Multi-Defender Security Games on Networks

Andrew Smith, Yevgeniy Vorobeychik, and Joshua Letchford
Sandia National Laboratories*
Livermore, CA

ABSTRACT

Stackelberg security game models and associated computational tools have seen deployment in a number of high-consequence security settings, such as LAX canine patrols and Federal Air Marshal Service. This deployment across essentially independent agencies raises a natural question: what global impact does the resulting strategic interaction among the defenders, each using a similar model, have? We address this question in two ways. First, we demonstrate that the most common solution concept of Strong Stackelberg equilibrium (SSE) can result in significant under-investment in security entirely because SSE presupposes a single defender. Second, we propose a framework based on a different solution concept which incorporates a model of interdependencies among targets, and show that in this framework defenders tend to over-defend, even under significant positive externalities of increased defense.

1. INTRODUCTION

Security, physical and cyber, has come to the forefront of national attention, particularly after 9/11. Among the variety of approaches that are used to tackle security problems, from risk analysis to red teaming, game theory has had a non-trivial impact, with tools based on game theoretic analysis having been deployed in LAX aiport to schedule canine patrols [13, 6, 14], by Federal Air Marshall Service (FAMS) to schedule the air marshals [9, 7, 5], and by the US Coast Guard to schedule boat patrols [15]. All of these deployments, and numerous other related efforts, have cast security as a Stackelberg game between a single defender and an attacker, in which the defender leads (i.e., acts first), choosing a probability distribution over defense actions, and the attacker, upon learning this probability distribution, chooses a response [4]. In many cases, the attacker is modeled as a rational agent who selects an optimal response and, in the many applications that compute a Strong Stackelberg equilibrium, an attacker is often assumed to break ties in the defender's favor [13, 10]. A crucial assumption that all these efforts have in common is that they assume a single defender. In practice, numerous parties are responsible for security; indeed, the fact that the basic framework has been deployed by different entities and agencies makes this manifest already. If security decisions made by different parties were entirely independent, both from the defender's and the attacker's perspective, a single-defender model would be entirely satisfactory. However, the assets protected by different entities are typically interdependent, or, more generally, have value to others who are not involved in security decisions. Additionally, attackers, insofar as they may target different sectors under the charge of different defenders, are resource constrained, implicitly coupling otherwise independent targets.

We extend the standard computational Stackelberg game framework to analyze games with multiple defenders. One key reason for using the Stackelberg game framework as a point of departure is that as single-defender approaches are increasingly deployed by the different parties, it is important to anticipate the joint defense that emerges in the long-run as a result. Additionally, unlike other multi-defender models (e.g., [11, 3, 1, 2]), our approach maintains the typical complexity of *individual defender* decision process in the multi-defender framework, with each defender responsible for securing *many*, possibly interdependent, targets.

Our setup gives rise to two competing externalities of security decisions: a *positive* externality, where greater security implies reduced contagion risk to other defenders, and a *negative* externality, which arises because high security by one player pushes the attacker to attack someone else's assets. Our main analytic contribution is therefore to study the impact of these competing effects of defense on the resulting Nash equilibrium outcomes. Overall, our results suggest that the negative externality dominates, and defenders tend to over-invest in security. However, the impact on defense outcomes and welfare (relative to optimal) is substantial only when the security decisions are significantly decentralized, and in the intermediate cases, joint defense decisions are nearly optimal.

2. PRELIMINARIES

Our point of departure is the relatively mature literature on *single-defender* Stackelberg security games. A (single-defender) Stackelberg security game consists of two players, the leader (defender) and the follower (attacker), and a set of possible targets. The leader can decide upon a randomized policy of defending the targets, possibly with limited defense resources. The follower (attacker) is assumed to observe the randomized policy of the leader, but not the realized defense

*Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

actions. Upon observing the leader's strategy, the follower chooses a target so as to maximize its expected utility.

In multi-defender security games (our model), multiple defenders simultaneously and non-cooperatively choose security policies over a collection of non-overlapping targets which they own. The attacker observes this joint policy of all defenders and attacks a single target that maximizes its expected utility. The restriction we impose in these games is that the set of targets is partitioned among the defenders; each defender can only choose defense configurations over the targets assigned to it. However, defenders may positively *value* any of the targets, including those that are outside of their domain of influence.

To formalize, let D be the set of all defenders, and suppose that a defender d can choose from a finite set O of security configurations for each target $t \in T_d$, where T_d is the set of targets under d's direct influence. Let T be the set of *all* targets, that is, $T = \cup_d T_d$, with $|T| = n$. A configuration $o \in O$ for target $t \in T_d$ incurs a cost $c_{o,t}$ to the defender d.

If the attacker happens to attack a particular target $t \in T$ while configuration o is in place, the expected value to a defender d is denoted by $U^d_{o,t}$, while the attacker's value is $V_{o,t}$. We assume throughout that each player's utility depends only on the target attacked and its security configuration [9, 12]. We denote by $q^d_{o,t}$ the probability that the defender d chooses o at target $t \in T_d$.

While the problem we study assumes that that the utility of any player for a given target depends only on its security configuration o, there is a rather natural way to model interdependencies while retaining this structure, proposed by Letchford and Vorobeychik [12]. Specifically, suppose that dependencies between targets are represented by a graph (T, E), with T the set of targets (nodes) as above, and E the set of edges (t, t'), where an edge from t to t' means that a successful attack on t may have impact on t'. Each target has associated with it a worth, w^d_t, for the defender d, which is the loss to d if t is affected (e.g., compromised, broken). The security configuration determines the probability $z_{o,t}(t)$ that target t is affected if the attacker attacks it *directly* and the defense configuration is o. We model the interdependencies between the nodes as independent cascade contagion [8, 12]. The contagion proceeds starting at an attacked node t, affecting its network neighbors t' each with probability $p_{t,t'}$; the contagion then spreads from the newly affected nodes t' to their neighbors, and so on. The contagion can only occur one time along any network edge, and once a node is affected it stays affected through the diffusion process.

3. COMPUTING A DEFENDER'S BEST RESPONSE

A crucial step in computing (or approximating) a Nash equilibrium of a game is to consider the problem of computing a best response for an arbitrary player (in our case, defender, since the attacker's best response is straightforward). Appealing to the standard computational methods for the now common Strong Stackelberg equilibrium is tempting in this case, given the rich literature on the two-player Stackelberg security games that use that solution concept [10]. We now describe the problem with this concept in the multi-defender setting, and propose an alternative solution concept that is more appropriate to our setting.

3.1 The Weakness of Strong Stackelberg Equilibrium

By far the most important solution concept in Stackelberg security games is a *Strong Stackelberg equilibrium (SSE)* [10]. A SSE is characterized by an assumption that the attacker breaks ties in defender's favor. When there is a single defender, this is well defined, and quite reasonable when the defender can commit to a mixed strategy: a slight adjustment in the defense policy will force the attacker to strictly prefer the desired option, with little loss to the defender. As we now illustrate, however, SSE is fundamentally problematic in a multi-defender context, because the notion of "breaking ties in defender's favor" is no longer well defined in general, as we must specify *which* defender will receive the favor.

To see concretely what goes wrong, consider the example in Figure 1. In this example, there are two defenders, one

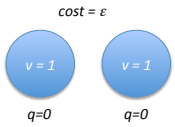

Figure 1: Example of a problem with a SSE in a multi-defender setting.

who defends the target on the left, while the other defends the target on the right. Both defenders value their respective targets at 1, and have no value for the counterpart's target. The cost of defending each target is $0 < \epsilon \ll 1$. Now, consider a strategy profile in which $q_t = 0$ for both targets t, and let us focus on the best response of the first (left) defender. If this defender attempts to compute an SSE by fixing the strategy of the second player, he perceives his utility under the current strategy profile to be 1, since he would assume that the attacker breaks ties in his favor and, thus, attacks the defender on the right. By the same logic, the defender on the right will assume that the attacker will attack his counterpart, and perceive $q_t = 0$ to be the best response. Since the attacker actually attacks one of them, the best response of the defender being attacked is to defend with a small probability, pushing the attacker towards the other target. What goes wrong here is that both players assume that the attacker attacks the other (breaks ties in their favor), which is inconsistent with the assumption that the attacker will certainly attack *some* target.

One possible resolution of the problem with SSE is to constrain that, in computing a player's best response, the attacker *strictly prefers* a single target over any other. This resolution would be convenient, as it would require only a minor modification to the standard formulations for computing a defender's best response based on the SSE solution concept. However, we now illustrate that this gives rise to a different problem: the utility of a defender is undefined when the attacker has multiple best responses. To see what can go wrong in this case, consider the above example, but suppose that each defender also has a value of 2 for the counterpart's target, and consider a best response of the left defender when $q_t = 0$ for both targets t. Since the constraint of the best response computation of a player is that the attacker has a strict preference for some single target, $q_t = 0$ is not feasible, as it causes the attacker to be indif-

ferent between the two targets. The defender's best *feasible* choice will make the attacker strictly prefer to attack the other player's target, a clear loss compared to staying with $q_t = 0$.

3.2 ASE: An Alternative Solution Concept

Since the classic (two-player) SSE solution concept used in Stackelberg security games does not conceptually extend to be an individual defender best-response problem in the multi-defender setting, we need to consider an alternative. One option is to compute an arbitrary subgame perfect equilibrium. However, we wish to impose a natural constraint on the solution concept that the attacker's best response be computed consistently for any joint defense policy, just as it is in a SSE (in other words, we wish to fix a tie-breaking rule). One natural tie-breaking rule is that the attacker chooses a target uniformly at random from the set of all best responses [1]. We call the corresponding solution concept (which is a refinement of the subgame perfect equilibrium of our game) the *Average-case Stackelberg Equilibrium (ASE)*. The crucial property of this solution concept that we desire is that the attacker's behavior presumed by a defender's best response problem is independent of that defender's identity, a property that SSE violates.

3.3 Computing ASE

While ASE seems a very natural alternative to SSE even in two-player security games, we are not aware of any proposals for computing it. Below, we present the first (to our knowledge) MILP formulation for computing ASE, which in our case would compute a best response for an arbitrary defender d when the strategies of all others, $q^{-d}_{t,o}$, are fixed.

$$\max_{a,q^d,s,u,v} \; u - \sum_{t\in T_d}\sum_{o\in O} c_{t,o} q^d_{t,o} \tag{1}$$

s.t.

$$0 \le q^d_{t,o} \le 1 \qquad \forall\, t \in T_d,\ \forall o \tag{2}$$

$$\sum_{o\in O} q^d_{t,o} = 1 \qquad \forall t \in T_d \tag{3}$$

$$a_t \in \{0,1\} \qquad \forall\, t \in T \tag{4}$$

$$\sum_{t\in T} a_t \ge 1 \tag{5}$$

$$0 \le v - \sum_o q^d_{t,o} V_{t,o} \le (1-a_t)M \qquad \forall\, t \in T_d \tag{6}$$

$$0 \le v - \sum_o q^{-d}_{t,o} V_{t,o} \le (1-a_t)M \qquad \forall\, t \in T_{-d} \tag{7}$$

$$s_t = v - \sum_o q^d_{t,o} V_{t,o} \qquad \forall\, t \in T_d \tag{8}$$

$$s_t = v - \sum_o q^{-d}_{t,o} V_{t,o} \qquad \forall\, t \in T_{-d} \tag{9}$$

$$a_t + M s_t \ge 1 \qquad \forall\, t \in T \tag{10}$$

$$u = f(q,a), \tag{11}$$

where M is a very large number and

$$f(q,a) = \frac{\sum_{t\in T_i} a_t \sum_{o\in O} q^i_{t,o} U_{t,o} + \sum_{t\in T_{-i}} a_t \sum_{o\in O} q^{-i}_{t,o} U_{t,o}}{\sum_{t\in T} a_t}.$$

While constraint 11 is non-linear, we can linearize it using McCormick inequalities.

4. APPROXIMATING NASH EQUILIBRIA

Previously, Vorobeychik and Wellman [16] presented a convergent equilibrium approximation algorithm based on *simulated annealing (SA)* that would be applicable in our setting. They additionally showed in simulation that SA is actually outperformed by a simple heuristic based on *iterated best response (IBR)* dynamics. Here, we interpret IBR as a greedy equilibrium approximation heuristic, with the property that if the starting point is a Nash equilibrium, IBR will never deviate from it (i.e., Nash equilibrium is a fixed point). Clearly, then, the choice of a starting point can be significant for the performance of IBR, making it natural to consider coupling it with random restarts. Our main contribution in this section is to present evidence that IBR with random restarts is a highly effective equilibrium approximation approach in our setting (and outperforms several alternatives). This is both of broad significance, and of particular importance in our setting, as we use this algorithm for our analyses below.

We compare the following Nash equilibrium approximation algorithms executed for 1000 iterations: random search (RS), which simply generates 1000 strategy profiles randomly, computes the game theoretic regret of each, and chooses a profile with the smallest regret; simulated annealing (SA), with the temperature exponentially increasing with iterations; and iterated best response (IBR) with no restarts. We also include in the comparison two additional variations of IBR: the first uses SA for the first 100 iterations, and then switches to IBR for the remainder (starting with the best approximation produced by SA); the second is IBR with random restarts, which we term RIBR. We execute our comparison on games with 2 players and 10 targets and games with 5 players and 20 targets. In all cases, targets are divided evenly among the players, and values over the targets are generated uniformly at random. The cost of defense is fixed at $c = 0.2$, and the targets are assumed to be independent (but players may have values for targets under the control of other defenders). Figure 2 demonstrate that in both settings, RIBR outperforms other alternatives.

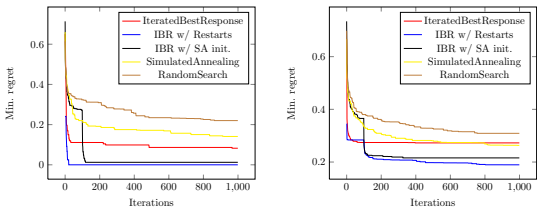

Figure 2: Comparison of algorithms. Left: $|D| = 5$ and $|T| = 20$. Right: $|D| = 2$ and $|T| = 10$.

5. ANALYSIS OF MULTI-DEFENDER SECURITY GAMES ON NETWORKS

Now that we have a method for approximately computing Nash equilibria in multi-defender security games, we proceed with analyzing instances of such games, in which the players' utilities for targets are derived based on the model of interdependencies described earlier, taking the worth of each target (asset) to be 1 to the player who owns it.

Our setting is a 8×8 grid, which we symmetrically divide among the $|D|$ players. Throughout, we assume that a

defender has only two defense options, to defend, and not, with the cost of defense being a uniform c (independent of players and targets), while no defense is free. In the experiments we vary three parameters: the number of players, $|D|$, the cost of defense, c, and the probability of cascade via an edge, p (identical for all edges).

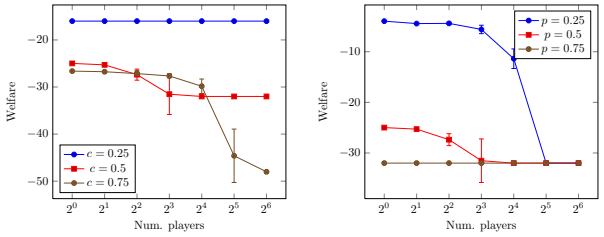

Figure 3: Average welfare of all players on a 8 x 8 grid.

Figure 3 presents welfare (total expected utility of the defenders) as a function of the problem parameters, with the horizontal axis corresponding to the number of defenders. The general trend is not in itself surprising: welfare decreases as the number of players increases (clearly, a single player computes the globally optimal solution). Another general and intuitive feature exhibited by Figure 3, right, is that increasing p decreases welfare across the board. Let us focus now on a few surprises. First, note that for intermediate numbers of players (between 2 and 4 in the medium-cost/p case, and 2-16 when the cost is high or p is low) the welfare curve is relatively little impacted by decentralization: equilibrium solutions are relatively close to optimal. When the resource control is significantly decentralized, however, externalities tend to lead to significant welfare loss. Second, with high decentralization (32-64 players), the value of p (i.e., the extent of positive security externalities) has no impact on welfare.

To understand these phenomena, we consider what happens to the average defense (i.e., average q_t over all targets t) in equilibrium, shown in Figure 4. Both as we vary defense costs, c, and the extent of network externalities, p, increasing the number of players leads to higher security investment, and when each player controls a single node in the grid, all players fully protect their nodes, resulting in security allocation well in excess of optimal. To sum up, the negative externalities appear to dominate, and players over-defend, although the impact only becomes substantial when defense decisions are highly decentralized.

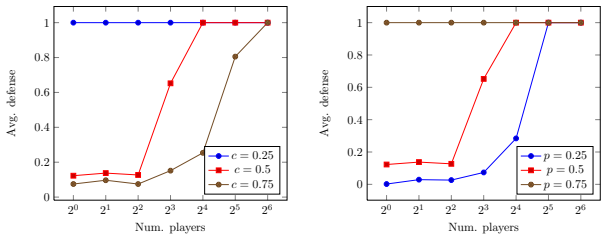

Figure 4: Average defense strategies for all players on an 8 x 8 grid.

6. REFERENCES

[1] Yoram Bachrach, Moez Draief, and Sanjeev Goyal. Security games with contagion, 2012. working paper.

[2] Tamer Basar and Geert Jan Olsder. *Dynamic Noncooperative Game Theory*. 2nd edition, 1999.

[3] Hau Chan, Michael Ceyko, and Luis E. Ortiz. Interdependent defense games: Modeling interdependent security under deliberate attack. In *Twenty-Eighth Conference on Uncertainty in Artificial Intelligence*, pages 152–162, 2012.

[4] Vincent Conitzer and Tuomas Sandholm. Computing the optimal strategy to commit to. In *Proceedings of the 7th ACM conference on Electronic commerce*, EC '06, pages 82–90, New York, NY, USA, 2006. ACM.

[5] Manish Jain, Erim Kardes, Christopher Kiekintveld, Milind Tambe, and Fernando Ordonez. Security games with arbitrary schedules: A branch and price approach. In *Twenty-Fourth National Conference on Artificial Intelligence*, 2010.

[6] Manish Jain, James Pita, Milind Tambe, Fernando Ordóñez, Praveen Paruchuri, and Sarit Kraus. Bayesian stackelberg games and their application for security at los angeles international airport. *SIGecom Exch.*, 7:10:1–10:3, June 2008.

[7] Manish Jain, Jason Tsai, James Pita, Christopher Kiekintveld, Shyamsunder Rathi, Milind Tambe, and Fernando Ordóñez. Software assistants for randomized patrol planning for the lax airport police and the federal air marshal service. *Interfaces*, 40:267–290, July 2010.

[8] David Kempe, Jon M. Kleinberg, and Éva Tardos. Maximizing the spread of influence in a social network. In *Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 137–146, 2003.

[9] Christopher Kiekintveld, Manish Jain, Jason Tsai, James Pita, Fernando Ordóñez, and Milind Tambe. Computing optimal randomized resource allocations for massive security games. In *Proceedings of the Eighth International Conference on Autonomous Agents and Multiagent Systems*, 2009.

[10] Dmytro Korzhyk, Zhengyu Yin, Christopher Kiekintveld, Vincent Conitzer, and Milind Tambe. Stackelberg vs. nash in security games: An extended investigation of interchangeability, equivalence, and uniqueness. *Journal of Artificial Intelligence Research*, 41:297–327, 2011.

[11] Howard Kunreuther and Geoffrey Heal. Interdependent security. *Journal of Risk and Uncertainty*, 26(2-3):231–249, 2003.

[12] Joshua Letchford and Yevgeniy Vorobeychik. Computing optimal security strategies for interdependent assets. In *Conference on Uncertainty in Artificial Intelligence*, pages 459–468, 2012.

[13] Praveen Paruchuri, Jonathan P. Pearce, Janusz Marecki, Milind Tambe, Fernando Ordóñez, and Sarit Kraus. Playing games with security: An efficient exact algorithm for Bayesian Stackelberg games. In *Proceedings of the Seventh International Conference on Autonomous Agents and Multiagent Systems*, pages 895–902, 2008.

[14] James Pita, Manish Jain, Fernando Ordóñez, Christopher Portway, Milind Tambe, Craig Western, Praveen Paruchuri, and Sarit Kraus. Using game theory for los angeles airport security. *AI Magazine*, 30(1):43–57, 2009.

[15] Eric Shieh, Rong Yang, Milind Tambe, Craig Baldwin, Joseph DiRenzo, Ben Maule, and Garrett Meyer. PROTECT: A deployed game theoretic system to protect the ports of the United States. In *Proceedings of the Eleventh International Conference on Autonomous Agents and Multiagent Systems*, pages 13–20, 2012.

[16] Yevgeniy Vorobeychik and Michael P. Wellman. Stochastic search methods for nash equilibrium approximation in simulation-based games. In *Seventh International Conference on Autonomous Agents and Multiagent Systems*, pages 1055–1062, 2008.