Risk-Averse Strategies for Security Games with Execution and Observational Uncertainty

Zhengyu Yin and **Manish Jain** and **Milind Tambe**
Computer Science Department
Univerisity of Southern California
Los Angeles, CA 90089
{zhengyuy, manish.jain, tambe}@usc.edu

Fernando Ordóñez
Industrial and Systems Engineering
University of Southern California
Los Angeles, CA 90089
Industrial Engineering Department
University of Chile (Santiago)
fordon@usc.edu

Abstract

Attacker-defender Stackelberg games have become a popular game-theoretic approach for security with deployments for LAX Police, the FAMS and the TSA. Unfortunately, most of the existing solution approaches do not model two key uncertainties of the real-world: there may be noise in the defender's execution of the suggested mixed strategy and/or the observations made by an attacker can be noisy. In this paper, we provide a framework to model these uncertainties, and demonstrate that previous strategies perform poorly in such uncertain settings. We also provide RECON, a novel algorithm that computes strategies for the defender that are robust to such uncertainties, and provide heuristics that further improve RECON's efficiency.

Introduction

The use of game-theoretic concepts has allowed security forces to exert maximum leverage with limited resources. Indeed, game-theoretic scheduling softwares have been assisting the LAX police, the Federal Air Marshals service, and are under consideration by the TSA (Jain et al. 2010). They have been studied for patrolling (Agmon et al. 2008; Basilico, Gatti, and Amigoni 2009) and routing in networks (Kodialam and Lakshman 2003).At the backbone of these applications are attacker-defender Stackelberg games. The solution concept is to compute a strong Stackelberg equilibrium (SSE) (von Stengel and Zamir 2004; Conitzer and Sandholm 2006); specifically, the optimal mixed strategy for the defender. It is then assumed that the defender perfectly executes her SSE strategy and the attacker perfectly observes the strategy before choosing his action.

Unfortunately, in the real-world, execution and observation is not perfect due to unforeseen circumstances and/or human errors. For example, a canine unit protecting a terminal at LAX may be urgently called off to another assignment or alternatively a new unit could become available. Similarly, the attacker's observations can be noisy: he may occasionally not observe an officer patrolling a target, or mistake a passing car as a security patrol. Thus, in real-world deployments, the defender may have a noisy execution and the attacker's observations may be even more noisy. A naïve defender strategy can be arbitrarily bad in such uncertain domains, and thus, it is important to develop risk-averse solution techniques that can be used by deployed security systems like ARMOR and IRIS (Jain et al. 2010).

This paper models execution and observation uncertainty, and presents efficient solution techniques to compute risk-averse defender strategies. Previous work has failed to provide such efficient solution algorithms. While the COBRA algorithm (Pita et al. 2010) focuses on human subjects' inference when faced with limited observations of defender strategies, it does not consider errors in such observations. In contrast, Yin et. al (2010) consider the limiting case where an attacker has no observations and thus investigate the equivalence of Stackelberg vs Nash equilibria. Even earlier investigations have emphasized the value of commitment to mixed strategies in Stackelberg games in the presence of noise (van Damme and Hurkens 1997). Outside of Stackelberg games, models for execution uncertainty have been separately developed (Archibald and Shoham 2009). Our research complements these efforts by providing a unified efficient algorithm that addresses both execution and observation uncertainties; in this way it also complements other research that addresses payoff uncertainties in such games (Jain et al. 2010).

This paper provides three key contributions: (1) We provide RECON, a mixed-integer linear program that computes the risk-averse strategy for the defender in domains with execution and observation uncertainty. RECON assumes that nature chooses noise to maximally reduce defenders utility, and RECON maximizes against this worst case. (2) We provide two novel heuristics that speed up the computation of RECON by orders of magnitude. (3) We present experimental results that demonstrate the superiority of RECON in uncertain domains where existing algorithms perform poorly.

Background and Notation

A Stackelberg security game is a game between two players: a defender and an attacker. The defender wishes to deploy up to γ security resources to protect a set of targets T from the attacker. Each player has a set of *pure strategies*: the defender can *cover* a set of targets, and the attacker can *attack* one target. The payoffs for each player depend on the target attacked, and whether or not the attack was successful. $U_d^u(t_i)$ and $U_d^c(t_i)$ represent the utilities for the defender when t_i, the target attacked, was uncovered and covered re-

spectively. The attacker's utilities are denoted similarly by $U_a^c(t_i)$ and $U_a^u(t_i)$. We use $\Delta U_d(t_i) = U_d^c(t_i) - U_d^u(t_i)$ to denote the difference between defender's covered and uncovered utilities. Similarly, $\Delta U_a(t_i) = U_a^u(t_i) - U_a^c(t_i)$. A key property of security games is that $\Delta U_d(t_i) > 0$ and $\Delta U_a(t_i) > 0$. Example payoffs for a game with two targets, t_1 and t_2, and one resource are given in Table 1.

Target	$U_d^c(t_i)$	$U_d^u(t_i)$	$U_a^c(t_i)$	$U_a^u(t_i)$
t_1	10	0	-1	1
t_2	0	-10	-1	1

Table 1: Example game with two targets

A strategy profile $\langle \mathbf{x}, t_i \rangle$ for this game is a mixed strategy $\mathbf{x}$ for the defender, and the attacked target t_i. The mixed strategy $\mathbf{x} = \langle x_i \rangle$ is a vector of probabilities of defender *coverage* over all targets (Yin et al. 2010), such that the sum total of coverage is not more than the number of available resources γ. For example, a mixed strategy for the defender can be .25 coverage on t_1 and .75 coverage on t_2. We allow y_i, the defender's actual coverage on t_i, to vary from the intended coverage x_i by the amount α_i, that is, $|y_i - x_i| \leq \alpha_i$. Thus, if we set $\alpha_1 = 0.1$, it would mean that $0.15 \leq y_1 \leq 0.35$. Additionally, we assume that the attacker wouldn't necessarily observe the actual implemented mixed strategy of the defender; instead the attacker's perceived coverage for t_i, denoted by z_i, can vary by β_i from the implemented coverage y_i. Therefore, $|z_i - y_i| \leq \beta_i$. Thus, in our example, if y_1 was 0.3 and β_1 was set to 0.05, then $0.25 \leq z_1 \leq 0.35$. Table 2 summarizes the notation used in this paper.

For example, at LAX, ARMOR might generate a schedule for two canines to patrol Terminals 1, 2, 3, 4 with probabilities of 0.2, 0.8, 0.5, 0.5 respectively. However, a last-minute cargo inspection may require a canine unit to be called away from, say, Terminal 2 in its particular patrol, or an extra canine unit may become available by chance and get sent to Terminal 3. Additionally, an attacker may fail to observe a canine patrol on a terminal, or he may mistake an officer walking across as engaged in a patrol. Since each target is patrolled and observed independently, we assume that both execution and observation noise are independent per target.

Again, consider the example in Table 1, suppose the defender has one resource. The SSE strategy for the defender would be protecting t_1 and t_2 with 0.5 probability each, making them indifferent for the attacker. The attacker breaks ties in defender's favor and chooses t_1 to attack, giving the defender an expected utility of 5. This SSE strategy is not robust to any noise – by deducting an infinitesimal amount of coverage probability from t_2, the attacker's best response changes to t_2, reducing the defender's expected utility to -5. We compute a risk-averse strategy, which provides the defender the maximum worst-case expected utility. For example, assuming no execution error and 0.1 observational uncertainty ($\alpha = 0$ and $\beta = 0.1$), the optimal risk-averse defender strategy is to protect t_1 with $0.4 - \epsilon$ probability and t_2 with $0.6 + \epsilon$ probability so that even in the worst-case, the attacker would choose t_1, giving the defender an expected utility of 4. Finding the optimal risk-averse strategy for large games remains difficult, as it is essentially a *bi-level* programming problem (Bard 2006).

Variable	Definition
T	Set of targets
$U_d^u(t_i)$	Defender's payoff if target t_i is uncovered
$U_d^c(t_i)$	Defender's payoff if target t_i is covered
$U_a^u(t_i)$	Attacker's payoff if target t_i is uncovered
$U_a^c(t_i)$	Attacker's payoff if target t_i is covered
γ	Number of defender resources
x_i	Defender's intended coverage of target t_i
y_i	Defender's actual coverage of target t_i
z_i	Attacker's observed coverage of target t_i
$\Delta U_d(t_i)$	$\Delta U_d(t_i) = U_d^c(t_i) - U_d^u(t_i)$
$\Delta U_a(t_i)$	$\Delta U_a(t_i) = U_a^u(t_i) - U_a^c(t_i)$
$D_i(x_i)$	Defender's expected utility for target t_i $D_i(x_i) = U_d^u(t_i) + \Delta U_d(t_i)x_i$
$A_i(x_i)$	Attacker's expected utility for target t_i $A_i(x_i) = U_a^u(t_i) - \Delta U_a(t_i)x_i$
α_i	Maximum execution error for target t_i
β_i	Maximum observation error for target t_i

Table 2: Notation

Problem Statement

The objective is to find the optimal risk-averse strategy $\mathbf{x}$, maximizing the worst-case defender utility, $U_d^*(\mathbf{x})$ (Constraint (1) and (2)). Given a fixed maximum execution and observation noise, $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ respectively, $U_d^*(\mathbf{x})$ is computed by the minimization problem from Constraint (3) to (6).

$$\max_{\mathbf{x}} \quad U_d^*(\mathbf{x}) \tag{1}$$

$$\text{s.t.} \quad \sum_{t_i \in T} x_i \leq \gamma,\ 0 \leq x_i \leq 1 \tag{2}$$

$$U_d^*(\mathbf{x}) = \min_{\mathbf{y},\mathbf{z},t_j} \quad D_j(y_j) \tag{3}$$

$$\text{s.t.} \quad t_j \in \arg\max_{t_i \in T} A_i(z_i) \tag{4}$$

$$-\alpha_i \leq y_i - x_i \leq \alpha_i,\ 0 \leq y_i \leq 1 \tag{5}$$

$$-\beta_i \leq z_i - y_i \leq \beta_i,\ 0 \leq z_i \leq 1 \tag{6}$$

The overall problem is a bi-level programming problem. For a fixed defender strategy $\mathbf{x}$, the *second-level* problem from Constraint (3) to (6) computes the worst-case defender's executed coverage $\mathbf{y}$, the attacker's observed coverage $\mathbf{z}$, and the target attacked t_j. $\langle \mathbf{y}, \mathbf{z}, t_j \rangle$ is chosen such that the defender's expected utility $D_j(y_j)$ (see Table 2) is minimized, given that the attacker maximizes his believed utility[1] $A_j(z_j)$ (Constraint (4)). This robust optimization is similar in spirit to Aghassi and Bertsimas (2006), although that is in the context of simultaneous move games.

This also highlights the need to separately model both execution and observation noise. Indeed a problem with un-

[1]The attacker's believed utility is computed using the strategy observed by the attacker, and it may not be achieved, since $\mathbf{z}$ can be different from $\mathbf{y}$, which can be different from $\mathbf{x}$.

certainty defined as $\langle \boldsymbol{\alpha}, \boldsymbol{\beta} \rangle$ is different from a problem with $\langle \boldsymbol{\alpha}' = 0, \boldsymbol{\beta}' = \boldsymbol{\alpha} + \boldsymbol{\beta} \rangle$ (or vice-versa), since the defender utility is different in the two problems. Other key properties of our approach include the solution of the above problem is an SSE if $\boldsymbol{\alpha} = \boldsymbol{\beta} = \mathbf{0}$. Furthermore, a MAXIMIN strategy is obtained when $\boldsymbol{\beta} = \mathbf{1}$ with $\boldsymbol{\alpha} = \mathbf{0}$, since $\mathbf{z}$ can be arbitrary. Finally, $\boldsymbol{\alpha} = \mathbf{1}$ implies that the execution of the defender is independent of $\mathbf{x}$ and thus, any feasible $\mathbf{x}$ is optimal.

Approach

We present RECON, *Risk-averse Execution Considering Observational Noise*, a mixed-integer linear programming (MILP) formulation to compute the risk-averse defender strategy in the presence of execution and observation noise. It encodes the necessary and sufficient conditions of the second-level problem (Constraint (4)) as linear constraints. The intuition behind these constraints is to identify $S(\mathbf{x})$, the *best-response* action set for the attacker given a strategy $\mathbf{x}$, and then break ties against the defender. Additionally, RECON represents the variables $\mathbf{y}$ and $\mathbf{z}$ in terms of the variable $\mathbf{x}$ – it reduces the bi-level optimization problem to a single-level optimization problem. We first define the term *inducible target* and then the associated necessary/sufficient conditions of the second level problem.

Definition 1. *We say a target t_j is weakly inducible by a mixed strategy $\mathbf{x}$ if there exists a strategy $\mathbf{z}$ with $0 \leq z_i \leq 1$ and $|z_i - x_i| \leq \alpha_i + \beta_i$ for all $t_i \in T$, such that t_j is the best response to $\mathbf{z}$ for the attacker, i.e., $t_j = \arg\max_{t_i \in T} A_i(z_i)$.*

Additionally, we define the upper and lower bounds on the utility the attacker may believe to obtain for the strategy profile $\langle \mathbf{x}, t_i \rangle$. These bounds will then be used to determine the *best response* set of the attacker.

Definition 2. *For the strategy profile $\langle \mathbf{x}, t_i \rangle$, the upper bound of attacker's believed utility is given by $A_i^+(x_i)$, which would be reached when the attacker's observed coverage of t_i reaches the lower bound $\max\{0, x_i - \alpha_i - \beta_i\}$.*

$$A_i^+(x_i) = \min\{U_a^u(t_i), A_i(x_i - \alpha_i - \beta_i)\} \tag{7}$$

Similarly, we denote the lower bound of attacker's believed utility of attacking target t_i by $A_i^-(x_i)$, which is reached when the attacker's observed coverage probability on t_i reaches the upper bound $\min\{1, x_i + \alpha_i + \beta_i\}$.

$$A_i^-(x_i) = \max\{U_a^c(t_i), A_i(x_i + \alpha_i + \beta_i)\} \tag{8}$$

Lemma 1. *A target t_j is weakly inducible by $\mathbf{x}$ if and only if $A_j^+(x_j) \geq \max_{t_i \in T} A_i^-(x_i)$.*

Proof. If t_j is weakly inducible, consider $\mathbf{z}$ such that $t_j = \arg\max_{t_i \in T} A_i(z_i)$. Since $z_j \geq \max\{0, x_j - \alpha_j - \beta_j\}$ and for all $t_i \neq t_j$, $z_i \leq \min\{1, x_i + \alpha_i + \beta_i\}$, we have:

$$A_j^+(x_j) = \min\{U_a^u(t_j), A_j(x_j - \alpha_j - \beta_j)\} \geq A_j(z_j)$$
$$\geq A_i(z_i) \geq \max\{U_a^c(t_i), A_i(x_i + \alpha_i + \beta_i)\} = A_i^-(x_i).$$

On the other hand, if $A_j^+(x_j) \geq A_i^-(x_i)$ for all $t_i \in T$, we can let $z_j = \max\{0, x_i - \alpha_j - \beta_j\}$ and $z_i = \min\{1, x_i + \alpha_i + \beta_i\}$ for all $t_i \neq t_j$, which satisfies $t_j = \arg\max_{t_i \in T} A_i(z_i)$. This implies t_j is weakly inducible. □

We also define $D_i^-(x_i)$, the lower bound on the defender's expected utility for the strategy profile $\langle \mathbf{x}, t_i \rangle$. This lower bound is used to determine the defender's worst-case expected utility.

Definition 3. *For the strategy profile $\langle \mathbf{x}, t_i \rangle$, $D_i^-(x_i)$ is achieved when the defender's implemented coverage on t_i reaches the lower bound $\max\{0, x_i - \alpha_i\}$, and is given by:*

$$D_i^-(x_i) = \max\{U_d^u(t_i), D_i(x_i - \alpha_i)\} \tag{9}$$

Lemma 2. *Let $S(\mathbf{x})$ be the set of all targets that are weakly inducible by $\mathbf{x}$, then $U_d^*(\mathbf{x}) = \min_{t_i \in S(\mathbf{x})} D_i^-(x_i)$.*

Proof. A target not in $S(\mathbf{x})$ cannot be attacked, since it is not the best response of the attacker for any feasible $\mathbf{z}$. Additionally, for any target t_i in $S(\mathbf{x})$, the minimum utility of the defender is $D_i^-(x_i)$. Therefore, $U_d^*(\mathbf{x}) \geq \min_{t_i \in S(\mathbf{x})} D_i^-(x_i)$.

Additionally, we prove $U_d^*(\mathbf{x}) \leq \min_{t_i \in S(\mathbf{x})} D_i^-(x_i)$ by showing there exist $\langle \mathbf{y}, \mathbf{z}, t_j \rangle$ satisfying Constraint (4) to (6) with $D_j(y_j) = \min_{t_i \in S(\mathbf{x})} D_i^-(x_i)$. To this end, we choose $t_j = \arg\min_{t_i \in S(\mathbf{x})} D_i^-(x_i)$, $y_j = \max\{0, x_j - \alpha_j\}$, $z_j = \max\{0, x_j - \alpha_j - \beta_j\}$, and $y_i = \min\{1, x_i + \alpha_i\}$, $z_i = \min\{1, x_i + \alpha_i + \beta_i\}$ for all $t_i \neq t_j$. $\mathbf{y}$ and $\mathbf{z}$ satisfy Constraint (5) and (6) by construction. And since t_j is weakly inducible, we have for all $t_i \neq t_j$, $A_j(z_j) = A_j^+(x_j) \geq A_i^-(x_i) = A_i(z_i)$, implying $t_j = \arg\max_{t_i \in T} A_i(z_i)$. □

Lemma (1) and (2) are the necessary and sufficient conditions for the second level optimization problem, reducing the bi-level optimization problem into a single level MILP.

RECON MILP

Now we present the MILP formulation for RECON. It maximizes the defender utility, denoted as v_d. v_a represents the *highest lower-bound* on the believed utility of the attacker ($A_i^+(x_i)$), given in Constraint (11). The binary variable q_i is 1 if the target t_i is weakly inducible; it is 0 otherwise. Constraint (12) says that $q_i = 1$ if $A_i^+(x_i) \geq v_a$ (ϵ is a small positive constant which ensures that $q_i = 1$ when $A_i^+(x_i) = v_a$) and together with Constraint (11), encodes Lemma 1. The constraint that $q_i = 0$ if $A_i^+(x_i) < v_a$ could be added to RECON, however, it is redundant since the defender wants to set $q_i = 0$ in order to maximize v_d. Constraint (13) says that the defender utility v_d is less than $D_i^-(x_i)$ for all inducible targets, thereby implementing Lemma 2. Constraint (14) ensures that the allocated resources are no more than the number of available resources γ, maintaining feasibility.

$$\max_{\mathbf{x}, \mathbf{q}, v_d, v_a} \quad v_d \tag{10}$$

$$\text{s.t.} \quad v_a = \max_{t_i \in T} A_i^-(x_i) \tag{11}$$

$$A_i^+(x_i) \leq v_a + q_i M - \epsilon \tag{12}$$

$$v_d \leq D_i^-(x_i) + (1 - q_i) M \tag{13}$$

$$\sum_i x_i \leq \gamma \tag{14}$$

$$x_i \in [0, 1] \tag{15}$$

$$q_i \in \{0, 1\} \tag{16}$$

The max function in Constraint (11) can be formulated using $|T|$ binary variables, $\{h_i\}$, in the following manner:

$$A_i^-(x_i) \le v_a \le A_i^-(x_i) + (1-h_i)M \tag{17}$$

$$\sum_{t_i \in T} h_i = 1, \quad h_i \in \{0,1\} \tag{18}$$

The min operation in $A_i^+(x_i)$ is also implemented similarly. For example, Equation (7) can be encoded as:

$$U_a^u(t_i) - (1-\nu_i)M \le A_i^+ \le U_a^u(t_i)$$
$$A_i(x_i - \alpha_i - \beta_i) - \nu_i M \le A_i^+ \le A_i(x_i - \alpha_i - \beta_i)$$
$$\nu_i \in \{0,1\}$$

We omit the details for expanding $A_i^-(x_i)$ and $D_i^-(x_i)$, they can be encoded in a similar fashion.

Speeding up

As described above, RECON uses a MILP formulation to compute the risk-averse strategy for the defender. Integer variables increase the complexity of the linear programming problem; indeed solving integer programs is NP-hard. MILP solvers internally use branch-and-bound to evaluate integer assignments. Availability of good lower bounds implies that less combinations of integer assignments (branch-and-bound nodes) need to be evaluated. This is the intuition behind speeding up the execution of RECON. We provide two methods, a-RECON and i-RECON, to generate lower bounds.

a-RECON: a-RECON solves a *restricted* version of RECON. This restricted version has lower number of integer variables, and thus generates solutions faster. It replaces $A_i^+(x_i)$ by $A_i(x_i - \alpha_i - \beta_i)$ and $D_i^-(x_i)$ by $D_i(x_i - \alpha_i)$, thereby rewriting Constraints (12) and (13) as follows:

$$A_i(x_i - \alpha_i - \beta_i) \le v_a + q_i M - \epsilon \tag{19}$$

$$v_d \le D_i(x_i - \alpha_i) + (1-q_i)M \tag{20}$$

a-RECON is indeed more *restricted* – the LHS of Constraint (19) in a-RECON is *no less than* the LHS of Constraint (12) in RECON; and the RHS of Constraint (20) is *no greater than* the RHS of Constraint (13) in a-RECON. Therefore, any solution generated by a-RECON is feasible in RECON, and acts as a lower bound. We do not restrict $A_i^-(x_i)$ since the RHS of Constraint (17) is non-trivial for only one target.

i-RECON: i-RECON uses an iterative method to obtain monotonically increasing lower bounds $v_d^{(k)}$ of RECON. Using the insight that Constraint (19) is *binding* only when $q_i = 0$, and (20) when $q_i = 1$, i-RECON rewrites Constraints (19) and (20) as follows:

$$x_i \ge \begin{cases} \tau_{a,i}(v_a) = \frac{U_a^u(t_i) - v_a + \epsilon}{\Delta U_a(t_i)} + \alpha_i + \beta_i & \text{if } q_i = 0 \\ \tau_{d,i}(v_d) = \frac{v_d - U_d^u(t_i)}{\Delta U_d(t_i)} + \alpha_i & \text{if } q_i = 1 \end{cases} \tag{21}$$

which says that $q_i = 0$ implies $x_i \ge \tau_{a,i}(v_a)$ and $q_i = 1$ implies $x_i \ge \tau_{d,i}(v_d)$.[2] Constraint (21) is equivalent to:

$$\begin{aligned} x_i \ge\ & \min\{\tau_{d,i}(v_d), \tau_{a,i}(v_a)\} \\ =\ & \tau_{d,i}(v_d) + \min\{0, \tau_{a,i}(v_a) - \tau_{d,i}(v_d)\} \end{aligned} \tag{22}$$

[2]This is *not* equivalent to the unconditional equation $x_i \ge \max\{\tau_{a,i}(v_a), \tau_{d,i}(v_d)\}$.

Algorithm 1: Pseudo code of i-RECON

```
k = 0, v_d^(0) = v_a^(0) = −∞;
while |v_a^(k+1) − v_a^(k)| ≤ η and |v_d^(k+1) − v_d^(k)| ≤ η do
    v_a^(k+1) = Solve(A-LP (v_d^(k), v_a^(k)));
    v_d^(k+1) = Solve(D-LP (v_d^(k), v_a^(k)));
    k = k + 1;
end
```

The equivalence between Constraint (21) and (22) can be verified as follows: $\langle \mathbf{x}, v_d, v_a \rangle$ from any feasible solution $\langle \mathbf{x}, \mathbf{q}, v_d, v_a \rangle$ of (21) is trivially feasible in (22). On the other hand, given a feasible solution $\langle \mathbf{x}, v_d, v_a \rangle$ to Constraint (22), we choose $q_i = 1$ if $x_i \ge \tau_{d,i}(v_d)$ and 0 otherwise, and thus obtain a feasible solution to Constraint (21). Hence, we obtain an equivalent problem of a-RECON by replacing Constraints (12) and (13) by (22). In the k^{th} iteration, i-RECON substitutes $\tau_{d,i}(v_d) - \tau_{a,i}(v_a)$ by a constant, $\Delta\tau_i^{(k)}$, *restricting* Constraint (22). This value is updated in every iteration while maintaining a restriction of Constraint (22). Such a substitution reduces Constraint (22) to a linear constraint, implying that i-RECON performs a polynomial-time computation in every iteration.[3]

Observe that $\tau_{d,i}(v_d)$ is increasing in v_d where as $\tau_{a,i}(v_a)$ is decreasing in v_a (refer Constraint (21)), and hence $\tau_{d,i}(v_d) - \tau_{a,i}(v_a)$ is *increasing* in both v_d and v_a. i-RECON generates an increasing sequence of $\{\Delta\tau_i^{(k)} = \tau_{d,i}(v_d^{(k)}) - \tau_{a,i}(v_a^{(k)})\}$ by finding increasing sequences of $v_d^{(k)}$ and $v_a^{(k)}$. As we will show later, substituting $\tau_{d,i}(v_d) - \tau_{a,i}(v_a)$ with $\{\Delta\tau_i^{(k)}\}$ in Constraint (22) guarantees correctness. Since a higher value of $\Delta\tau_i^{(k)}$ implies a lower value of $\min\{0, -\Delta\tau_i^{(k)}\}$, a weaker restriction is imposed by Constraint (22), leading to a better lower bound $v_d^{(k+1)}$.

Given $v_d^{(k)}$ and $v_a^{(k)}$, i-RECON uses D-LP to compute the $v_d^{(k+1)}$, and A-LP to compute $v_a^{(k+1)}$. The pseudo-code for i-RECON is given in Algorithm 1. D-LP is the following maximization linear program, which returns the solution vector $\langle \mathbf{x}, v_d, \hat{v}_a \rangle$, such that v_d is the desired lower bound.

$$\max_{\mathbf{x}, v_d, \hat{v}_a} \ v_d$$
$$\text{s.t. Constraint}(11), (14) \text{ and } (15)$$
$$x_i \ge \tau_{d,i}(v_d) + \min\{0, -\Delta\tau_i^{(k)}\} \tag{23}$$
$$v_d \ge v_d^{(k)}; \quad \hat{v}_a \ge v_a^{(k)} \tag{24}$$

Constraint (24) is added to D-LP to ensure that we get a monotonically increasing solution in every iteration. Similarly, given $v_d^{(k)}$ and $v_a^{(k)}$, A-LP is the following minimization problem. It minimizes v_a to guarantee that Constraint (23) in D-LP remains a restriction to Constraint (22) for the next iteration, ensuring D-LP always provides a lower bound of RECON. More details are in Proposition 1 which proves the

[3]While the formulation has integer variables from Constraint (11), it can be considered as $2|T|$ LPs since there are only $2|T|$ distinct combinations of integer assignments.

correctness of i-RECON.

$$\min_{\mathbf{x}, v_d, v_a} \quad v_a$$
$$\text{s.t. Constraint } (11), (14) \text{ and } (15)$$
$$x_i \geq \tau_{a,i}(v_a) + \min\{\Delta\tau_i^{(k)}, 0\} \tag{25}$$
$$v_a \geq v_a^{(k)} \tag{26}$$

Proposition 1. *Both* D-LP *and* A-LP *are feasible and bounded for every iteration k until i-*RECON *converges.*

Proof. A-LP is bounded for every iteration because $v_a \geq \max_{t_i \in T} U_a^c(t_i)$ by Constraint (11). We prove the rest of the proposition using induction. First we establish that both D-LP and A-LP are feasible and bounded in the first iteration. In the first iteration, D-LP is feasible for any value of $x_i \geq 0$ when $v_d = \min_{t_i \in T}\{U_d^u(t_i) - \alpha_i \Delta U_d(t_i)\}$ (from Constraint (21)), and it is bounded since $\tau_{d,i}(v_d) \leq x_i \leq 1$ for all $t_i \in T$. In the same way, for A-LP, Constraint (25) becomes $x_i \geq -\infty$ in the first iteration. Thus, $v_a = \max_{t_i \in T} A_i^-(x_i) > -\infty$ is a feasible solution.

Assuming that D-LP and A-LP are feasible and bounded for iterations $1, 2, \ldots, k$, we now show that they remain bounded and feasible in iteration $k+1$. Firstly, D-LP is bounded in the $k+1^{\text{th}}$ iteration since $\tau_{d,i}(v_d) \leq 1 - \min\{0, -\Delta\tau_i^{(k)}\}$ for all $t_i \in T$. D-LP is feasible because the solution from the k^{th} iteration, $\langle \mathbf{x}^{(k)}, v_d^{(k)}, \hat{v}_a^{(k)} \rangle$, remains feasible. To see this, observe that since $\tau_{d,i}^{(k)}$ is increasing and $\tau_{a,i}^{(k)}$ is decreasing with k, thus we have $\Delta\tau_i^{(k)} \geq \Delta\tau_i^{(k-1)}$. Hence $\min\{0, -\Delta\tau_i^{(k-1)}\} \geq \min\{0, -\Delta\tau_i^{(k)}\}$, implying that $\langle \mathbf{x}^{(k)}, v_d^{(k)} \rangle$ satisfies Constraint (23). Moreover, Constraints (11), (14), (15) and (24) are trivially satisfied.

Similarly, A-LP is also feasible in the $k+1^{th}$ iteration since $\langle \mathbf{x}^{(k+1)}, v_d^{(k+1)}, \hat{v}_a^{(k+1)} \rangle$, the solution returned by D-LP in the $k+1^{th}$ iteration, satisfies all the constraints of A-LP. Firstly, Constraints (11), (14), (15) and (26) are trivially satisfied. Secondly, Constraint (25) is also satisfied since:

$$\tau_{d,i}(v_d^{(k+1)}) - \tau_{a,i}(\hat{v}_a^{(k+1)}) \geq \Delta\tau_i^{(k)}. \tag{27}$$

$$\begin{aligned}
x_i^{(k+1)} &\geq \tau_{d,i}(v_d^{(k+1)}) + \min\{0, -\Delta\tau_i^{(k)}\} && \text{from (23)} \\
&= \min\{\tau_{d,i}(v_d^{(k+1)}), \tau_{d,i}(v_d^{(k+1)}) - \Delta\tau_i^{(k)}\} \\
&\geq \min\{\tau_{d,i}(v_d^{(k+1)}), \tau_{a,i}(\hat{v}_a^{(k+1)})\} && \text{from (27)} \\
&= \tau_{a,i}(\hat{v}_a^{(k+1)}) + \min\{\tau_{d,i}(v_d^{(k+1)}) - \tau_{a,i}(\hat{v}_a^{(k+1)}), 0\} \\
&\geq \tau_{a,i}(\hat{v}_a^{(k+1)}) + \min\{\Delta\tau_i^{(k)}, 0\} && \text{from (27)}
\end{aligned}$$

Similarly, we can show that $\langle \mathbf{x}^{(k+1)}, v_d^{(k+1)}, \hat{v}_a^{(k+1)} \rangle$ is a feasible solution of a-RECON for any k using inequality (27), and hence, $v_d^{(k+1)}$ is a lower bound of RECON. Additionally, since the sequence $\{v_d^{(k)}\}$ is bounded and monotonically increasing, we can conclude that it converges. □

Experimental Results

We provide two sets of experimental results: (1) we compare the solution quality of RECON, ERASER, and COBRA under uncertainty: ERASER (Jain et al. 2010) is used to compute the SSE solution, where as COBRA (Pita et al. 2010) is one of the latest algorithms that addresses attacker's observational error.[4] (2) We provide the runtime results of RECON, showing the effectiveness of the two heuristics a-RECON and i-RECON. In all test instances, we set the number of defender resources to 20% of the number of targets. Payoffs $U_d^c(t_i)$ and $U_a^u(t_i)$ are chosen uniformly randomly from 1 to 10 while $U_d^u(t_i)$ and $U_a^c(t_i)$ are chosen uniformly randomly from -10 to -1. The results were obtained using CPLEX on a standard 2.8GHz machine with 2GB main memory, and averaged over 39 trials. All comparison results are statistically significant under t-test ($p < 0.05$).

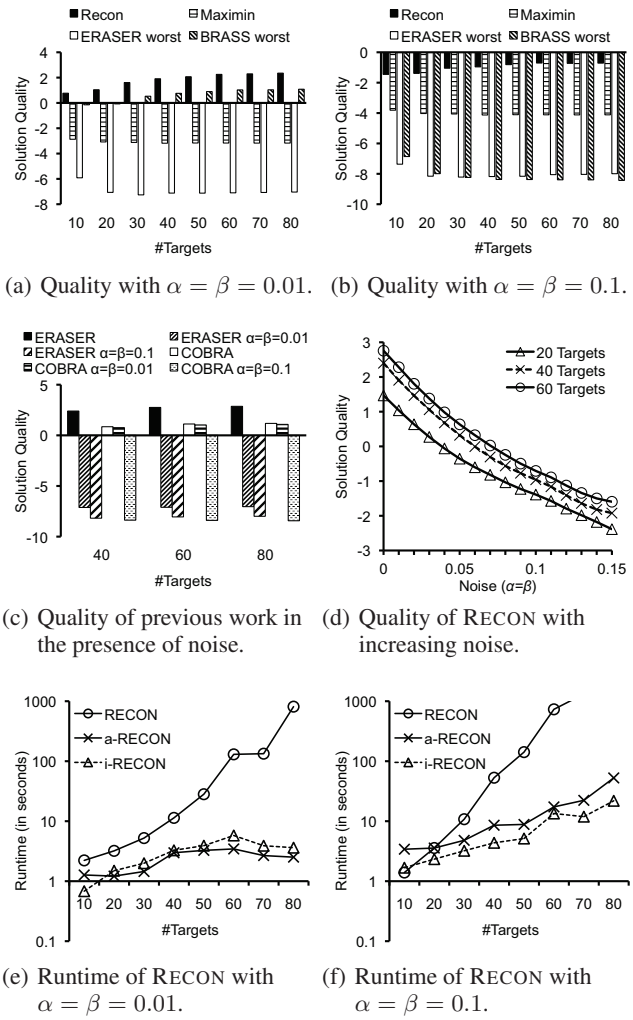

(a) Quality with $\alpha = \beta = 0.01$. (b) Quality with $\alpha = \beta = 0.1$.

(c) Quality of previous work in the presence of noise. (d) Quality of RECON with increasing noise.

(e) Runtime of RECON with $\alpha = \beta = 0.01$. (f) Runtime of RECON with $\alpha = \beta = 0.1$.

Figure 1: Experimental results.

Measuring robustness of a given strategy: Given a defender mixed strategy $\mathbf{x}$, a maximum execution error $\boldsymbol{\alpha}$, and a maximum possible observation error $\boldsymbol{\beta}$, the worst-case defender utility is computed using the second-level optimization problem given in Constraints (3) to (6). Figure 1(a) and

[4]The bounded rationality parameter ϵ in COBRA is set to 2 as suggested by Pita et. al (2010). The bias parameter α is set to 1 since our experiments are not tested against human subjects.

Figure 1(b) presents the comparisons between the worst-case utilities of RECON, ERASER and COBRA under two uncertainty settings – low uncertainty ($\boldsymbol{\alpha} = \boldsymbol{\beta} = \mathbf{0.01}$) and high uncertainty ($\boldsymbol{\alpha} = \boldsymbol{\beta} = \mathbf{0.1}$). Also, MAXIMIN utility is provided as a benchmark. Here x-axis shows the number of targets and y-axis shows the defender's worst-case utility. RECON significantly outperforms MAXIMIN, ERASER and COBRA in both uncertainty settings. For example, in high uncertainty setting, for 80 targets, RECON on average provides a worst-case utility of -0.7, significantly better than MAXIMIN (-4.1), ERASER (-8.0) and COBRA (-8.4).

Next, in Figure 1(c), we present the ideal defender utilities of ERASER and COBRA assuming no execution and observational uncertainties, comparing to their worst-case utilities (computed as described above). Again, x-axis is the number of targets and y-axis is the defender's worst-case utility. As we can see, ERASER is not robust – even 0.01 noise reduces the solution quality significantly. For instance, for 80 targets with low uncertainty, ERASER on average has a worst-case utility of -7.0 as opposed to an ideal utility of 2.9. Similarly, COBRA is not robust to large amount of noise (0.1) although it is robust when noise is low (0.01). Again, for 80 targets, COBRA on average has an ideal utility of 1.2, however, its worst-case utility drops to -7.0 in high uncertainty setting.

Finally, in Figure 1(d), we show the quality of RECON with increasing noise from $\boldsymbol{\alpha} = \boldsymbol{\beta} = \mathbf{0}$ to $\boldsymbol{\alpha} = \boldsymbol{\beta} = \mathbf{0.15}$. The x-axis shows the amount of noise while the y-axis shows the defender's utility returned by RECON. The three lines represent 20, 40, and 60 targets respectively. As we can see, while ERASER and COBRA cannot adapt to noise even when bounds on noise are known *a-priori*, RECON is more robust and provides significantly higher defender utility.

Runtime results of RECON: Figures 1(e) and 1(f) show the runtime results of the three variants of RECON— RECON without any lower bounds, and with lower bounds provided by a-RECON and i-RECON respectively. The x-axis shows the number of targets and the y-axis (in logarithmic scale) shows the total runtime in seconds. Both a-RECON and i-RECON heuristics help reduce the total runtime significantly in both uncertainty settings – the speedup is of orders of magnitude in games with large number of targets. For instance, for cases with 80 targets and high uncertainty, RECON without heuristic lower bounds takes $3,948$ seconds, whereas RECON with a-RECON lower bound takes a total runtime of 52 seconds and RECON with i-RECON lower bound takes a total runtime of 22 seconds.

Conclusions

Game-theoretic scheduling assistants are now being used daily to schedule checkpoints, patrols and other security activities by agencies such as LAX police, FAMS and the TSA. Augmenting the game-theoretic framework to handle the fundamental challenge of uncertainty is pivotal to increase the value of such scheduling assistants. In this paper, we have presented RECON, a new model that computes risk-averse strategies for the defender in the presence of execution and observation uncertainty. Our experimental results show that RECON is robust to such noise where the performance of existing algorithms can be arbitrarily bad. Additionally, we have provided two heuristics, a-RECON and i-RECON, that further speed up the performance of RECON. This research complements other research focused on handling other types of uncertainty such as in payoff and decision making (Kiekintveld, Tambe, and Marecki 2010), and could ultimately be part of a single unified robust algorithm.

Acknowledgement

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. F. Ordóñez was also supported by FONDECYT through Grant 1090630. We thank Matthew Brown for detailed comments and suggestions.

References

Aghassi, M., and Bertsimas, D. 2006. Robust game theory. *Math. Program.* 107:231–273.

Agmon, N.; Sadov, V.; Kaminka, G. A.; and Kraus, S. 2008. The Impact of Adversarial Knowledge on Adversarial Planning in Perimeter Patrol. In *AAMAS*, volume 1.

Archibald, C., and Shoham, Y. 2009. Modeling billiards games. In *AAMAS*.

Bard, J. F. 2006. *Practical Bilevel Optimization: Algorithms and Applications (Nonconvex Optimization and Its Applications)*. Springer-Verlag New York, Inc.

Basilico, N.; Gatti, N.; and Amigoni, F. 2009. Leader-follower strategies for robotic patrolling in environments with arbitrary topologies. In *AAMAS*.

Conitzer, V., and Sandholm, T. 2006. Computing the optimal strategy to commit to. In *ACM EC-06*, 82–90.

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 Marshals Service. *Interfaces* 40:267–290.

Kiekintveld, C.; Tambe, M.; and Marecki, J. 2010. Robust bayesian methods for stackelberg security games (extended abstract). In *AAMAS Short Paper*.

Kodialam, M., and Lakshman, T. 2003. Detecting network intrusions via sampling: A game theoretic approach. In *INFOCOM*.

Pita, J.; Jain, M.; Tambe, M.; Ordóñez, F.; and Kraus, S. 2010. Robust solutions to stackelberg games: Addressing bounded rationality and limited observations in human cognition. *AIJ* 174:1142–1171.

van Damme, E., and Hurkens, S. 1997. Games with imperfectly observable commitment. *Games and Economic Behavior* 21(1-2):282 – 308.

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

Yin, Z.; Korzhyk, D.; Kiekintveld, C.; Conitzer, V.; and Tambe, M. 2010. Stackelberg vs. Nash in security games: interchangeability, equivalence, and uniqueness. In *AAMAS*.