SAND2016-9222C

Network Security Games with Probabilistic Evasion

Paper 1525

Abstract

Stackelberg or defender-attacker games have recently become one of the main tools used to model security decisions in adversarial settings such as network security games. In these games the adversarial nature of these interactions leads to a great deal of uncertainty that has not been successfully captured in existing network models. To address this, we propose a model of attack interdiction in network settings that takes into account outcome uncertainty for the attacker and give a double oracle formulation for solving games in this setting. Finally, we show both theoretically and experimentally that ignoring this uncertainty has the potential to significantly degrade solution quality.

Introduction

In recent years, game theoretic techniques have been used to model adaptive intelligent adversaries in a wide range of critical infrastructure settings, from airports (Jain et al. 2010) to ports (Fang, Jiang, and Tambe 2013). Network security games (NSG) focus on a subset of these problems, namely those that are most naturally modeled as a graphical network, such as placing check points on road networks (Jain, Conitzer, and Tambe 2013) or choosing patrol routes in the waterways near our major cities (Vorobeychik, An, and Tambe 2012). One important element of NSGs is the asymmetry between the evader (attacker) and the interdictor (defender). It is commonly assumed that evader's ability to choose the day of attack and observe past security choices forces the interdictor to *commit* to a defense strategy. However, introducing randomness into the decision space of the interdictor adds uncertainty on the side of the evader, as the evader is not certain exactly how these random choices will be evaluated on the day of their attack. This is modeled as the interdictor committing to a mixed strategy and the evader best responding to the aggregate strategy, rather than the individual pure strategies that make up the support.

While uncertainty has been explored in the standard Stackelberg model (Nguyen, Jiang, and Tambe 2014), much of the work in network security games (Letchford and Vorobeychik 2013; Jain, Conitzer, and Tambe 2013) has been primarily focused on settings where interdiction success is binary; once an edge or a target is defended, the interdictor is guaranteed an interdiction if the evader chooses a strategy that uses the defended edge or target. However, in a real world setting, interdiction is rarely guaranteed. These models miss an important aspect in not modeling uncertainty in the success of the interdiction. In general, evader actions often have some risk associated with them and defense mitigations rarely are 100% effective. As an example, consider the task of deciding where to place radiological sensors to detect and interdict the smuggling of radioactive materials such as dirty bombs. These sensors are well known for having both false positives and false negatives (Cochran and McKinzie 2008), thus a binary success model here is overly optimistic. In fact it is possible to construct a game where this loss is unbounded. In Preliminaries we show that the potential loss to the interdictor for making this optimistic assumption is unbounded and in Computational Experiments we further explore this expected loss experimentally on random graphs.

In this work we relax this problematic assumption. In particular, we propose a model where an evader is interested in traversing from one node within a set of source node to one node within a set of target nodes, each with an associated payoff. Each arc in the graph has a baseline evasion probability, which represents the chance of successfully traversing the arc even if no additional resources are invested in that arc. Moreover, we assume that the defender, subject to a limited budget, has the ability to impose additional security on a subset of these arcs, further reducing the probability of evasion within this subset.

In addition to the work discussed above, there has been extensive work in a number of related domains. One area of research (Letchford and Vorobeychik 2012; Tsai et al. 2013) has focused on settings with probabilistic success rates, but without a focused adversary. Instead, the adversary resembles an epidemiological process, interested in spreading across the network. A second area of research focuses on patrolling (Basilico, Nittis, and Gatti 2016) or hider-seeker (Halvorson, Conitzer, and Parr 2009) games, where both the adversary and the defender are mobile. However, the uncertainty considered here generally focuses on the alarms that an adversary might trip, such as the alarm failing to go off or giving spatially uncertain data. Finally, there has been other work exploring uncertainty in generic

Stackelberg games, such as adding uncertainty over previous actions in repeated games (Nguyen, Alpcan, and Telekom 2008) or assuming that it may not be optimal (or even possible) for the attacker to perform perfect surveillance before attacking (An et al. 2012).

Preliminaries

We consider a *Probabilistic Network Security Game* (PNSG) on a digraph $D(N, A)$ with node set N and arc set A, where the leader plays the role of the network interdictor and the follower plays the role of a network evader. The interdictor's goal is to intercept/capture the evader, whereas the evader's goal, in direct contrast, is to successfully traverse the network, starting at a source node and terminating at a target node, without being captured. Each target node is associated with a payoff and the objective of the evader is to maximize the expected payoff. The expected payoff is the product of the expected probability of success in navigating a path to a target and that target node's payoff. The interdictor's diametrically opposed objective is to minimize the evader's expected payoff.

In this work, we deviate from the standard assumption that defenses are perfect (i.e., an evader traversing a defended arc will be interdicted with certainty). We assume that each arc $(u, v) \in A$ has a baseline evasion probability p_{uv} and a defended evasion probability p'_{uv}, with $1 \geq p_{uv} \geq p'_{uv} \geq 0$. The standard binary variant (SBV) of NSG, with $p_{uv} = 1$ and $p'_{uv} = 0$ for all $(u, v) \in A$, is thus a special case of PNSG. We also assume that arc evasion probabilities are independent and, therefore, the probability of successfully traversing a path j is given by the product of the evasion probabilities for each arc in the path. The utility to the evader for traversing path j with target node t_j is the product of the path evasion probability times the payoff of the target node ω_{t_j}. Conversely, the expected payoff to the interdictor is the product of the path evasion probability times the payoff (loss) of the target node $-\omega_{t_j}$.

A simple example with two edge-disjoint paths are shown in Figure 1. For each arc, two evasion probabilities (baseline and defended) are shown. In Table 1, example evader payoffs for two specific paths are shown under three cases: no defense (none), arc $(3, t_1)$ defended, and arcs $(3, t_1), (1, 4)$ defended. The interdictor allocates defensive resources at cost c_{uv} for all $(u, v) \in A$ subject to budget constraint Γ_b. $\boldsymbol{X} = \{\boldsymbol{x}^1, \cdots, \boldsymbol{x}^n\}$ is a restricted set of defense allocations and $\boldsymbol{x}^i \in \{0, 1\}^{|A|}$ for all $i = 1, \cdots, n$ is a valid defense allocation, with $x^i_{uv} = 1$ if arc (u, v) is defended and 0 otherwise. Similarly, the evader chooses a path from a source node to a target node to maximize expected payoff, which is a product of path evasion probability times the payoff of the target node. $\boldsymbol{Y} = \{\boldsymbol{y}^1, \cdots, \boldsymbol{y}^m\}$ is a restricted set of paths and $\boldsymbol{y}^j \in \{0, 1\}^{|A|}$ for all $j = 1, \cdots, m$ prescribes a valid path with $y^j_{uv} = 1$ if arc (u, v) is in path j and 0 otherwise.

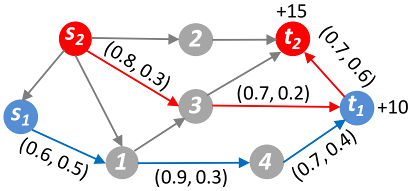

Figure 1: PNSG with two source nodes $\{s_1, s_2\}$ and two target nodes $\{t_1, t_2\}$ with payoffs 10 and 15, respectively.

Table 1: Payoffs of path 1 and path 2 under different defense allocations

		Expected Payoff		
$\omega_{t\cdot}$	path	none	$\{(3, t_1)\}$	$\{(3, t_1), (1, 4)\}$
+10	$(s_1, 1), (1, 4), (4, t_1)$	3.78	3.78	1.26
+15	$(s_2, 3), (3, t_1), (t_1, t_2)$	5.88	1.68	1.68

Theorem 1 *There exists PNSG where optimal solution to the SBV version of game is arbitrarily worse than the optimal solution to the PNSG.*

Proof We can translate any graph with non-binary edge probabilities to one in SBV by transforming every arc (u, v) : $p_{uv} = 1$ and $p'_{uv} = 0$. It is sufficient to show that there exists a game where the optimal solution for this transformed game (evaluated against the original game) is arbitrarily worse than the optimal solution computed for the original game. Consider the game pictured in Figure 2. For the SBV version of this game, the optimal defense solution under a budget of 2 is to defend (s_1, t_1) and $(s_1, 1)$, as this appears to fully defends all targets. However, this still has an expected payoff of $(1 - \epsilon)$ in the actual graph. For the PNSG version of the game, the optimal solution is to defend edge $(1, t_3)$ 100% of the time and to defend edges $(1, t_2)$ and (s_1, t_1) each 50% of the time (which is equivalent to the following mixed strategy: .5 $((1, t_3)$,$(1, t_2))$ + .5 $((1, t_3)$,$(s_1, t_1)))$. This fully defends target t_3 and prevents attacks against each of the other targets half of the time, leading to an expected loss of $\frac{\epsilon}{2}$. Thus, if we consider the ratio of expected loss between these two methods: $\frac{1-\epsilon}{\frac{\epsilon}{2}} \rightarrow \infty$ as $\epsilon \rightarrow 0$. □

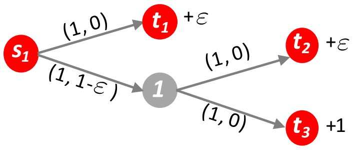

Figure 2: Simple counter example

Double Oracle Algorithm for PNSG

We employ a double oracle algorithm which is initialized with arbitrary sets of defense allocations $\boldsymbol{X}$ and paths $\boldsymbol{Y}$. At each iteration, we solve an LP $Minimax(\boldsymbol{X}, \boldsymbol{Y})$, which returns equilibrium strategies $\boldsymbol{d}$ and $\boldsymbol{a}$ for the interdictor and the evader, respectively. The *Interdictor's Best Response* oracle, $IBR(\boldsymbol{a})$, is solved to generate a pure strategy best response against evader mixed strategy $\boldsymbol{a}$. Analogously, the *Evader's Best Response* oracle, $EBR(\boldsymbol{d})$, is solved to generate a pure strategy best response against defender mixed strategy $\boldsymbol{d}$. We note that at each iteration, $IBR(\boldsymbol{a})$ and $EBR(\boldsymbol{d})$ *implicitly* search for pure strategy best responses amongst all valid interdictor defense allocations and evader paths, respectively. The algorithm terminates when at a given iteration, the two best response oracles cannot produce a solution (payoff) that is better than those prescribed by $Minimax(\boldsymbol{X}, \boldsymbol{Y})$. The validity of the best-response double oracle algorithm was established by (McMahan and G.J. Gordon 2003) and we refer the reader to convergence proof therein.

Minimax Formulation

Given the current sets of interdictor's defense allocations $\boldsymbol{X}$ and evader's paths $\boldsymbol{Y}$, an equilibrium for the restricted game can be found by solving a linear program (LP) that computes the maximin strategy for the interdictor. u^* represents the utility for the interdictor and $\boldsymbol{d} = (d^1, \cdots, d^n)$ represents the interdictor's mixed strategy over $\boldsymbol{X}$.

$$\max_{u^*, \boldsymbol{d} \geq \boldsymbol{0}} \quad u^* \tag{1a}$$

$$\text{s.t.} \quad u^* \leq -\sum_{i=1}^{n} u(\boldsymbol{x}^i, \boldsymbol{y}^j) d^i \quad \forall j = 1, \cdots, m \tag{1b}$$

$$\boldsymbol{1}^\top \boldsymbol{d} = 1 \tag{1c}$$

$-u(\boldsymbol{x}^i, \boldsymbol{y}^j)$ is the interdictor's payoff when playing $\boldsymbol{x}^i$ against $\boldsymbol{y}^j$, where

$$u(\boldsymbol{x}^i, \boldsymbol{y}^j) = \omega_{t_j} \prod_{(u,v) \in A_j} \max \left\{ p_{uv}(1 - x^i_{uv}), p'_{uv} \right\} \tag{2}$$

In (2), ω_{t_j} is the payoff associated with the target node t_j. Since success probabilities are independent, the probability of successfully traversing the path is the product of the independent arc evasion probabilities. For each arc $(u, v) \in A_j$ (i.e. arcs in path j), $\max\{p_{uv}(1 - x^i_{uv}), p'_{uv}\}$ provides the success probability of traversing arc (u, v) given interdictor's defense allocation x^i_{uv}. If $x^i_{uv} = 1$, then $\max\{p_{uv}(1 - x^i_{uv}), p'_{uv}\} = p'_{uv}$ and p_{uv} otherwise. For each evader path in $\boldsymbol{Y}$, (1b) computes mixed strategy $\boldsymbol{d}$'s payoff against that path. The evader's mixed strategy $\boldsymbol{a}$ can be easily extracted from the solutions of (1) as they are associated with the dual variables corresponding to constraints (1b).

Interdictor's Best Response Oracle (IBR)

The objective of IBR is to find the best pure strategy defense allocation $\boldsymbol{x}$ over arc set A, bounded by defense budget Γ_b, which minimizes the evader's expected payoff under mixed strategy $\boldsymbol{a}$. The IBR oracle is defined on an augmented network with two auxiliary nodes, a starting node s and a terminal node t. For fixed defense allocation $\boldsymbol{x}$, the evader solves a path selection problem to find the path that maximizes expected payoff. In order to formulate this path selection problem as an LP (i.e., *maximum flow problem*), we introduce an arc from node s to every source node and an arc from every target node to node t. For each $(u, v) \in A$, we define two flow variables f_{uv} and f'_{uv} with probabilities p_{uv} and p'_{uv} to represent the baseline and defended case, respectively. In Figure (3), (a) shows a simple path from s to t and (b) shows the same path represented in the augmented network with arc $(1, 4)$ defended. If an arc (u, v) is defended (i.e., $x_{uv} = 1$), we restrict flow on f_{uv} to zero. If (u, v) is undefended, flows are permitted on both f_{uv} and f'_{uv}; however, objective pressure to maximize flow will always induce flows on f_{uv}, which has a higher evasion probability. Unlike standard network flow models in which flow in equal flow out, in our probabilistic network, arc inflows are adjusted by their evasion probabilities. In Figure (3) (c), sending one unit of flow into node s_1 will result in $1 \times 0.6 \times 0.3 \times 0.7 = 0.126$ unit of flow out of target node t_1, which is exactly the evasion probability of the path. IBR can thus be formulated as the following bilevel program, in which the interdictor (upper-level) selects the optimal pure strategy best response defense allocation and the evader (lower-level) solves a set of maximum flow problems (weighted by $\boldsymbol{a}$) over the augmented network.

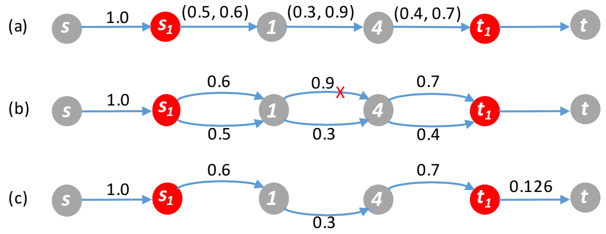

Figure 3: (a) A simple path (b) Augmented network representation (c) Probabilistically adjusted network flow

$$\min_{\boldsymbol{x}} \max_{\boldsymbol{f}, \boldsymbol{f'} \geq \boldsymbol{0}} \sum_{j=1}^{m} a_j \omega_{t_j} f^j_{t_j t} \tag{3a}$$

$$\text{s.t.} \sum_{(u,v) \in A_j} (f^j_{uv} + f'^j_{uv}) - \sum_{(v,u) \in A_j} (p_{vu} f^j_{vu} + p'_{vu} f'^j_{vu}) = b_u \ \forall u \in N \setminus t, \forall j \tag{3b}$$

$$f^j_{uv} \leq 1 - x_{uv} \quad \forall (u, v) \in A_j, \forall j \tag{3c}$$

$$\sum_{(u,v) \in A} c_{uv} x_{uv} \leq \Gamma_b \tag{3d}$$

The interdictor's objective (3a) is to minimize the expected payoff of the evader, which is a weighted sum of the evader's utility over mixed strategy $\boldsymbol{a}$. Constraints (3b) are nodal balanced constraints where in-flows have been adjusted to reflect arc evasion probabilities. b_u represents the supply or demand at node u (where $b_u = 1$ if $u = s$ and 0 otherwise). (3c) restricts flows on f_{uv} to be zero for all defended arcs. Finally, (3d) is the constraint on the overall defense budget.

Lemma 1 *For a fixed defense allocation $\boldsymbol{x}$, the inner maximization problem of (3) correctly computes the evader's expected utility for playing mixed strategy $\boldsymbol{a}$ over $\boldsymbol{Y}$.*

Proof It suffices to show that for a given path j prescribed by arcs $(u, v) \in A_j$, we have $f^j_{t_j t} = \prod_{(u,v) \in A_j} \max \left\{ p_{uv}(1 - x^i_{uv}), p'_{uv} \right\}$. This equality is trivially satisfied because for any arc $(u, v) \in A_j$, constraints

(3c) ensure that flows occur only on the lower probability defended arc f'_{uv} if $x_{uv} = 1$. On the contrary if $x_{uv} = 0$, non-zero flows are permitted on both f_{uv} and f_{uv}; however, objective pressure to maximize flow will ensure usage of f_{uv}, which has the higher evasion probability p_{uv}. Finally, nodal balance constraints (3b) ensures that flows *into* arcs $(u, v) \in A_j$ are proportionally discounted by the associated evasion probability, p_{uv} or p'_{uv}. Thus, a one unit injection into node s will result in an outflow of $\prod_{(u,v)\in A_j} \max \left\{ p_{uv}(1 - x^i_{uv}), p'_{uv} \right\}$ on arc $f^j_{t_j,t}$. □

Theorem 2 *The bilevel program* (3) *correctly computes a best-response for the interdictor under the probabilistic setting.*

Proof The upper-level problem searches over all valid defense allocation $\boldsymbol{x}$ with total cost less than or equal to Γ_b (3d). Thus, it suffices to show that for any fixed interdictor's defense allocation $\boldsymbol{x}$, the lower-level correctly computes the evader's expected payoff. This follows directly from Lemma 1. □

(3) is a bilevel program that cannot be solved directly, therefore, we solve a mixed-integer linear programming (MILP) reformulation of IBR instead.

Theorem 3 *Bilevel program* (3) *has an equivalent MILP reformulation.*

Proof The proof is by construction. Given a fixed upper-level decision $\boldsymbol{x}$, the lower-level problem of (3) is a feasible and bounded LP. By strong duality, we can replace the inner maximization problem with its equivalent dual minimization problem to arrive at the following single-level program.

$$\begin{aligned} \min_{\boldsymbol{x},\boldsymbol{\alpha},\boldsymbol{\beta}\geq\mathbf{0}} \; & \sum_{j=1}^{m} b_s \alpha^j_s + \sum_{j=1}^{m} \sum_{(u,v)\in A_j} (1 - x_{uv})\beta^j_{uv} \\ \text{s.t. } & \alpha^j_u - p_{vu}\alpha^j_v + \beta^j_{uv} \geq 0 \, \forall (u, v) \in A_j \setminus (t_j, t), \forall j \\ & \alpha^j_{t_j} - \alpha^j_t + \beta^j_{t_j t} \geq a_j \omega_{t_j} \; \forall j \qquad (4) \\ & \alpha^j_u - p'_{uv}\alpha^j_v \geq 0 \, \forall (u, v) \in A_j, \forall j \\ & \boldsymbol{x} \in \{0, 1\}^{|A|} \end{aligned}$$

where $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are dual variables associated with constraints (3b) and (3c). The objective of (4) contains bilinear terms, which are products of non-negative $\boldsymbol{\beta}$ and binary $\boldsymbol{x}$ variables. We introduce a new non-negative variable $\boldsymbol{\gamma}$ and the following disjunctive constraints to linearize these bilinear terms.

$$\beta^j_{uv} - M x_{uv} \leq \gamma^j_{uv} \qquad (5a)$$

$$\gamma^j_{uv} \leq \beta^j_{uv} + M x_{uv} \qquad (5b)$$

$$\gamma^j_{uv} \leq M(1 - x_{uv}) \qquad (5c)$$

In (5), M is a coefficient chosen to be sufficiently large to upper bound dual variables $\boldsymbol{\beta}$. By substituting γ^j_{uv} for $\beta^j_{uv} x_{uv}$ terms in the objective and introducing constraints (5) for all $(u, v) \in A$ and $j = 1, \cdots, m$, we arrive at the desired single-level MILP reformulation. □

Evader's Best Response Oracle (EBR)

The objective of EBR is to find the best pure strategy path $\boldsymbol{y}$, starting at a source node and ending at a target node, that maximizes the evader's expected payoff given the defender's mixed strategy $\boldsymbol{d}$ over $\boldsymbol{X}$. EBR can be formulated as a MILP composed of a set of weighted maximum-flow problems.

$$\max_{\boldsymbol{y}\in\{0,1\}^{|A|},\boldsymbol{f},\boldsymbol{f'}\geq\mathbf{0}} \sum_{i=1}^{n} d_i \Big(\sum_{(u,t)\in A} \omega_u f^i_{ut} \Big) \qquad (6a)$$

$$\text{s.t.} \sum_{(u,v)\in A} y_{uv} - \sum_{(v,u)\in A} y_{vu} = b_u \; \forall u \in N \setminus t \qquad (6b)$$

$$\sum_{(u,v)\in A} (f^i_{uv} + f'^i_{uv}) - \sum_{(v,u)\in A} (p_{vu} f^i_{vu} + p'_{vu} f'^i_{vu}) \qquad (6c)$$

$$= b_u \; \forall u \in N \setminus t, \forall i$$

$$f^i_{uv} + f'^i_{uv} \leq y_{uv} \quad \forall (u, v) \in A, \forall i \qquad (6d)$$

$$f^i_{uv} \leq 1 - x^i_{uv} \quad \forall (u, v) \in A, \forall i \qquad (6e)$$

The objective of the evader is the select a path $\boldsymbol{y}$ that maximize the evader's utility, which is the sum of the utility of the chosen path $\boldsymbol{y}$ weighted against the interdictor's mixed strategy $\boldsymbol{d}$ over $\boldsymbol{X}$. Constraints (6b) are nodal balance constraints for *path selection*, which prescribes $\boldsymbol{y}$. (6c) are probabilistically adjusted flow balance constraints under interdictor's defense allocations (i.e., for $i = 1, \cdots, n$). Constraints (6d) and (6e) are constraints restricting flows to arcs in the selected path $\boldsymbol{y}$ and not "shut-off", respectively. In summary, the selection of the best pure strategy response path $\boldsymbol{y}$ restricts flows to arcs in that path, but for each defense allocation i, the flows on that path is specifically adjusted to account for the defense in place.

Theorem 4 *MILP* (6) *correctly computes a best-response for the evader under the probabilistic setting.*

Proof Constraints (6b) define a valid path prescribed by binary variables $\boldsymbol{y}$, and constraints (6d) ensure that flows are permitted only on arcs prescribed by $\boldsymbol{y}$ (restricting flows to the chosen path). Then, it suffices to show that for any fixed $\boldsymbol{y}$, (6) correctly computes the evader's expected payoff. This follows directly from Lemma 1. □

Computational Experiments

We now present computational results on some randomly generated instances to demonstrate the effectiveness of the proposed models and algorithms. All experiments were run with CPLEX 12.5 on a machine with 4-16 core 2.70GHz processors and 512 GB RAM using a maximum of 16 threads and a max cutoff of 1 hour. We evaluated our formulations on two different graph structures, Erdős-Rényi (ER) random graphs with $p = .25$, and a directed grid-like setting where we divided the nodes into a series of bins. Each node can only have arcs to adjacent bins, and the source and target nodes are on opposite sides of the grid. This second structure generates graphs where all possible paths are of equal length, avoiding generating graphs with small number

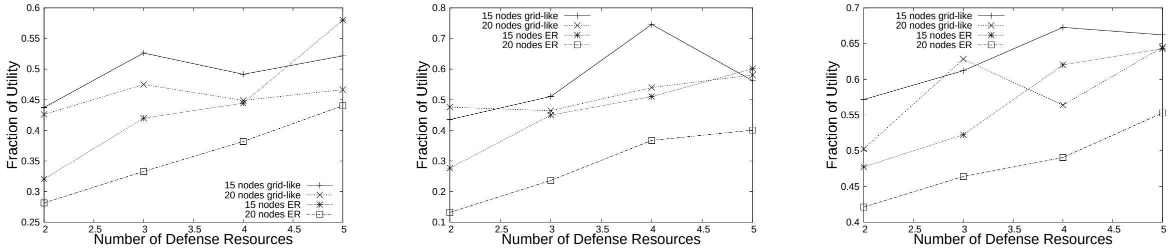

(a) $p: .8, p': .6$

(b) $p: .4, p': .2$

(c) $p: .8, p': .4$

Figure 4: Solution quality calculated against SBV as a fraction of possible utility gain

of clearly optimal paths to defend. Unless otherwise mentioned, the probability of the attacker successfully transitioning an undefended (defended) arc varies between .8 and .9 (.6 and .7), and target utilities vary between 10 and 20. All results reported are averaged from 30 runs.

Solution Quality

First, to capture how the solutions generated against PNSG and SBV on random graphs compare, we measured the difference in solution quality between these two models on graphs with uncertainty (as on graphs without uncertainty both algorithms would find identical solutions). To do this, we define the utility guaranteed to the defender by the following three defense strategies as:

- U(Null): Zero defense resources evaluated against PNSG
- U(PNSG): The optimal defense allocation calculated against PNSG and evaluated against PNSG
- U(SBV): The optimal defense allocation calculated against SBV and evaluated against PNSG

We calculated the fraction of the possible utility gain over the null setting that the binary setting achieved: $\frac{U(SBV)-U(Null)}{U(PNSG)-U(Null)}$. Figure 4a shows how this values varies with the number of defense resources for both grid-like and ER random graphs. We found that on average, the optimal defense allocation calculated against SBV only seems to capture about half of the potential utility gain and seems to perform worse as the size of the graph increases. Additionally, while the efficency of the solution generated via SBV seems to improve significantly as we increase the number of defense resources in the ER graphs, we do not see a similar improvement in the grid-like case.

Next, we examined how the quality of the solution found via SBV changed as we varied the baseline utility level without defense (p) and with defense (p') (with the same .1 noise levels). The above corresponds to values of (p: .8, p': .6). We considered p values of $\{.8,.6,.4\}$ and p' values of $\{p-.1, p-.2, p-.4\}$. Figure 4b demonstrates how these utilities change when we change p to .4, but hold $p-p'$ constant. We see that while grid-like results did not significantly change, however the solutions calculated against SBV seem significantly worse on ER graphs with smaller number of resources. This is likely caused by an increased importance in path length in expected utility, which is not captured properly in the SBV abstraction. Figure 4c shows what happens when we hold p constant at .8 but instead decrease p' to .4. Not suprisingly, since this setting is a better approximation of the binary case, the SBV solutions peformed almost 10% better uniformly across all data points. The remaining combinations of p and p' appear in Figures 5-6.

Runtime

Figures 7 (8) shows how runtime increases as we increase the number of nodes in the grid-like (ER) graph and the number of defense resources. Unsurprisingly, runtime increases with both number of nodes and number of defense resources. Perhaps more interesting is how much quickly the runtime increases in the grid-like graphs. This can intuitively be explained by the fact that there is higher variance in path quality in the ER graph (as the grid-like graph forces all paths to be of the same length). In fact, a large number of these paths will never be a best response for the adversary under any defense strategy. This reduces the number of iterations and the number of paths that needs to be considered.

Conclusion

We have shown that it is crucial to consider uncertainties of interactions between the evader and the interdictor in adversarial settings. This work is a significant advance towards addressing uncertainties in NSG. Specifically, we present novel models and algorithms for NSG with probabilistic evasion to account for the fact that in most cases the success of defense mitigation against specific attacks is not a binary outcome. Additionally, we demonstrated that solutions generated for the SBV (i.e., ignoring uncertainties) are significantly worst then solutions generated for the PNSG (i.e., with uncertainties). This is shown to be true both theoretically and experimentally, which highlights the significance in modeling uncertainty for NSG. Our computational results show that while we can solve medium sized instances, considerable room for improvement still exists. Developing heuristics or approximations to use as oracles has been shown to work in related problems and is a promising direction for future work.

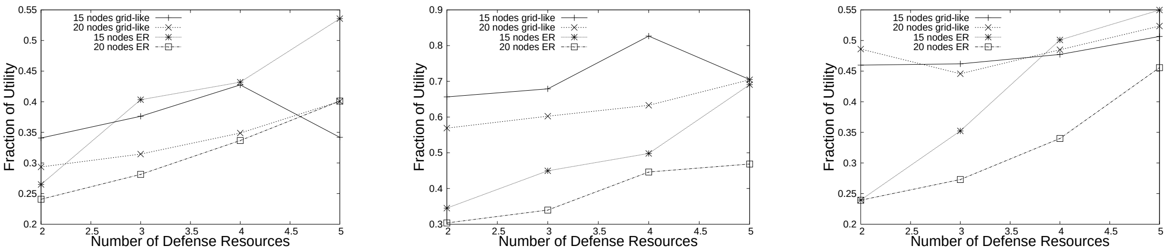

(a) p: .8, p': .7

(b) p: .6, p': .2

(c) p: .6, p': .4

Figure 5: Solution quality calculated against SBV as a fraction of possible utility gain

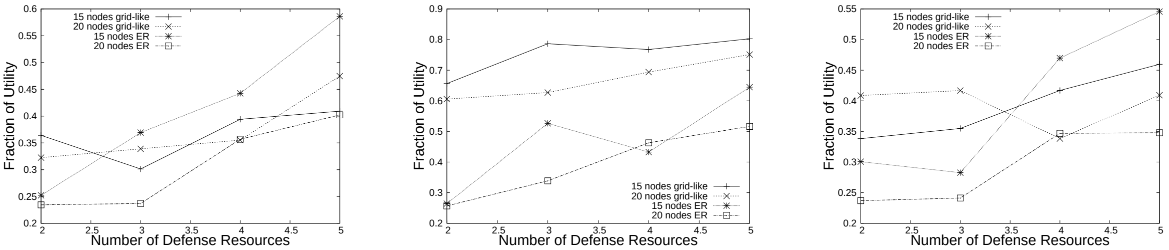

(a) p: .6, p': .5

(b) p: .4, p': 0

(c) p: .4, p': .3

Figure 6: Solution quality calculated against SBV as a fraction of possible utility gain

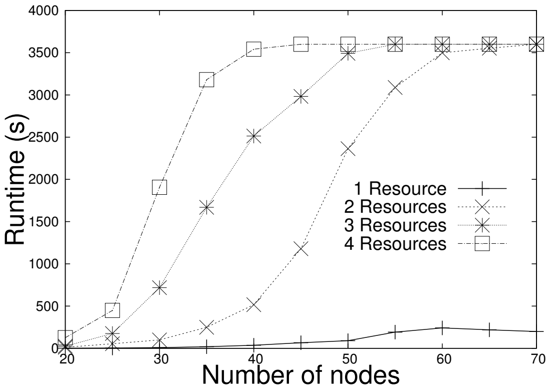

Figure 7: Runtime Results: Grid-like graphs

Acknowledgment Sandia National Laboratories is a multi-mission 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.

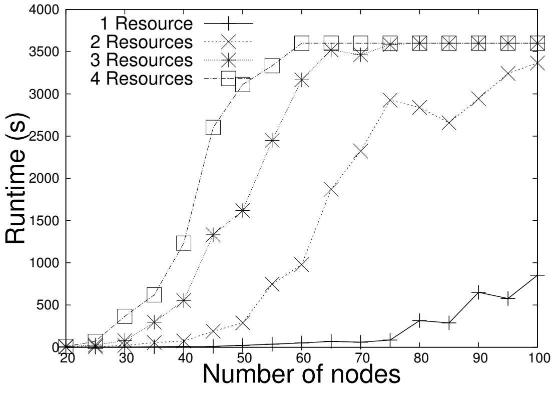

Figure 8: Runtime Results: Erdős-Rényi graphs

References

[An et al. 2012] An, B.; Kempe, D.; Kiekintveld, C.; Shieh, E.; Singh, S.; Tambe, M.; and Vorobeychik, Y. 2012. Security games with limited surveillance. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, AAAI'12, 1241–1248. AAAI Press.

[Basilico, Nittis, and Gatti 2016] Basilico, N.; Nittis, G. D.; and Gatti, N. 2016. A security game combining patrolling and alarm-triggered responses under spatial and detection uncertainties.

[Cochran and McKinzie 2008] Cochran, T., and McKinzie, M. 2008. Detecting nuclear smuggling. *Scientific American* 298:98–104.

[Fang, Jiang, and Tambe 2013] Fang, F.; Jiang, A. X.; and Tambe, M. 2013. Optimal patrol strategy for protecting moving targets with multiple mobile resources. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS '13, 957–964. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

[Halvorson, Conitzer, and Parr 2009] Halvorson, E.; Conitzer, V.; and Parr, R. 2009. Multi-step multi-sensor hider-seeker games. In *Proceedings of the 21st International Jont Conference on Artifical Intelligence*, IJCAI'09, 159–166. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.

[Jain et al. 2010] Jain, M.; Tsai, J.; Pita, J.; Kiekintveld, C.; Rathi, S.; Tambe, M.; and Ordóòez, F. 2010. Software assistants for randomized patrol planning for the lax airport police and the federal air marshal service. *Interfaces* 40(4):267–290.

[Jain, Conitzer, and Tambe 2013] Jain, M.; Conitzer, V.; and Tambe, M. 2013. Security scheduling for real-world networks. *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

[Letchford and Vorobeychik 2012] Letchford, J., and Vorobeychik, Y. 2012. Computing optimal security strategies for interdependent assets. In *The Conference on Uncertainty in Artificial Intelligence*, UAI.

[Letchford and Vorobeychik 2013] Letchford, J., and Vorobeychik, Y. 2013. Optimal interdiction of attack plans. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS '13, 199–206. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

[McMahan and G.J. Gordon 2003] McMahan, H., and G.J. Gordon, A. B. 2003. Planning int he presence of cost functions controlled by an adversary. In *International Conference on Machine Learning*, ICML '03, 98–104.

[Nguyen, Alpcan, and Telekom 2008] Nguyen, K. C.; Alpcan, T.; and Telekom, D. 2008. Security games with incomplete information. Technical report.

[Nguyen, Jiang, and Tambe 2014] Nguyen, T. H.; Jiang, A. X.; and Tambe, M. 2014. Stop the compartmentalization: Unified robust algorithms for handling uncertainties in security games. 317–324.

[Tsai et al. 2013] Tsai, J.; Qian, Y.; Vorobeychik, Y.; Kiekintveld, C.; and Tambe, M. 2013. Bayesian security games for controlling contagion. In *Proceedings of the Workshop on Multiagent Interaction Networks*, MAIN.

[Vorobeychik, An, and Tambe 2012] Vorobeychik, Y.; An, B.; and Tambe, M. 2012. Adversarial patrolling games. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 3*, AAMAS '12, 1307–1308. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.