

Solving Security Games on Graphs via Marginal Probabilities

Joshua Letchford and Vincent Conitzer

Duke University
Department of Computer Science
Durham, NC 27708, USA
{jcl,conitzer}@cs.duke.edu

Abstract

Security games involving the allocation of multiple security resources to defend multiple targets generally have an exponential number of pure strategies for the defender. One method that has been successful in addressing this computational issue is to instead directly compute the marginal probabilities with which the individual resources are assigned (first pursued by Kiekintveld et al. (2009)). However, in sufficiently general settings, there exist games where these marginal solutions are not implementable, that is, they do not correspond to any mixed strategy of the defender.

In this paper, we examine security games where the defender tries to monitor the vertices of a graph, and we show how the type of graph, the type of schedules, and the type of defender resources affect the applicability of this approach. In some settings, we show the approach is applicable and give a polynomial-time algorithm for computing an optimal defender strategy; in other settings, we give counterexample games that demonstrate that the approach does not work, and prove NP-hardness results for computing an optimal defender strategy.

Introduction

Algorithms for computing game-theoretic solutions are an important component of noncooperative multiagent systems. Such algorithms have recently started to be deployed in real-world security applications. They are used for the placement of checkpoints and canine units at Los Angeles International Airport (Jain et al. 2008; Pita et al. 2009), the assignment of Federal Air Marshals to international flights (Tsai et al. 2009), and the choice of patrol routes for the US Coast Guard in Boston (Shieh et al. 2012). Another application that is currently being pursued is the choice of schedules for fare inspectors on the Los Angeles Metro Rail system (Yin et al. 2012). In all these applications, choices are made in a randomized fashion because otherwise it would be easy to circumvent the security measures. Game theory provides a systematic approach to randomizing in an intelligent way. Specifically, the focus has been on computing *Stackelberg mixed strategies*, which are distributions that are optimal to play when the other player observes the distribution before playing.

All the security games above involve the allocation of multiple security resources to multiple targets (or, more generally, a single resource can be assigned to multiple targets, in which case the subset is referred to as a *schedule*). This results in exponentially many pure strategies for the defender, because every assignment of resources is a pure strategy. As a result, general linear or mixed integer program formulations for the Stackelberg problem (Conitzer and Sandholm 2006; Paruchuri et al. 2008; von Stengel and Zamir 2010) are exponential in size. A natural question to ask, however, is whether we really need to have a probability variable for every complete assignment of resources. Can we not restrict our attention simply to the *marginal* probability that a particular resource is assigned to a particular schedule? This would result in only a polynomial number of variables (one for each resource-schedule pair). In the context of security games, this approach was first pursued by Kiekintveld et al. (Kiekintveld et al. 2009), showing that indeed it is possible to find polynomial-size formulations that only refer to the marginal probabilities. Unfortunately, there are examples where this approach fails, in that the marginal probabilities returned do not correspond to any mixed strategy (distribution over assignments). On the other hand, Korzhyk et al. (Korzhyk, Conitzer, and Parr 2010) showed that the Birkhoff-von Neumann (BvN) theorem can be applied to show that under certain conditions, such a mixed strategy is guaranteed to exist and can be found efficiently. This work focused primarily on the relationship between the size of schedules, the applicability of the BvN theorem, and the complexity of computing Stackelberg strategies. For example, they showed that if schedules have size 1—so that resources are assigned to single targets—the BvN theorem applies, resulting in a polynomial-size linear program formulation for finding the marginal probabilities, together with an efficient algorithm for finding a corresponding mixed strategy.¹

In contrast, in this paper we focus not on the size of schedules, but rather on the structure of their relationship to each other. Specifically, many security games are naturally represented by a graph whose vertices are the targets and whose

¹A nonconstructive direct proof of the existence of a corresponding mixed strategy in this case was already given by Kiekintveld et al. (Kiekintveld et al. 2009).

edges are related to which schedules are feasible. For example, the vertices might represent stations of a subway system, and feasible schedules may be paths along the edges of the subway system graphs. Our approach is similar to that of Korzhyk et al. at a high level, but it turns out that for one of these problems, we need a generalization of the BvN theorem that was given by Budish et al. (Budish et al. 2013). This generalization allows us to characterize a new class of games where a mixed strategy corresponding to the marginal probabilities is guaranteed to exist, and gives us an algorithm for finding it. In another case, we show how to compute optimal marginal probabilities in a specific way, and show how to directly obtain a mixed strategy that corresponds to these marginal probabilities.

Background

We first review the standard algorithm for computing a Stackelberg mixed strategy in two-player normal-form games (Conitzer and Sandholm 2006; von Stengel and Zamir 2010). This algorithm creates a separate LP for each follower (attacker) pure strategy $t^* \in T$, which, in our case, is a target for the attacker to attack. An optimal solution of this LP gives the optimal leader (defender) strategy and utility, under the constraint that t^* is the attacker's best response. After solving this set of LPs, the solution of the LP with the highest objective value will be optimal overall. Letting $\alpha \in A$ denote the defender's pure strategies and U_d and U_a the defender and attacker's utility functions, the LP is as follows:

$$\begin{aligned} &\text{General LP} \\ &\text{maximize } \sum_{\alpha} p_{\alpha} U_d(\alpha, t^*) \\ &\text{subject to:} \\ &\forall t \in T : \sum_{\alpha} p_{\alpha} U_a(\alpha, t) \leq \sum_{\alpha} p_{\alpha} U_a(\alpha, t^*) \\ &\sum_{\alpha} p_{\alpha} = 1 \end{aligned}$$

Security games (Kiekintveld et al. 2009)

A security game has a set of targets T ($|T| = n$). A schedule $s \in S \subseteq 2^T$ consists of a subset of targets that can be simultaneously covered by a single defender resource. Defender resources (denoted by $\omega \in \Omega$) can be either homogeneous, meaning that any defender resource can cover any schedule, or heterogeneous, meaning that, for each resource ω , there is a set of schedules $A(\omega)$ that that resource can cover. A target t is said to be *covered* if a resource has been assigned to a schedule that covers it. Finally, the utility of the defender (attacker) is denoted as $U_d^c(t)$ ($U_a^c(t)$) if a target is covered and $U_d^u(t)$ ($U_a^u(t)$) if it is not covered. We have $U_d^c(t) \geq U_d^u(t)$ and $U_a^c(t) \leq U_a^u(t)$.

Compact security game LP

In the context of security games, the linear program presented above has exponentially many variables p_{α} (one for every assignment of resources). The following LP is an attempt to modify the above LP to instead use variables for the *marginal* probabilities $c_{\omega,s}$ that resource ω is assigned to schedule $s \in A(s)$. (A similar LP was given in (Korzhyk,

Conitzer, and Parr 2010), which in turn was a linear program reformulation of the MIP given in (Kiekintveld et al. 2009).) In this LP, c_t denotes the probability that target t is covered.

Compact LP

$$\begin{aligned} &\text{maximize } c_{t^*} U_d^c(t^*) + (1 - c_{t^*}) U_d^u(t^*) \\ &\text{subject to:} \\ &\forall \omega \in \Omega, \forall s \in A(\omega) : 0 \leq c_{\omega,s} \leq 1 \\ &\forall t \in T : c_t \leq \sum_{\omega \in \Omega, s \in A(\omega): t \in s} c_{\omega,s} \\ &\forall t \in T : c_t \leq 1 \\ &\forall \omega \in \Omega : \sum_{s \notin A(\omega)} c_{\omega,s} = 0 \\ &\forall \omega \in \Omega : \sum_{s \in A(\omega)} c_{\omega,s} \leq 1 \\ &\forall t \in T : c_t U_a^c(t) + (1 - c_t) U_a^u(t) \\ &\quad \leq c_{t^*} U_a^c(t^*) + (1 - c_{t^*}) U_a^u(t^*) \end{aligned}$$

The problematic part of this LP is that it is not clear that it makes sense to set $c_t = \sum_{\omega \in \Omega, s \in A(\omega): t \in s} c_{\omega,s}$; this calculation would be correct only if all the events where some ω covers some s with $t \in s$ are disjoint events. Indeed, in some games, the optimal marginal probabilities c_t cannot be achieved in any mixed strategy. In previous work, the Birkoff-von Neumann (Birkhoff 1946) theorem has been used to characterize some special cases where it is guaranteed that the marginal probabilities do correspond to some mixed strategy (Korzhyk, Conitzer, and Parr 2010). In this paper, we also have such a result, but we need a generalization of the BvN theorem that we discuss next. (Also note that we allow c_t to be lower than the total probability on schedules that cover t ; this corresponds to the common assumption that every subset of a schedule is also a schedule, i.e., we can always reduce our coverage on a target without affecting anything else.)

Bihierarchy extension of BvN Theorem (Budish et al. 2013)

For problems where there are two sets of objects X and Y (in our case, resources and schedules) and for every pair $(x, y) \in X \times Y$ a marginal probability $p_{x,y}$ (in our case, $c_{\omega,s}$) is given, this result gives a sufficient condition for there to exist a distribution over assignments of X to Y (in our case, a distribution over assignments of resources to schedules, i.e., a mixed strategy) that is consistent with the marginal probabilities.

In these problems, generally, for various subsets $Z \subseteq X \times Y$, there is a constraint of the form $\underline{qZ} \leq \sum_{(x,y) \in Z} p_{x,y} \leq \overline{qZ}$, where \underline{qZ} and \overline{qZ} are integers. Not only does this constraint hold on the probabilities, but we also would like it to hold for the realized assignments; that is, in such a realized assignment μ , we want $\sum_{(x,y) \in Z} b_{x,y} \in \{\underline{qZ}, \dots, \overline{qZ}\}$, where $b_{x,y} \in \{0, 1\}$ indicates whether x is matched to y in μ . (For example, in our problem, for each resource ω , we want $0 \leq \sum_{s \in A(\omega)} c_{\omega,s} \leq 1$, and moreover in the assignments over which we randomize, we need $\sum_{s \in A(\omega)} b_{\omega,s} \in \{0, 1\}$.) The result (Budish et al. 2013) gives a sufficient

condition on the family (or *constraint structure*) \mathcal{C} of sets Z (with $Z \in \mathcal{C}$) to guarantee that a distribution over partial matchings that has the right marginal probabilities and that satisfies all the constraints can be found, namely that \mathcal{C} is a bihierarchy, defined as follows. A constraint structure \mathcal{H} is a *hierarchy* if, for every pair of elements $Z, Z' \in \mathcal{H}$, we have $Z \subseteq Z'$, $Z' \subseteq Z$ or $Z \cap Z' = \emptyset$. Constraint structure \mathcal{B} is a *bihierarchy* if there exists hierarchies \mathcal{H}_1 and \mathcal{H}_2 such that $\mathcal{H}_1 \cup \mathcal{H}_2 = \mathcal{B}$ and $\mathcal{H}_1 \cap \mathcal{H}_2 = \emptyset$. (Budish et al. 2013) also gives an efficient flow-based algorithm for finding such a distribution.

Dimensions of the problem

In this section, we describe the different dimensions along which we vary the problem. The first dimension concerns whether the defender resources are homogeneous or heterogeneous. If we restrict ourselves to homogeneous defense resources, then any defender resource can cover any schedule.

The second dimension concerns the graph (whose vertices correspond to the targets and whose edges are related to the schedules). We consider:

- **Path graph.** A graph consisting of a single path.
- **Tree.** A graph with a single acyclic component.
- **General graph.** A general graph consisting of a single component.

In fact, for our positive results, we can also allow for multiple connected components in each case (e.g., forests instead of trees). We sometimes also require each component to have a distinguished root vertex.

The final dimension concerns how the graph restricts the possibilities for feasible schedules. We consider:

- **Edge.** Every schedule consists of two adjacent vertices in the graph.
- **Path.** Every schedule consists of a path in the graph. We also consider the further restrictions of: (1) only paths that pass through the root, and (2) only paths that start at the root.

We emphasize that not *every* subset of targets satisfying the requirement is a feasible schedule; rather only the converse is the case, that any subset *not* satisfying the requirement *cannot* be a feasible schedule. Our results are summarized in Figure 1.

Motivating examples. We now give two motivating examples for which we obtain positive results.

1. Consider a subway system without any cycles that is rooted at a central station. An inspector or guard can start a patrol from the central station, travel down a path, and then return along the same path. This corresponds to a tree graph with all feasible schedules being paths from the root.
2. Consider a high-speed rail line between two locations (such as the Japanese Shinkansen or the partially complete line between Beijing and Hong Kong). An inspector or guard can enter the train at some point, stay on it for

some number of stops, and get off. This corresponds to a path graph with all feasible schedules being (sub)paths.

Positive results

In this section, we cover the cases where we can solve for an optimal strategy in polynomial time via marginal probabilities:

1. Defender resources are heterogeneous, the graph is a set of rooted trees, and schedules correspond to paths starting at a root.
2. Defender resources are homogeneous, the graph is a path graph, and schedules correspond to paths.

Heterogeneous, set of trees, paths from the root

Theorem 1. *There exists a polynomial-time algorithm for finding the optimal Stackelberg strategy when defender resources are heterogeneous, the graph is a set of rooted trees, and schedules correspond to paths starting at a root.*

Proof. We start by solving the Compact LP above to obtain marginal probabilities. The following constraints on the $c_{\omega,s}$ variables hold:

1. $\forall t \in T : \lfloor \sum_{\omega \in \Omega, s \in A(\omega): t \in s} c_{\omega,s} \rfloor \leq \sum_{\omega \in \Omega, s \in A(\omega): t \in s} c_{\omega,s} \leq \lceil \sum_{\omega \in \Omega, s \in A(\omega): t \in s} c_{\omega,s} \rceil$
2. $\forall \omega \in \Omega : 0 \leq \sum_{s \in A(\omega)} c_{\omega,s} \leq 1$

The constraints under 2 clearly form a hierarchy, because their variable sets do not intersect with each other. When the graph is a set of rooted trees and schedules correspond to paths starting at a root, the constraints under 1 also form a hierarchy: if t is an ancestor of t' (i.e., t is on the path from the root to t'), then $\{(\omega, s) : s \in A(\omega) : t' \in s\} \subseteq \{(\omega, s) : s \in A(\omega) : t \in s\}$; if neither of t and t' is an ancestor of the other, then $\{(\omega, s) : s \in A(\omega) : t' \in s\} \cap \{(\omega, s) : s \in A(\omega) : t \in s\} = \emptyset$. Therefore, these constraints form a bihierarchy, and we can apply the algorithm from (Budish et al. 2013) to obtain a probability distribution over (deterministic) assignments of resources to schedules, where in each such assignment:

- each resource is used at most once (by the second constraint),
- every target for which $0 \leq \sum_{\omega \in \Omega, s \in A(\omega): t \in s} c_{\omega,s} < 1$ is covered either by 0 resources or by 1 resources, so that the coverage events are disjoint and the overall probability of covering t is in fact $\sum_{\omega \in \Omega, s \in A(\omega): t \in s} c_{\omega,s}$,
- every target for which $1 \leq \sum_{\omega \in \Omega, s \in A(\omega): t \in s} c_{\omega,s}$ is always covered by at least 1 resource.

We note that it is possible that $c_t < \sum_{\omega \in \Omega, s \in A(\omega): t \in s} c_{\omega,s}$; this can cause a problem in the case of $t = t^*$, which we may wish to defend less in order to lure the attacker to it. If so, we can simply move probability from a schedule s that places probability on t^* to schedule $s' = s \setminus \{t^*\}$ until the coverage on t^* reaches the desired

HOMOGENEOUS RESOURCES				
<i>Graph \ Schedule</i>	Edge	Path (start root)	Path (pass through root)	Path (general)
Path	Yes/P (KCP 2010)	Yes/P (Th 1)	Yes*/P (Th 2)	Yes*/P (Th 2)
Tree	Yes/P (KCP 2010)	Yes/P (Th 1)	No/? (Ex 3)	No/NP-h (Ex 3/(LCL 2006))
General	No/P (KCP 2010)	No/NP-h (Ex 2/Th 4)	No/NP-h (Ex 3/Th 4)	No/NP-h (Ex 3/(LCL 2006))

HETEROGENEOUS RESOURCES				
<i>Graph \ Schedule</i>	Edge	Path (start root)	Path (pass through root)	Path (general)
Path	No/NP-h (Ex 1/Th 3)	Yes/P (Th 1)	?	No/NP-h (Ex 1/Th 3)
Tree	No/NP-h (Ex 1/Th 3)	Yes/P (Th 1)	No/NP-h (Ex 3/Th 5)	No/NP-h (Ex 3/(LCL 2006))
General	No/NP-h (Ex 1/Th 3)	No/NP-h (Ex 2/Th 4)	No/NP-h (Ex 3/Th 5)	No/NP-h (Ex 3/(LCL 2006))

Figure 1: Summary of results. Every entry states, before the /, whether the Birkhoff-von Neumann property holds (“Yes” or “No”), that is, whether marginal probabilities can always be realized. (“Yes*” indicates that we only prove this for the marginal probabilities that result from our direct algorithm.) After the /, it states the complexity of finding an optimal Stackelberg strategy. All of the positive results hold even if the graph consists of multiple lines/trees/DAGs (that each have their own root in the rooted case), and all the negative results hold even if there is only one component.

value c_{t^*} .² (Note that we move this probability *after* obtaining a mixed strategy using the Budish et al. result.) By the constraints of the Compact LP, the resulting strategy will incentivize the attacker to attack t^* , and obtain the optimal value from this LP. \square

Homogeneous, path graph, general paths

Theorem 2. *There exists a polynomial-time algorithm for finding the optimal Stackelberg strategy when defender resources are homogeneous, the graph is a set of path graphs, and schedules correspond to (sub)paths.*

Proof. We start by solving the Compact LP above. However, for our argument, we will need a particular type of optimal solution for the marginal probabilities, which this LP is not guaranteed to give. The only parts of the LP solution that we will use are the attacker’s preferred target t^* and the utility that the attacker receives in this solution ($u_a = c_{t^*}^{\text{LP}} U_a^c(t^*) + (1 - c_{t^*}^{\text{LP}}) U_a^u(t^*)$). We now describe a method to generate a new marginal solution that results in the same utility for the defender as the LP solution and for which we can find a corresponding mixed strategy.

Let t_1, \dots, t_n be the targets in the order in which they appear in the path graph, from left to right. Initialize the current target coverage probabilities as $c_{t_j} = 0$ for all j . Now, for $j = 1$ to n , consider the schedule that covers t_j and extends as far as possible towards the right. If $c_{t_j} U_a^c(t_j) + (1 - c_{t_j}) U_a^u(t_j) > u_a$, then place on this schedule the amount of probability needed to make $c_{t_j} U_a^c(t_j) + (1 - c_{t_j}) U_a^u(t_j) = u_a$, and update the other targets’ coverage probabilities accordingly. We then assign the probability mass that we have now placed on this schedule to the first remaining available resource. If the schedule requires more probability mass than this resource has left, then cover as much as possible with this resource, and cover the remainder with the start of the next resource. Let $c_{\omega, s}$ be the

probability mass from resource ω that is assigned to schedule s at the end of the for loop. It is straightforward to see that this solution minimizes the number of resources needed to bring the attacker down to utility u_a . Hence, this solution uses at most the number of defender resources used in the LP.

Next, we need to guarantee that the attacker is incentivized to attack target t^* . If the attacker’s expected utility for attacking target t^* is u_a , then we are done with this step. Otherwise, starting with the last schedule s in the marginal solution that covers t^* , we can simply move probability from schedule s to schedule $s_0 = s \setminus \{t^*\}$ until $c_{t^*} U_a^c(t^*) + (1 - c_{t^*}) U_a^c(t^*) = u_a$. Since our marginal solution now has the same coverage on target t^* as the LP solution, and guarantees that each other target’s utility is at most u_a , it will give the defender the same expected utility.

Finally, we need to transform our marginal solution into a mixed strategy. This transformation can be intuitively described by the following process. Randomly draw a single number p uniformly from $[0, 1]$. For each resource ω , consider the schedule s that, in the sequential process described above, made ω ’s total assigned mass rise above p ; then, simply assign ω to this schedule s . The assignment chosen by this process corresponds to one pure strategy in the support of our mixed strategy, and the probability that this is chosen is the probability placed on this strategy in the mixed strategy. It is easy to see that this process indeed results in a marginal probability of $c_{\omega, s}$ that ω is assigned to s .³ \square

³Instead of performing this procedure explicitly when drawing an assignment, one can explicitly write the mixed strategy as follows. Consider all values $p_1, \dots, p_x \in [0, 1]$ at which one of the resources switches schedules. Then, for $i \in \{1, \dots, x + 1\}$, assign probability $p_i - p_{i-1}$ to the schedule that occurs when $p_{i-1} < p < p_i$, interpreting $p_0 = 0$ and $p_{x+1} = 1$. We note that $x + 1$, the total number of assignments with positive probability, can be at most $n + 1$: in the sequential process above, every one of the n vertices introduces at most one new schedule, and the very first one does not correspond to a cutoff point, so at the end of the sequential process there were at most n cutoff points; however, in the process of reducing the probability on t^* , we may have introduced a single additional cutoff point by splitting a schedule

²Dropping a target from a schedule in this way might correspond, for example, to an officer who is riding the train not inspecting a station when the train stops there.

Negative results

In this section, we give our negative results. We consider the following settings:

1. Defender resources are heterogeneous, the graph is a path graph, and schedules correspond to edges.
2. Defender resources are homogeneous, the graph is general, and schedules correspond to paths from the root.
3. The graph is a tree and schedules correspond to paths through the root.

In each of these settings, we first give an example instance where the optimal marginal solution does not correspond to any mixed strategy. (In setting 3, this example works even for homogeneous resources.) Then, we show that the problem of finding an optimal Stackelberg strategy is in fact NP-hard. (In setting 3, we are only able to show this hardness for heterogeneous resources.)

Heterogeneous resources, a path graph, schedules are edges

Counterexample 1. Consider the graph depicted in Figure 2a. There are two defender resources, one that can cover either edge (A, B) or edge (C, D) , and one that can cover either edge (B, C) or edge (D, E) . Suppose the utility functions are such that the defender would like to defend A with probability $1/2$, and C and D each with probability 1. (For example, suppose that the defender greatly prefers A to be attacked rather than any other target, the attacker is not interested in B or E , and if A is defended with probability $1/2$, the attacker would prefer attacking C or D unless these are defended with probability 1.) The numbers in the figure depict a marginal solution that achieves this. However, it is easy to see that there is no mixed strategy that achieves this: whenever A is defended by the first resource, it is impossible for the second resource to cover both C and D ; so, the latter two cannot be defended with probability 1 if A is defended with nonzero probability.

Theorem 3. Finding the optimal Stackelberg strategy is NP-hard when defender resources are heterogeneous, the graph is a path graph, and schedules correspond to edges.

Proof. We show this by a reduction from 3-COVER, in which we are given a set of elements $\{1, 2, \dots, m\}$ and n subsets of size 3, and the goal is to find $\frac{m}{3}$ sets that exactly cover all m elements.

Structure of the resources, path graph, and schedules.

We create one resource ω_j for every element $j \in \{1, 2, \dots, m\}$. For every one of the n subsets S_i , we create

s into s and $s \setminus \{t^*\}$, so $x \leq n$. All that remains to show is that a target either is never covered by more than one resource, or, if it sometimes is, then it is covered with probability 1. To show this, note that in the sequential process above, it is impossible that we add mass to a schedule containing resource t , then to a schedule not containing t , and then later again to a schedule containing t . Now, suppose that t is sometimes covered by two resources, ω_1 and ω_2 , where ω_1 is the earlier resource. Suppose this happens for $p = p_0$. Then, for $p > p_0$, ω_1 must cover t ; and for $p < p_0$, ω_2 must cover t .

three resources, $\omega_{S_i}^1, \omega_{S_i}^2, \omega_{S_i}^3$. Finally, we create $n - m/3$ interchangeable resources $\omega_\phi^1, \dots, \omega_\phi^{n-m/3}$.

For the sake of exposition, we in fact create multiple path graphs; it is easy to connect all these together into a single path graph using dummy edges. For each subset $S_i = \{j, k, l\}$, we create 4 path graphs $G_{S_i}^1, G_{S_i}^2, G_{S_i}^3, G_{S_i}^4$. We use G_{S_i} to denote the union of these four graphs. $G_{S_i}^1$ and $G_{S_i}^4$ have three vertices and two edges; $G_{S_i}^2$ and $G_{S_i}^3$ have four vertices and three edges. The utility functions are such that our goal is to cover every internal vertex (with two neighbors) with probability 1; the leaf vertices do not need to be defended. Any resource ω_ϕ^x can defend the leftmost edge of $G_{S_i}^1$; resource ω_j can defend the leftmost edge of $G_{S_i}^2$; resource ω_k can defend the leftmost edge of $G_{S_i}^3$; resource ω_l can defend the leftmost edge of $G_{S_i}^4$. Resource $\omega_{S_i}^1$ can cover either the rightmost edge of $G_{S_i}^1$ or the center edge of $G_{S_i}^2$; resource $\omega_{S_i}^2$ can cover either the rightmost edge of $G_{S_i}^2$ or the center edge of $G_{S_i}^3$; resource $\omega_{S_i}^3$ can cover either the rightmost edge of $G_{S_i}^3$ or the rightmost edge of $G_{S_i}^4$.

Proof of equivalence We now show that there exists a strategy that covers all the internal vertices with probability 1 if and only if there is a 3-cover. If there is a 3-cover, then consider the following pure strategy. If $S_i = \{j, k, l\}$ is in the cover, then we defend the leftmost edges of $G_{S_i}^2, G_{S_i}^3, G_{S_i}^4$ with resources $\omega_j, \omega_k, \omega_l$, respectively, and we defend the rightmost edges of $G_{S_i}^1, G_{S_i}^2, G_{S_i}^3$ with resources $\omega_{S_i}^1, \omega_{S_i}^2, \omega_{S_i}^3$, respectively, thereby covering all the internal vertices of these path graphs. If S_i is not in the cover, then we defend the leftmost edge of $G_{S_i}^1$ with some resource ω_ϕ^x , and we defend the middle edges of $G_{S_i}^2, G_{S_i}^3$ and the rightmost edge of $G_{S_i}^4$ with resources $\omega_{S_i}^1, \omega_{S_i}^2, \omega_{S_i}^3$, respectively, thereby covering all the internal vertices of these path graphs. Note that this uses exactly all the $n - m/3$ resources ω_ϕ^x . Thus, there exists a strategy that covers all the internal vertices with probability 1.

Conversely, suppose that there exists a strategy that covers all the internal vertices with probability 1. Any pure strategy on which this mixed strategy places positive probability must cover all the internal vertices; consider one such pure strategy. Consider some subset $S_i = \{j, k, l\}$ where this pure strategy does not allocate any of the resources ω_ϕ^x to $G_{S_i}^1$. To cover the center vertex of $G_{S_i}^1$, $\omega_{S_i}^1$ must be allocated to $G_{S_i}^1$. Then, to cover the left internal vertex of $G_{S_i}^2$, ω_j must be allocated to $G_{S_i}^2$. Then, to cover the right internal vertex of $G_{S_i}^2$, $\omega_{S_i}^2$ must be allocated to $G_{S_i}^2$. Then, to cover the left internal vertex of $G_{S_i}^3$, ω_k must be allocated to $G_{S_i}^3$. Then, to cover the right internal vertex of $G_{S_i}^3$, $\omega_{S_i}^3$ must be allocated to $G_{S_i}^3$. Then, to cover the center vertex of $G_{S_i}^4$, ω_l must be allocated to $G_{S_i}^4$. So, all of $\omega_j, \omega_k, \omega_l$ have been assigned to G_{S_i} (and cannot be used elsewhere). Now, because there are only $n - m/3$ resources ω_ϕ^x , there must be exactly $m/3$ subsets S_i for which none of the ω_ϕ^x resources are allocated to $G_{S_i}^1$. Moreover, these S_i cannot overlap on any element j , if they did, they would both have ω_j assigned




Figure 2

to them. Therefore, these S_i constitute a 3-cover. \square

Homogeneous, general graph, paths from the root

Counterexample 2. Consider the graph depicted in Figure 2b. The root is the node in the middle. There are 6 feasible schedules, each corresponding to a path from the root. They correspond to the (solid, dashed, dotted) edges to the (left, right) of the root. For example, one schedule covers the center (root) node, the top-left node, and the far-left node (the solid path to the left).

We have 3 homogeneous resources that we can assign to these schedules. The marginal solution picks each of these six possible schedules with a probability of $1/2$, resulting in a total defense probability of 1 for each target. However, it is easy to see that it is not in fact possible to cover all the targets all the time: for any assignment, either the left or the right side of the graph has only one resource assigned to it, and therefore an uncovered target.

Theorem 4. Finding the optimal Stackelberg strategy is NP-hard when defender resources are homogeneous, the graph is general, and schedules correspond to paths starting at a single root vertex.

This can be shown by reduction from 3-COVER; however, we omit this proof due to space considerations.

Tree graph, paths that pass through the root

For this case, our counterexample works even for homogeneous resources, but our NP-hardness result only works with heterogeneous resources; the complexity with homogeneous resources is open.

Counterexample 3. Consider the graph depicted in Figure 2c. The root of each tree is the middle vertex. There are six feasible schedules, each consisting of a path from one leaf vertex to another leaf vertex (covering both those leaves as well as the root of the corresponding tree). For example, one schedule covers targets A, D, and B.

We have 3 homogeneous resources that we can assign to these schedules. The marginal solution picks each of these six possible schedules with a probability of $1/2$, resulting in a total defense probability of 1 for each target. However, it is easy to see that it is not in fact possible to cover all the targets all the time: for any assignment, either the left or the right tree has only one resource assigned to it, and therefore an uncovered target.

While this example uses two trees, it is easy to turn it into an equivalent single-tree example, by merging D and W into a single root vertex.

Theorem 5. Finding the optimal Stackelberg strategy is NP-hard when defender resources are heterogeneous, the graph is a tree, and schedules correspond to paths that pass through the root.

This can be shown by reduction from tripartite matching; however, we omit this proof due to space considerations.

Conclusion

Many security games involve guarding the vertices of a graph by letting the defender resources travel across the edges of the graph. In this paper, we showed two cases where optimal Stackelberg strategies can be computed efficiently by first computing optimal marginal probabilities of assigning resources to schedules, and then computing a mixed strategy that is consistent with these marginal probabilities. In other cases, we showed that there does not necessarily exist a mixed strategy corresponding to the optimal marginal probabilities, and also gave explicit NP-hardness results for computing optimal Stackelberg strategies.

Overall, the results in this paper suggest that it is rare to find a case where there does not necessarily exist a mixed strategy corresponding to the optimal marginal probabilities, but nevertheless an optimal Stackelberg strategy can be found in polynomial time. (One such case does appear in (Korzhyk, Conitzer, and Parr 2010): homogeneous resources, arbitrary schedules with size 2.) Future research could focus on identifying other strategies where the constraint structure is a bihierarchy. Another interesting direction is the following. When we find that, for a specific game, there is no mixed strategy corresponding to the optimal marginal probabilities, can we incrementally add variables or constraints to the compact LP formulation to rule out the current marginal solution, until we find a solution that does correspond to a mixed strategy—but without blowing up to the fully general LP formulation?

Acknowledgements

The authors would like to thank Dmytro Korzhyk and Ronald Parr for helpful discussions and Albert Xin Jiang for pointing out the Budish et al. paper to us. We thank ARO and NSF for support under grants W911NF-12-1-0550, W911NF-11-1-0332, IIS-0953756, and CCF-1101659.

References

- Birkhoff, G. 1946. Tres observaciones sobre el algebra lineal. *Univ. Nac. Tucumn Rev, Ser. A, no. 5* 147–151.
- Budish, E.; Koo Che, Y.; Kojima, F.; and Milgrom, P. 2013. Designing random allocation mechanisms: Theory and applications. *American Economic Review*, 103(2): 585–623.
- Conitzer, V., and Sandholm, T. 2006. Computing the optimal strategy to commit to. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, 82–90.
- Jain, M.; Pita, J.; Tambe, M.; Ordóñez, F.; Paruchuri, P.; and Kraus, S. 2008. Bayesian Stackelberg games and their application for security at Los Angeles International Airport. *SIGecom Exch.* 7(2):1–3.
- Kiekintveld, C.; Jain, M.; Tsai, J.; Pita, J.; Ordóñez, F.; and Tambe, M. 2009. Computing optimal randomized resource allocations for massive security games. In *Proceedings of the Eighth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 689–696.
- Korzhyk, D.; Conitzer, V.; and Parr, R. 2010. Complexity of computing optimal Stackelberg strategies in security resource allocation games. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 805–810.
- Lin, G.; Cai, Z.; and Lin, D. 2006. Vertex covering by paths on trees with its applications in machine translation. *Inf. Process. Lett.* 97(2):73–81.
- Paruchuri, P.; Pearce, J. P.; Marecki, J.; Tambe, M.; Ordóñez, F.; and Kraus, S. 2008. Playing games for security: An efficient exact algorithm for solving Bayesian Stackelberg games. In *Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 895–902.
- Pita, J.; Jain, M.; Ordóñez, F.; Portway, C.; Tambe, M.; and Western, C. 2009. Using game theory for Los Angeles airport security. *AI Magazine* 30(1):43–57.
- Shieh, E.; An, B.; Yang, R.; Tambe, M.; Baldwin, C.; DiRenzo, J.; Maule, B.; and Meyer, G. 2012. PROTECT: A deployed game theoretic system to protect the ports of the United States. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 847–854.
- Tsai, J.; Rathi, S.; Kiekintveld, C.; Ordóñez, F.; and Tambe, M. 2009. IRIS - a tool for strategic security allocation in transportation networks. In *The Eighth International Conference on Autonomous Agents and Multiagent Systems - Industry Track*, 37–44.
- von Stengel, B., and Zamir, S. 2010. Leadership games with convex strategy sets. *Games and Economic Behavior* 69:446–457.
- Yin, Z.; Jiang, A. X.; Johnson, M.; Kiekintveld, C.; Leyton-Brown, K.; Sandholm, T.; Tambe, M.; and Sullivan, J. 2012. Trusts: Scheduling randomized patrols for fare inspection in transit systems. In *Proceedings of Innovative Applications of Artificial Intelligence (IAAI)*.