making (McKelvey and Palfrey 1995; Rogers, Palfrey, and Camerer 2009; Yang et al. 2011). A recent study demonstrated the use of QR as an effective prediction model of humans (Wright and Leyton-Brown 2010). An even more relevant study of the QR model was conducted by Yang et al. (Yang et al. 2011) in the context of security games where this model was shown to outperform competitors in modeling human subjects. Based on this evidence, PROTECT uses a QR model of a human adversary, i.e., in the Stackelberg game model, the attacker best-responds according to a QR model and the defender computes her optimal mixed patrolling strategy with this knowledge.

To apply the QR model in a Stackelberg framework, PROTECT employs an algorithm known as PASAQ (Yang, Tambe, and Ordonez 2012). PASAQ computes the optimal defender strategy (within a guaranteed error bound) given a QR model of the adversary by solving the following non-linear and non-convex optimization problem P , with Table 3

¹Creating optimal Stackelberg defender strategies that increase the attacker’s difficulty of surveillance is an open research issue in the literature; here we choose to maximize unpredictability as the first step.

t_i	Target i
R_i^d	Defender reward on covering t_i if it’s attacked
P_i^d	Defender penalty on not covering t_i if it’s attack
R_i^a	Attacker reward on attacking t_i if it’s not covered
P_i^a	Attacker penalty on attacking t_i if it’s covered
A_{ij}	Effectiveness probability of compact strategy Γ_j on t_i
a_j	Probability of choosing compact strategy Γ_j
J	Total number of compact strategies
x_i	Marginal coverage on t_i

Table 3: PASAQ notation as applied to PROTECT

listing the notation:

$$P : \begin{cases} \max_a \sum_{i=1}^T QR(i|\lambda)((R_i^d - P_i^d)x_i + P_i^d) \\ x_i = \sum_{j=1}^J a_j A_{ij}, \quad \forall i \\ \sum_{j=1}^J a_j = 1 \\ 0 \leq a_j \leq 1, \quad \forall j \end{cases}$$

The first line of the problem corresponds to the computation of the defender’s expected utility resulting from a combination of Equations 1 and 2 (in the sidebar). $QR(i|\lambda)$ is the probability of that the attacker using the QR model will attack target t_i . Unlike previous application (Jain et al. 2010; Kiekintveld, Marecki, and Tambe 2011; Paruchuri et al. 2008), x_i in this case not just summarizes presence or absence on a target, but also the effectiveness probability A_{ij} on the target as well. That is, the second line computes the marginal coverage on the targets based on the effectiveness factor A_{ij} and the probability of choosing compact strategy Γ_j , denoted as a_j .

As with all QR models, a value for λ is needed to represent the noise in the attacker’s strategy. Clearly, a λ value of 0 (uniform random) and ∞ (fully rational) are not reasonable. Given the payoff data for Boston, an attacker’s strategy with $\lambda = 4$ starts approaching a fully rational attacker — the probability of attack focuses on a single target. In addition, an attacker’s strategy with $\lambda = 0.5$ is similar to a fully random strategy that uniformly chooses a target to attack. USCG experts (with expertise in terrorist behavior modeling) suggested that we could use a broad range for representing possible λ values used by the attacker. Combining the above observations, it was determined that the attacker’s strategy is best modeled with a λ value that is in the range $[0.5, 4]$, rather than a single point estimate. A discrete sampling approach was used to determine a λ value that gives the highest average defender expected utility across attacker strategies within this range to get $\lambda = 1.5$. Specifically, the defender considers different assumptions of the attacker’s λ value and for each assumption about the λ value, the defender computes her expected utility against the attacker with different λ values within the range $[0.5, 4]$. We find

that when the defender assumes the attacker is using the QR model with $\lambda = 1.5$, the defender’s strategy leads to the highest defender expected utility when the attacker follows the QR model with a λ value uniformly randomly chosen from the range of $[0.5, 4]$. Selecting an appropriate value for λ remains a complex issue however, and it is a key agenda item for future work.

Evaluation

This section presents evaluations based on (i) experiments completed via simulations and (ii) real-world patrol data along with USCG analysis. All scenarios and experiments, including the payoff values and graph (composed of 9 patrol areas), were based off the Port of Boston. The defender’s payoff values have a range of $[-10,5]$ while the attacker’s payoff values have a range of $[-5,10]$. The game was modeled as a zero-sum game in which the attacker’s loss or gain is balanced precisely by the defender’s gain or loss. For PASAQ, the defender’s strategy is computed assuming that the attacker follows the QR model with $\lambda = 1.5$ as justified in the previous section. All experiments are run on a machine with an Intel Dual Core 1.4 GHz processor and 2 GB of RAM.

Memory and Run-time Analysis

This section presents the results based on simulation to show the efficiency in memory and run-time of the compact representation versus the full representation. In Figure 3(a), the x-axis is the maximum patrol time allowed and the y-axis is the memory needed to run PROTECT. In Figure 3(b), the x-axis is the maximum patrol time allowed and the y-axis is the run-time of PROTECT. The maximum patrol time allowed determines the number of combinations of patrol areas that can be visited — so the x-axis indicates a scale-up in the number of defender strategies. When the maximum patrol time is set to 90 minutes, the full representation takes 30 seconds and uses 540 MB of memory while the compact representation takes 11 seconds to run and requires 20 MB of memory. Due to the exponential increase in the memory and run-time that is needed for the full representation, it cannot be scaled up beyond 90 minutes.




Figure 3: Comparison of full vs. compact representation

Utility Analysis

It is useful to understand whether PROTECT using PASAQ with $\lambda = 1.5$ provides an advantage when compared to: (i) a uniform random defender’s strategy; (ii) a mixed strategy

with the assumption of the attacker attacking any target uniformly at random ($\lambda = 0$) or; (iii) a mixed strategy assuming a fully rational attacker ($\lambda = \infty$). The previously existing DOBSS algorithm was used for $\lambda = \infty$ (Paruchuri et al. 2008). Additionally, comparison with the $\lambda = \infty$ approach is important because of the extensive use of this assumption in previous applications (for our zero-sum case, DOBSS is equivalent to minimax but the utility does not change). Typically, we may not have an estimate of the exact value of the attacker’s λ value, only a possible range. Therefore, ideally we would wish to show that PROTECT (using $\lambda = 1.5$ in computing the optimal defender strategy) provides an advantage over a range of λ values assumed for the attacker (not just over a point estimate) in his best-response, justifying our use of the PASAQ algorithm. In other words, we are distinguishing between (i) the actual λ value employed by the attacker in best-responding, and (ii) the λ assumed by PASAQ in computing the defender’s optimal mixed strategy. The point is to see how sensitive the choice of (ii) is, with respect to prevailing uncertainty about (i).




Figure 4: Defender’s Expected Utility when varying λ for attacker’s strategy(color)

To achieve this, we compute the average defender utility of the four approaches above as the λ value of the attacker’s strategy changes from $[0, 6]$, which subsumes the range $[0.5, 4]$ of reasonable attacker strategies. In Figure 4, the y-axis represents the defender’s expected utility and the x-axis is the λ value that is used for the attacker’s strategy. Both uniform random strategies perform well when the attacker’s strategy is based on $\lambda = 0$. However, as λ increases, both strategies quickly drop to a very low defender expected utility. In contrast, the PASAQ strategy with $\lambda = 1.5$ provides a higher expected utility than that assuming a fully rational attacker over a range of attacker λ values (and indeed over the range of interest), not just at $\lambda = 1.5$.

Robustness Analysis

In the real world, observation, execution, and payoffs, are not always perfect due to the following: noise in the attacker’s surveillance of the defender’s patrols, the many tasks and responsibilities of the USCG where the crew may be pulled off a patrol, and limited knowledge of the attacker’s payoff values. Our hypothesis is that PASAQ with $\lambda = 1.5$ is more robust to such noise than a defender strategy which assumes full rationality of the attacker such as DOBSS (Paruchuri et al. 2008), i.e., PASAQ’s expected defender

utility will not degrade as much as DOBSS over the range of attacker λ of interest. This is illustrated by comparing both PASAQ and DOBSS against observation, execution, and payoff noise (Kiekintveld, Marecki, and Tambe 2011; Korzhyk, Conitzer, and Parr 2011; Yin et al. 2011). Intuitively, the QR model is more robust than models assuming perfect rationality since the QR model assumes that the attacker may attack multiple targets with positive probabilities, rather than attacking a single target in the model assuming perfect rationality of the adversaries. Such intuition has been verified in other contexts (Rogers, Palfrey, and Camerer 2009). (A comparison of the uniform random strategies was not included due to its poor performance shown in Figure 4.) All experiments were run generating 200 samples with added noise and averaging over all the samples.

Figure 5 shows the performance of different strategies while considering execution noise. The y-axis represents the defender's expected utility and the x-axis is the attacker's λ value. If the defender covered a target with probability p , this probability now changes to be in $[p - x, p + x]$ where x is the noise. The low execution error corresponds to $x = 0.1$ whereas high error corresponds to $x = 0.2$. In the experiments, the attacker best-responds to the mixed strategy with added noise. The key takeaway here is that execution error leads to PASAQ dominating DOBSS over all tested values of λ , further strengthening the reason to use PASAQ rather than a full-rationality model. For both algorithms, the defender's expected utility decreases as more execution error is added because the defender's strategy is impacted by the additional error. When execution error is added, PASAQ dominates DOBSS because the latter seeks to maximize the minimum defender's expected utility so multiple targets will have the same minimum defender utility. For DOBSS, when execution error is added, there is a greater probability that one of these targets will have less coverage, resulting in a lower defender's expected utility. For PASAQ, typically only one target has the minimum defender expected utility. As a result changes in coverage do not impact it as much as DOBSS. As execution error increases, the advantage in the defender's expected utility of PASAQ over DOBSS increases even more. This section only shows the execution noise results; the details of the observation and payoff noise results can be found in (Shieh et al. 2012).




Figure 5: Defender's expected utility: Execution noise(color)

USCG Real-World Evaluation

In addition to the data made available from simulations, the USCG conducted its own real-world evaluation of PROTECT. With permission, some aspects of the evaluation are presented in this paper.

Real-world scheduling data: Unlike prior publications of real-world applications of game theory for security, a key novelty of this paper is the inclusion of actual data from USCG patrols before and after the deployment of PROTECT at the Port of Boston. Figure 6 and Figure 7 show the frequency of visits by USCG to different patrol areas over a number of weeks. Figure 6 shows pre-PROTECT patrol visits per day by area and Figure 7 shows post-PROTECT patrol visits per day by area. The x-axis is the day of the week, and the y-axis is the number of times a patrol area is visited for a given day of the week. The y-axis is intentionally blurred for security reasons as this is real data from Boston. There are more lines in Figure 6 than in Figure 7 because during the implementation of PROTECT, new patrol areas were formed which contained more targets and thus fewer patrol areas in the post-PROTECT figure. Figure 6 depicts a definite pattern in the patrols. While there is a spike in patrols executed on Day 5, there is a dearth of patrols on Day 2. Besides this pattern, the lines in Figure 6 intersect, indicating that some days, a higher value target was visited more often while on other days it was visited less often. This means that there was not a consistently high frequency of coverage of higher value targets before PROTECT.

In Figure 7, we notice that the pattern of low patrols on Day 2 (from Figure 6) disappears. Furthermore, lines do not frequently intersect, i.e., higher valued targets are visited consistently across the week. The top line in Figure 7 is the base patrol area and is visited at a higher rate than all other patrol areas.




Figure 6: Patrol visits per day by area - pre-PROTECT(Color), one line per area

Adversary Perspective Teams(APT): To obtain a better understanding of how the adversary views the potential targets in the port, the USCG created the Adversarial Perspective Team (APT), a mock attacker team. The APT provides assessments from the terrorist perspective and as a secondary function, assesses the effectiveness of the patrol activities before and after deployment of PROTECT. In their evaluation, the APT incorporates the adversary's known intent, capabilities, skills, commitment, resources, and cultural influences. In addition, it screens attack possibilities




Figure 7: Patrol visits per day by area - post-PROTECT(Color), one line per area

and assists in identifying the level of deterrence projected at and perceived by the adversary. For the purposes of this research, the adversary is defined as an individual(s) with ties to al-Qa'ida or its affiliates.

The APT conducted a pre- and post-PROTECT assessment of the system's impact on an adversary's deterrence at the Port of Boston. This analysis uncovered a positive trend where the effectiveness of deterrence increased from the pre- to post- PROTECT observations.

Additional Real-world Indicators: The use of PROTECT and APT's improved guidance given to boat crews on how to conduct the patrol jointly provided a noticeable increase in the quality and effectiveness of the patrols. Prior to implementing PROTECT, there were no documented reports of illicit activity. After implementation, USCG crews, reported more illicit activities within the port (therefore justify the effectiveness of the PROTECT model) and provided a noticeable "on the water" presence with industry port partners commenting, "the Coast Guard seems to be everywhere, all the time." With no actual increase in the number of resources applied, and therefore no increase in capital or operating costs, these outcomes support the practical application of game theory in the maritime security environment.

Outcomes after Boston Implementation

After evaluating the performance and impact of PROTECT at Boston, the USCG viewed this system as a success. As a result, PROTECT is now getting deployed in the Port of New York. We were presented an award for the work on the PROTECT system for the Boston Harbor which reflects USCG's recognition of the impact and value of PROTECT.

Lessons Learned: Putting Theory into Practice

Developing the PROTECT model was a collaborative effort involving university researchers and USCG personnel representing decision makers, planners and operators. Building on the lessons reported in (Pita et al. 2011) for working with security organizations, we informed the USCG of (i) the assumptions underlying the game-theoretic approaches, e.g., full adversary rationality, and strengths and limitations of different algorithms — rather than pre-selecting a simple heuristic approach; (ii) the need to define and collect correct inputs for model development and; (iii) a fundamental understanding of how the inputs affect the results. We gained

three new insights involving real-world applied research; (i) unforeseen positive benefits because security agencies were compelled to reexamine their assumptions; (ii) requirement to work with multiple teams in a security organization at multiple levels of their hierarchy and; (iii) need to prepare answers to end-user practical questions not always directly related to the "meaty" research problems.

The first insight came about when USCG was compelled to reassess their operational assumptions as a result of working through the research problem. A positive result of this reexamination prompted USCG to develop new PWCS mission tactics, techniques and procedures. Through the iterative development process, USCG reassessed the reasons why boat crews performed certain activities and whether they were sufficient. For example, instead of "covered" vs "not covered" as the only two possibilities at a patrol point, there are now multiple sets of activities at each patrol point.

The second insight is that applied research requires the research team to collaborate with planners and operators on the multiple levels of a security organization to ensure the model accounts for all aspects of a complex real world environment. Initially when we started working on PROTECT, the focus was on patrolling each individual target. This appeared to micromanage the activities of boat crews, and it was through their input that individual targets were grouped into patrol areas associated with a PWCS patrol. On the other hand, input from USCG headquarters and the APT mentioned earlier, led to other changes in PROTECT, e.g., departing from a fully rational model of an adversary to a QR model.

The third insight is the need to develop answers to end-user questions which are not always related to the "meaty" research question but are related to the larger knowledge domain on which the research depends. One example of the need to explain results involved the user citing that one patrol area was being repeated and hence, randomization did not seem to occur. After assessing this concern, we determined that the cause for the repeated visits to a patrol area was its high reward — order of magnitude greater than the rarely visited patrol areas. PROTECT correctly assigned patrol schedules that covered the more "important" patrol areas more frequently. In another example, the user noted that PROTECT did not assign any patrols to start at 4:00 AM or 4:00 PM over a 60 day test period. They expected patrols would be scheduled to start at any hour of the day, leading them to ask if there was a problem with the program. This required us to develop a layman's briefing on probabilities, randomness, and sampling. With 60 patrol schedules, a few start hours may not be chosen given our uniform random sampling of the start time. These practitioner-based issues demonstrate the need for researchers to not only be conversant in the algorithms and math behind the research, but also be able to explain from a user's perspective how solutions are accurate. An inability to address these issues would result in a lack of real-world user confidence in the model.

Summary and Related Work

This paper reports on PROTECT, a game-theoretic system deployed by the USCG in the Port of Boston since April

2011 for scheduling their patrols. USCG has deemed the deployment of PROTECT in Boston a success and efforts are underway to deploy PROTECT in the Port of New York, and to other ports in the United States. PROTECT uses an attacker-defender Stackelberg game model, and includes five key innovations.

First, PROTECT moves away from the assumption of perfect adversary rationality seen in previous work, relying instead on a quantal response (QR) model of the adversary's behavior. While the QR model has been extensively studied in the realm of behavioral game theory, to the best of our knowledge, this is its first real-world deployment. Second, to improve PROTECT's efficiency, we generate a novel compact representation of the defender's strategy space, exploiting equivalence and dominance. Third, the paper shows how to practically model a real-world (maritime) patrolling problem as a Stackelberg game. Fourth, we provide experimental results illustrating that PROTECT's QR model of the adversary is better able to handle real-world uncertainties than a perfect rationality model. Finally, for the first time in a security application evaluation, we use real-world data: (i) providing a comparison of human-generated security schedules versus those generated via a game-theoretic algorithm and; (ii) results from an APT's analysis of the impact of the PROTECT system. The paper also outlined the insights from the project which include the ancillary benefits due to a review of assumptions made by security agencies, and the need for knowledge to answer questions not directly related to the research problem.

As a result, PROTECT has advanced the state of the art beyond previous applications of game theory for security. Prior applications mentioned earlier, including ARMOR, IRIS or GUARDS (Jain et al. 2010; Pita et al. 2011), have each provided unique contributions in applying novel game-theoretic algorithms and techniques. Interestingly, these applications have revolved around airport and air-transportation security. PROTECT's novelty is not only its application domain in maritime patrolling, but also in the five key innovations mentioned above, particularly its emphasis on moving away from the assumption of perfect rationality by using the QR model.

In addition to game-theoretic applications, the issue of patrolling has received significant attention in the multiagent systems literature. These include patrol work done by robots primarily for perimeter patrols that have been addressed in arbitrary topologies (Basilico, Gatti, and Amigoni 2009), maritime patrols in simulations for deterring pirate attacks (Vanek et al. 2011), and in research looking at the impact of uncertainty in adversarial behavior (Agmon et al. 2009). PROTECT differs from these approaches in its use of a QR model of a human adversary in a game theoretic setting, and in being a deployed application. Building on this initial success of PROTECT, we hope to deploy it at more and much larger-sized ports. In so doing, in the future, we will consider significantly more complex attacker strategies, including potential real-time surveillance and coordinated attacks.

Acknowledgments

We thank the USCG offices, and particularly sector Boston, for their exceptional collaboration. The views expressed herein are those of the author(s) and are not to be construed as official or reflecting the views of the Commandant or of the U.S. Coast Guard. This research was supported by the United States Department of Homeland Security through the National Center for Risk and Economic Analysis of Terrorism Events (CREATE) under award number 2010-ST-061-RE0001. This paper extends and modifies our previous conference publication (Shieh et al. 2012).

References

- Agmon, N.; Kraus, S.; Kaminka, G. A.; and Sadvov, V. 2009. Adversarial uncertainty in multi-robot patrol. In *Proc. of The 21st International Joint Conference on Artificial Intelligence (IJCAI)*, 1811–1817.
- Basilico, N.; Gatti, N.; and Amigoni, F. 2009. Leader-follower strategies for robotic patrolling in environments with arbitrary topologies. In *Proc. of The 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 500–503.
- Blair, D. 2010. Annual threat assessment of the us intelligence community for the senate select committee on intelligence. http://www.dni.gov/testimonies/20100202_testimony.pdf.
- Camerer, C. F. 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press.
- Conitzer, V., and Sandholm, T. 2006. Computing the optimal strategy to commit to. In *Proc. of the 7th ACM conference on Electronic commerce*, 82–90.
- Dozier, K. 2011. Bin laden trove of documents sharpen US aim. http://www.msnbc.msn.com/id/43331634/ns/us_news-security/t/bin-laden-trove-documents-sharpen-us-aim/.
- Fudenberg, D., and Tirole, J. 1991. *Game Theory*. MIT Press.
- Gilpin, A. 2009. *Algorithms for abstracting and solving imperfect information games*. Ph.D. Dissertation, Carnegie Mellon University.
- Jain, M.; Tsai, J.; Pita, J.; Kiekintveld, C.; Rathi, S.; Tambe, M.; and Ordóñez, F. 2010. Software assistants for randomized patrol planning for the LAX airport police and the federal air marshal service. *Interfaces* 40:267–290.
- Kiekintveld, C.; Marecki, J.; and Tambe, M. 2011. Approximation methods for infinite bayesian Stackelberg games: modeling distributional uncertainty. In *Proc. of The 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 1005–1012.
- Korzhyk, D.; Conitzer, V.; and Parr, R. 2011. Solving Stackelberg games with uncertain observability. In *Proc. of The 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 1013–1020.
- McKelvey, R. D., and Palfrey, T. R. 1995. Quantal response equilibria for normal form games. *Games and Economic Behavior* 10(1):6–38.

Paruchuri, P.; Pearce, J. P.; Marecki, J.; Tambe, M.; Ordóñez, F.; and Kraus, S. 2008. 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)*, 895–902.

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

Rogers, B. W.; Palfrey, T. R.; and Camerer, C. F. 2009. Heterogeneous quantal response equilibrium and cognitive hierarchies. *Journal of Economic Theory* 144:1440–1467.

Shieh, E.; An, B.; Yang, R.; Tambe, M.; Baldwin, C.; DiRenzo, J.; Maule, B.; and Meyer, G. 2012. 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)*.

Tambe, M. 2011. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press.

Vanek, O.; Jakob, M.; Hrstka, O.; and Pechoucek, M. 2011. Using multi-agent simulation to improve the security of maritime transit. In *Proc. of 12th International Workshop on Multi-Agent-Based Simulation (MABS)*.

von Stengel, B., and Zamir, S. 2004. Leadership with commitment to mixed strategies. Technical Report LSE-CDAM-2004-01, CDAM Research Report.

Wright, J., and Leyton-Brown, K. 2010. Beyond equilibrium: Predicting human behavior in normal form games. In *Proc. of The 24th AAAI Conference on Artificial Intelligence (AAAI)*, 901–907.

Yang, R.; Kiekintveld, C.; Ordóñez, F.; Tambe, M.; and John, R. 2011. Improving resource allocation strategy against human adversaries in security games. In *Proc. of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, 458–464.

Yang, R.; Tambe, M.; and Ordóñez, F. 2012. Computing optimal strategy against quantal response in security games. In *Proc. of The 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

Yin, Z.; Korzhyk, D.; Kiekintveld, C.; Conitzer, V.; ; and Tambe, M. 2010. Stackelberg vs. nash in security games: interchangeability, equivalence, and uniqueness. In *Proc. of The 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 1139–1146.

Yin, Z.; Jain, M.; Tambe, M.; and Ordóñez, F. 2011. Risk-averse strategies for security games with execution and observational uncertainty. In *Proc. of The 25th AAAI Conference on Artificial Intelligence (AAAI)*, 758–763.

Young, S., and Orchard, D. 2011. Remembering 9/11: Protecting America's ports. <http://coastguard.dodlive.mil/2011/09/remembering-911-protecting-americas-ports/>.

Appendices – Mini-tutorials in Separate Boxes

Stackelberg Game

A generic Stackelberg game has two players, a *leader*, and a *follower* (Fudenberg and Tirole 1991). A leader commits to a strategy first, and then a follower optimizes her reward, considering the action chosen by the leader (von Stengel and Zamir 2004). The two players in a Stackelberg game need not represent individuals, but could also be groups that cooperate to execute a joint strategy, such as a police force or a terrorist organization. Each player has a set of possible *pure strategies*, or the actions that they can execute. A *mixed strategy* allows a player to play a probability distribution over pure strategies. Payoffs for each player are defined over all possible pure-strategy outcomes for both the players. The payoff functions are extended to mixed strategies by taking the expectation over pure-strategy outcomes. The follower can observe the leader’s strategy, and then act in a way to optimize his own payoffs.

To see the advantage of being the leader in a Stackelberg game, consider the game with the payoff as shown in Table 4. The leader is the row player and the follower is the column player. The only pure-strategy Nash equilibrium for this game is when the leader plays *a* and the follower plays *c* which gives the leader a payoff of 3; in fact, for the leader, playing *b* is strictly dominated.

However, in the simultaneous game if the leader can commit to playing *b* before the follower chooses his strategy, then the leader will obtain a payoff of 4, since the follower would then play *d* to ensure a higher payoff for himself. If the leader commits to a mixed strategy of playing *a* and *b* with equal (0.5) probability, then the follower will play *d*, leading to a higher expected payoff for the leader of 4.5. As we can see from this example, the equilibrium strategy in the Stackelberg game can be in fact different from the Nash equilibria.

	c	d
a	3,1	5, 0
b	2,0	4, 2

Table 4: Payoff table for example Stackelberg game.

Stackelberg games are used to model the attacker-defender strategic interaction in security domains and this class of Stackelberg games (with certain restrictions on payoffs (Yin et al. 2010)) is called *Stackelberg security games*. In the Stackelberg security game framework, the security force (defender) is modeled as the leader and the terrorist adversary (attacker) is in the role of the follower. The defender commits to a mixed (randomized) strategy, whereas the attacker conducts surveillance of these mixed strategies and responds with a pure strategy of an attack on a target. Thus, the Stackelberg game framework is a natural approximation of the real-world security scenarios. In contrast, the surveillance activity of the attacker cannot be modeled in the simultaneous move games with the Nash equilibrium solution concept. The objective is to find the optimal mixed strategy for the defender. See (Tambe 2011) for a more detailed introduction to research on Stackelberg security games.

Quantal Response

Quantal Response Equilibrium is an important model in behavior game theory that has received widespread support in the literature in terms of its superior ability to model human behavior in simultaneous-move games (McKelvey and Palfrey 1995) and is the baseline model of many studies (Wright and Leyton-Brown 2010). It suggests that instead of strictly maximizing utility, individuals respond stochastically in games: the chance of selecting a non-optimal strategy increases as the cost of such an error decreases. However, the applicability of these models to Stackelberg security games had not been explored previously. In Stackelberg security games, we assume that the attacker acts with bounded rationality; the defender is assisted by software and thus we compute the defender’s optimal rational strategy (Yang et al. 2011).

Given the strategy of the defender, the Quantal Best Response of the attacker is defined as

$$q_i = \frac{e^{\lambda G_i^a(x_i)}}{\sum_{j=1}^T e^{\lambda G_j^a(x_i)}} \quad (2)$$

The parameter λ represents the amount of noise in the attacker’s strategy. λ can range from 0 to ∞ with a value of 0 representing a uniform random probability over attacker strategies while a value of ∞ representing a perfectly rational attacker. q_i corresponds to the probability that the attacker chooses a target i ; $G_i^a(x_i)$ corresponds to the attacker’s expected utility of attacking target i given x_i , the probability that the defender covers target i ; and T is the total number of targets.

Consider the Stackelberg game with the payoffs shown in Table 4. Assume that the leader commits to playing *b*. The follower obtains a payoff of 0 by playing *c* and obtains a payoff of 2 by playing *d*. A rational attacker will play *d* to maximize his payoff. However, the quantal best response of the attacker would be playing *c* with probability $\frac{e^0}{e^0 + e^{2\lambda}}$ and playing *d* with probability $\frac{e^{2\lambda}}{e^0 + e^{2\lambda}}$.