

Approximation Algorithm for Security Games with Costly Resources

Sayan Bhattacharya, Vincent Conitzer, and Kamesh Munagala

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

Abstract. In recent years, algorithms for computing game-theoretic solutions have been developed for real-world security domains. These games are between a defender, who must allocate her resources to defend potential targets, and an attacker, who chooses a target to attack. Existing work has assumed the set of defender’s resources to be fixed. This assumption precludes the effective use of approximation algorithms, since a slight change in the defender’s allocation strategy can result in a massive change in her utility. In contrast, we consider a model where resources are obtained at a cost, initiating the study of the following optimization problem: Minimize the total cost of the purchased resources, given that every target has to be defended with at least a certain probability. We give an efficient logarithmic approximation algorithm for this problem.

1 Introduction

Taking a game as input and computing a solution of it is one of the core problems of algorithmic game theory. To be more precise, it is a collection of problems, one for each combination of a representation scheme and a solution concept. Perhaps the best-known example is the problem of computing a Nash equilibrium of a normal-form game, which is now known to be PPAD-complete for two-player games [2, 3] and FIXP-complete for games with three or more players [5].

In contrast, this paper deals with an alternative solution concept, corresponding to a “Stackelberg” model. In this model, there are two players, a “leader” and a “follower”. The leader first commits to a mixed strategy, and the follower responds only after observing the commitment. Recently, algorithms for computing optimal leader strategies have been developed for various real-world security domains, including US airports [11, 12], the US Federal Air Marshals [13], and the US Coast Guard [1].

Motivated by such applications, Kiekintveld *et al.* [8] proposed a general model of security games. The players are a *defender* and an *attacker*. There is a set of *targets* that the attacker may want to attack (in the Federal Air Marshal example, these would be individual flights), and the defender has a set of *resources* (Federal Air Marshals). The defender can assign each resource to a *schedule*, which consists of a subset of the targets (for example, a tour of multiple flights that a single Federal Air Marshal can take). In general, not every resource

can be assigned to every schedule. If a resource is assigned to a schedule, then it defends all the targets contained in that schedule.

A pure (resp., mixed) strategy for the defender is a deterministic (resp., randomized) assignment of the resources to schedules. The typical motivation given for the Stackelberg model, where the defender (the leader) commits to a mixed strategy and the attacker (the follower) subsequently best-responds, is as follows. Every day, the defender draws an assignment from her distribution. Over the course of time, the attacker observes the realized assignments and eventually learns the probabilities with which each target is defended on any given day. Then, the attacker decides to attack a single target that maximizes her expected utility. The utilities of both players usually depend on: (a) the target that is attacked, and (b) whether or not that target is defended on the day of the attack. In the existing literature on this topic, the typical problem is to find the optimal mixed strategy for the defender to commit to, that is, the one that will maximize the defender’s expected utility.

In this paper, we explore the setting where the resources can be purchased by the defender at a cost. To be more specific, we consider the following problem. Some resources are to be (deterministically) purchased in advance. Next, they are to be randomly assigned to schedules, ensuring that each target is defended with at least a certain probability. The objective is to minimize the total cost of the purchased resources.

For example, suppose that the Federal Air Marshal Service can hire a set of Marshals on a contractual basis for one year, thereby incurring a certain cost. During each day of the following year, the hired Marshals are (randomly) assigned to the schedules, according to the mixed strategy chosen. It is important to emphasize that we require a set of resources to be purchased *deterministically*, after which these resources can be randomly assigned to schedules. It is impractical to randomize the resource supply itself each day, since, for example, security personnel cannot be hired as day laborers.

Existing literature [8, 9] on security games assumes a fixed set of resources. In contrast, our model assumes that the resources are available in unlimited supply, but they can be purchased at a cost. There is a connection between these two settings: Given a fixed set of resources, there is an efficient algorithm to compute the optimal commitment strategy for a defender *if and only if* we can decide in polynomial time whether each target can be defended with a certain probability. The proof of this claim appears in Appendix A.1.

Our Results and Techniques. In Section 2, we formally state our problem and show that it is a generalization of Set Cover, and hence the problem is unlikely to admit an approximation ratio better than $O(\log |\mathcal{T}|)$, where $|\mathcal{T}|$ is the number of targets [6]. Section 3 gives an $O(\log |\mathcal{T}| + \log |\mathcal{S}|)$ approximation algorithm for this problem, where $|\mathcal{S}|$ is the number of schedules. The algorithm partitions the set of targets into two subsets, depending on whether their required coverage probabilities are *small* or *big*, and separately deals with these two subsets.

In Section 3.1, we show that the subproblem of defending the targets with big probabilities admits an LP relaxation (LP (3)), and it can be converted to the

LP relaxation for Set Cover (LP (6)) while losing at most a constant factor in the objective value. Thus, a greedy algorithm (Figure 2) gives a good approximation for this case. Unfortunately, this analysis cannot be extended to defend the targets with small probabilities, because the gap between the optimal objective value of LP (3) and the cost of the optimal solution can become arbitrarily large.

Section 3.2 considers the remaining subproblem—to defend the targets with small probabilities. For this, we present an exponentially sized covering *integer* program (IP (9)) that bounds the cost (Lemma 3) of the optimal solution. Interestingly, the LP relaxation of IP (9) can again be far away from the optimal solution, and thus, it is necessary to impose the integrality constraints.

IP (9) has the additional property that all of its constraint data are integral. Under such circumstances, a simple greedy heuristic (see Dobson [4]) gives a logarithmic approximation to the integer optimum (Lemma 4). The idea is to treat every variable as a set that *covers* part of the constraints, and during each iteration, increment the variable that gives maximum coverage per unit cost. However, IP (9) has exponentially many variables, and hence, we cannot directly apply Dobson’s algorithm to solve it in polynomial time. Instead, we show that the subroutine of selecting the variable with maximum coverage per unit cost is exactly equivalent to maximizing a submodular function subject to a budget constraint. Hence, we can use another greedy algorithm [10] to implement this subroutine upto a constant factor approximation. Accordingly, we solve IP (9), and its outcome determines the resources we purchase and the way in which we randomly assign them to the schedules (Theorem 2).

Remark. All the missing proofs appear in the Appendix.

2 Notations and Preliminaries

There are a set of targets \mathcal{T} , a set of schedules \mathcal{S} , and a set of resource-types Θ . There is an *unlimited supply* of resources of each type. Any resource of type $\theta \in \Theta$ has cost $c(\theta)$. In addition, there is a subset $S(\theta) \subseteq \mathcal{S}$ such that any resource of type θ can be *assigned* to at most one schedule in $S(\theta)$. The type of a resource r is denoted by $\theta_r \in \Theta$. Whenever some resource is assigned to schedule $s \in \mathcal{S}$, it *defends* all targets in the subset $T(s) \subseteq \mathcal{T}$. Furthermore, each target $t \in \mathcal{T}$ has a *threshold requirement* $0 \leq q_t \leq 1$. In the SECURITY GAME problem, we want to (deterministically) purchase some resources, and randomly assign them to schedules so that every target $t \in \mathcal{T}$ is defended with probability at least q_t . We want to minimize the sum of the costs of the purchased resources.¹

Consider an example (see Figure 1). There are 4 targets, 3 schedules, and 2 different types of resources. Target t_1 (resp. t_4) needs to be defended with

¹ Note that the unlimited availability of resources guarantees the existence of a feasible solution satisfying all the threshold requirements: simply purchase a sufficient number of resources of each type. This assertion holds provided each target can be defended by some schedule, and for each schedule, there is some resource that can be assigned to it. We will make these assumptions without any loss of generality.

probability $2/3$ (resp. $1/3$), whereas each of the targets in $\{t_2, t_3\}$ has a threshold requirement of 1. Any resource of type θ_1 costs 2, and any resource of type θ_2 costs 3. A resource of type θ_1 (resp. θ_2) can be assigned to at most one schedule in the set $\{s_1, s_2\}$ (resp. $\{s_2, s_3\}$). Finally, whenever it has some resource assigned to it, schedule s_1 defends target t_1 ; schedule s_2 defends both the targets t_2, t_3 ; and schedule s_3 defends both the targets t_3, t_4 . In terms of notations, we have

$$\begin{aligned} \mathcal{T} &= \{t_1, t_2, t_3, t_4\}, \quad \mathcal{S} = \{s_1, s_2, s_3\}, \quad \Theta = \{\theta_1, \theta_2\} \\ q_{t_1} &= 2/3, \quad q_{t_2} = q_{t_3} = 1, \quad q_{t_4} = 1/3 \\ c(\theta_1) &= 2, \quad c(\theta_2) = 3 \\ S(\theta_1) &= \{s_1, s_2\}, \quad S(\theta_2) = \{s_2, s_3\} \\ T(s_1) &= \{t_1\}, \quad T(s_2) = \{t_2, t_3\}, \quad T(s_3) = \{t_3, t_4\} \end{aligned}$$




Fig. 1. An instance of the SECURITY-GAME problem.

The optimal solution will purchase one resource r_1 of type θ_1 and one resource r_2 of type θ_2 so that $\theta_{r_1} = \theta_1$, $\theta_{r_2} = \theta_2$; thereby incurring a total cost of $2 + 3 = 5$. Next, with probability $2/3$, it will *simultaneously* assign resource r_1 to schedule s_1 and resource r_2 to schedule s_2 ; and with the remaining probability $1 - 2/3 = 1/3$, it will *simultaneously* assign resource r_1 to schedule s_2 and resource r_2 to schedule s_3 . It is important to note how the optimal solution correlates the random assignments of the resources to schedules.

The SECURITY-GAME problem is a generalization of SET-COVER. Consider an instance of the SECURITY-GAME problem where the threshold requirements of all the targets are equal to 1, and there is only one resource-type. In this

case, the task of finding the optimal solution is equivalent to finding a minimum cardinality subset of schedules to defend all the targets, which is exactly the SET-COVER problem. As a consequence, the SECURITY-GAME problem is NP-hard and unless NP has slightly superpolynomial time algorithms, we cannot even approximate it to a factor better than $O(\log |\mathcal{T}|)$ [6]. In the next section, we give an $O(\log |\mathcal{T}| + \log |\mathcal{S}|)$ approximation algorithm.

3 Approximation Algorithm

First, we partition the set of targets into two groups depending on their threshold requirements. Define

$$\mathcal{T}_{big} = \{t \in \mathcal{T} : 1/e < q_t \leq 1\} \quad (1)$$

$$\mathcal{T}_{small} = \{t \in \mathcal{T} : q_t \leq 1/e\} \quad (2)$$

We deal with the two subsets separately. In Section 3.1, we consider the subproblem where we need to defend *only* the subset \mathcal{T}_{big} of targets, and give an $O(\log |\mathcal{T}|)$ approximation algorithm for this task. On the other hand, Section 3.2 deals with the subproblem where we have to defend *only* the targets $t \in \mathcal{T}_{small}$. For this task, we present an $O(\log |\mathcal{T}| + \log |\mathcal{S}|)$ approximation algorithm. Finally, we take the union of the two solutions, and this results in an $O(\log |\mathcal{T}| + \log |\mathcal{S}|)$ approximation for the SECURITY-GAME problem.

3.1 Targets with Big Threshold Requirements

Let OPT_{big} denote the minimum-cost solution that only defends the targets in the subset $\mathcal{T}_{big} \subseteq \mathcal{T}$ according to their threshold requirements, and ignores the remaining targets in $\mathcal{T}_{small} = \mathcal{T} \setminus \mathcal{T}_{big}$. We now derive an LP-relaxation.

$$\min \sum_{\theta \in \Theta} c(\theta) \sum_{s \in S(\theta)} w(\theta, s) \quad (3)$$

$$\text{s.t.} \quad \sum_{\theta \in \Theta} \sum_{s \in S(\theta): t \in T(s)} w(\theta, s) \geq q_t, \quad \forall t \in \mathcal{T}_{big} \quad (4)$$

$$w(\theta, s) \geq 0, \quad \forall \theta \in \Theta, s \in S(\theta) \quad (5)$$

Let \mathcal{R}^* be the set of resources purchased by OPT_{big} . For all $r \in \mathcal{R}^*, s \in S(\theta_r)$, let y_{rs}^* be the probability that OPT_{big} assigns resource r to schedule s . Recall that the type of a resource r is denoted by θ_r . Define

$$w^*(\theta, s) = \sum_{r \in \mathcal{R}: \theta_r = \theta} y_{rs}^*$$

It is easy to verify that the $w^*(\theta, s)$ values are a feasible solution to LP (3). Constraint (4) holds since OPT_{big} defends every $t \in \mathcal{T}_{big}$ with probability at least q_t and by the union bound, the left hand side of the constraint is an

overestimate of the probability with which target t is defended. The relaxation assumes that the resources can be purchased partially and hence, the objective value is at most the total cost incurred by OPT_{big} . This leads us to Lemma 1.

Lemma 1. *LP (3) gives a lower bound on the total cost incurred by OPT_{big} .*

Consider the following linear program (6). It replaces the right hand of Constraint (4) by 1. Recall that all targets in the subset $\mathcal{T}_{big} \subseteq \mathcal{T}$ have $q_t \geq 1/e$. Therefore, if we solve LP (3) optimally and multiply every $w(\theta, s)$ by e , then we get a feasible solution to LP (6). However, the objective value also increases by a factor of e . Combining this observation with Lemma 1, we obtain Fact 1.

$$\min \sum_{\theta \in \Theta} c(\theta) \sum_{s \in S(\theta)} w(\theta, s) \quad (6)$$

$$\text{s.t.} \quad \sum_{\theta \in \Theta} \sum_{s \in S(\theta): t \in T(s)} w(\theta, s) \geq 1, \quad \forall t \in \mathcal{T}_{big} \quad (7)$$

$$w(\theta, s) \geq 0, \quad \forall \theta \in \Theta, s \in S(\theta) \quad (8)$$

Fact 1 *The optimal objective value of LP (6) is at most $O(1)$ (specifically, e) times the cost incurred by OPT_{big} .*

We note that LP (6) is an LP-relaxation for the SET-COVER problem, where the targets $t \in \mathcal{T}_{big}$ behave like elements that have to be covered, and each pair (θ, s) acts like a set $T(s) \cap \mathcal{T}_{big}$ having a cost $c(\theta)$. Hence the greedy algorithm described in Figure 2 gives a $O(\log |\mathcal{T}|)$ approximation [7] to LP (6). Hence, Fact 1 implies the following Theorem 1.

```

Initialize  $D \leftarrow \emptyset, F \leftarrow \emptyset.$ 
While  $D \neq \mathcal{T}_{big}$  do
  Find an ordered pair  $(\theta', s') \in \arg \max_{\theta \in \Theta, s \in S(\theta)} |(T(s) \cap \mathcal{T}_{big}) \setminus D|/c(\theta)$ 
   $F \leftarrow F \cup \{(\theta', s')\}$ 
   $D \leftarrow D \cup (T(s') \cap \mathcal{T}_{big})$ 
For all  $(\theta, s) \in F$  do
  Buy a resource of type  $\theta$ , and deterministically assign it to schedule  $s.$ 

```

Fig. 2. Greedy algorithm for targets with big threshold requirements. It defends every target $t \in \mathcal{T}_{big}$ with probability 1.

Theorem 1. *The greedy algorithm described in Figure (2) gives an $O(\log |\mathcal{T}|)$ approximation to OPT_{big} .*

3.2 Targets with Small Threshold Requirements

Let SG_{small} be the problem where we must defend each target $t \in \mathcal{T}_{small}$ with probability at least q_t , but we are free to ignore the remaining targets in \mathcal{T}_{big} . A *solution* to the SG_{small} problem purchases some resources and randomly assigns them to schedules so that each target $t \in \mathcal{T}_{small}$ is defended with probability q_t .

Definition 1. *Given any solution to the SG_{small} problem, every purchased resource r can be associated with an assignment vector. It is a vector with $|\mathcal{S}|$ components, where the value of component s equals y_{rs} if $s \in S(\theta_r)$ and is zero otherwise. Here, y_{rs} is the probability that resource r is assigned to schedule s .*

Lemma 2 shows that we can restrict our attention to a subset of feasible solutions to the SG_{small} problem. Specifically, we focus on those solutions where the values of all components of the assignment vectors come from a discrete set.

Lemma 2. *Let OPT be the minimum-cost solution to the SG_{small} problem. There exists a solution DISCRETE-OPT to the SG_{small} problem such that:*

1. *The cost incurred by DISCRETE-OPT is at most 2 times the cost of OPT.*
2. *For every resource purchased by DISCRETE-OPT, all the components of the corresponding assignment vector are integral multiples of $1/|\mathcal{S}|^2$.*

Fix any target $t \in \mathcal{T}_{small}$ with $0 < q_t < 1/|\mathcal{S}|^2$. The solution DISCRETE-OPT defends this target according to its threshold requirement. Thus, DISCRETE-OPT purchases some resource of type θ , and assigns it with *non-zero* probability to some schedule $s \in S(\theta)$ such that $t \in T(s)$. However, the probability of assigning the resource to schedule s is an integral multiple of $1/|\mathcal{S}|^2$ (Lemma 2), and hence the target t is defended with probability at least $1/|\mathcal{S}|^2$.

Corollary 1. *The solution DISCRETE-OPT defends each target $t \in \mathcal{T}_{small}$ with probability at least $\min(q_t, 1/|\mathcal{S}|^2)$.*

Let \mathcal{P} denote the set of all assignment vectors that have been discretized according to Lemma 2. More formally, the set \mathcal{P} consists of all possible $|\mathcal{S}|$ -tuples where the value of each component is an integral multiple of $1/|\mathcal{S}|^2$, and the sum of the values of all the components is at most 1. For all $\mathbf{p} \in \mathcal{P}$, $s \in \mathcal{S}$, let $\mathbf{p}(s)$ denote the component of the assignment vector \mathbf{p} corresponding to schedule s .

Suppose that a resource r cannot be assigned to some schedule s , that is, $s \notin S(\theta_r)$, and a valid solution assigns the resource r to different schedules with probabilities that are given by the components of the vector $\mathbf{p} \in \mathcal{P}$. In this case, we must have $\mathbf{p}(s) = 0$. Definition 2 formalizes this concept.

Definition 2. *An assignment vector $\mathbf{p} \in \mathcal{P}$ is feasible for a resource-type $\theta \in \Theta$ if $\mathbf{p}(s) = 0$ for all schedules $s \in \mathcal{S} \setminus S(\theta)$. Define \mathcal{P}_θ to be the set of all feasible assignment vectors for resource type $\theta \in \Theta$.*

Now we present an *Integer Program* to bound the cost of DISCRETE-OPT.

$$\min \sum_{\theta \in \Theta, \mathbf{p} \in \mathcal{P}_\theta} c(\theta)x(\theta, \mathbf{p}) \quad (9)$$

$$\text{s.t.} \quad \sum_{\theta \in \Theta, \mathbf{p} \in \mathcal{P}_\theta} \sum_{s: t \in T(s)} \eta(\mathbf{p}, s) x(\theta, \mathbf{p}) \geq \lambda_t, \quad \forall t \in \mathcal{T}_{small} \quad (10)$$

$$x(\theta, \mathbf{p}) \in \mathbb{N}, \quad \forall \theta \in \Theta, \mathbf{p} \in \mathcal{P}_\theta \quad (11)$$

IP (9) introduces some new notation. In Constraint (11), the set of all non-negative integers is denoted by \mathbb{N} . In Constraint (10), we have

$$\lambda_t = \lceil e \times q_t \times |\mathcal{S}|^2 \rceil \quad \text{for all } t \in \mathcal{T}_{small} \quad (12)$$

$$\eta(\mathbf{p}, s) = \mathbf{p}(s) \times |\mathcal{S}|^2 \quad \text{for all } \mathbf{p} \in \mathcal{P}, s \in \mathcal{S} \quad (13)$$

By definition, each $\mathbf{p}(s) \in [0, 1]$ is an integral multiple of $1/|\mathcal{S}|^2$, and $q_t \in [0, 1/e]$ for all $t \in \mathcal{T}_{small}$. These observations lead to the the following fact.

Fact 2 *IP (9) is a covering integer program where all the coefficients in the constraint data, that is, all the values of $\eta(\mathbf{p}, s)$ and λ_t , are integers lying between 0 and $|\mathcal{S}|^2$.*

Lemma 3. *The optimal objective value of the Integer Program (9) is at most 4 times the cost incurred by DISCRETE-OPT.*

We now describe some intuitions behind IP (9). The variable $x(\theta, \mathbf{p})$ denotes the number of purchased resources that satisfy *both* of the following conditions.

1. The resource is of type $\theta \in \Theta$, and
2. For all $s \in \mathcal{S}$, the probability that the resource is assigned to schedule s is given by $\mathbf{p}(s)$.

Each resource of type θ costs $c(\theta)$. Therefore, summing over all possible resource types and feasible assignment vectors, we see that the total cost is given by the objective function. We now proceed towards verifying Constraint (10). Applying union-bound, we can show that the left hand side of Constraint (10) is at least $|\mathcal{S}|^2$ times the probability of defending target t . Glossing over some of the details, the constraint holds since, the right hand side is roughly equal to $|\mathcal{S}|^2$ times the probability of defending target t .²

If a covering *integer* program has a constraint matrix with integral entries, then a simple greedy heuristic returns a logarithmic approximation to the *integral* optimum (Dobson [4]). Fact 2 tells us that we can apply Dobson’s heuristic to IP (9). A simple intuition behind the algorithm (Figure 3) comes from an alternate way of viewing the problem: Each target t has to be *covered* by a

² To be more precise, the RHS is equal to $\lceil e \times q_t \times |\mathcal{S}|^2 \rceil$. While converting the IP solution to a feasible (random) assignment of resources to schedules, the final algorithm (Figure 4) loses a factor of e in the probability of defending any target $t \in \mathcal{T}_{small}$.

1. Let $\boldsymbol{\delta}$ denote a vector with $|\mathcal{T}_{small}|$ components, where $\boldsymbol{\delta}(t)$ gives the value of the component corresponding to target $t \in \mathcal{T}_{small}$.
2. FOR ALL $t \in \mathcal{T}_{small}$: Initialize $\boldsymbol{\delta}(t) \leftarrow \lambda_t$.
3. FOR ALL $\theta \in \Theta$, $\mathbf{p} \in \mathcal{P}_\theta$: Initialize $X(\theta, \mathbf{p}) \leftarrow 0$.
4. WHILE $\boldsymbol{\delta} \neq 0$ DO
 5. FOR ALL $\mathbf{p} \in \mathcal{P}$, and $t \in \mathcal{T}_{small}$

$$\Delta_{Cov}(\mathbf{p}, \boldsymbol{\delta}, t) \leftarrow \min \left(\boldsymbol{\delta}(t), \sum_{s: t \in T(s)} \mathbf{p}(s) \times |\mathcal{S}|^2 \right).$$
 6. FOR ALL $\mathbf{p} \in \mathcal{P}$: $\Delta_{Cov}(\mathbf{p}, \boldsymbol{\delta}) \leftarrow \sum_{t \in \mathcal{T}_{small}} \Delta_{Cov}(\mathbf{p}, \boldsymbol{\delta}, t)$.
 7. Find some $(\theta, \mathbf{p}) \in \arg \max_{\theta \in \Theta, \mathbf{p} \in \mathcal{P}_\theta} \{ \Delta_{Cov}(\mathbf{p}, \boldsymbol{\delta}) / c(\theta) \}$.
 8. $X(\theta, \mathbf{p}) \leftarrow X(\theta, \mathbf{p}) + 1$.
 9. FOR ALL $t \in \mathcal{T}_{small}$: $\boldsymbol{\delta}(t) \leftarrow \boldsymbol{\delta}(t) - \Delta_{Cov}(\mathbf{p}, \boldsymbol{\delta}, t)$.
10. Return the $X(\theta, \mathbf{p})$ values for all $\theta \in \Theta$, $\mathbf{p} \in \mathcal{P}_\theta$.

Fig. 3. Dobson’s Algorithm applied to LP (9)

threshold amount λ_t . The total *coverage* required is $\sum_{t \in \mathcal{T}_{small}} \lambda_t$. This coverage can be achieved by incrementing the *columns* $\{(\theta, \mathbf{p})\}$, where each column (θ, \mathbf{p}) corresponds to the variable $x(\theta, \mathbf{p})$. We want to increment the columns so that the required coverage is attained at minimum cost.

At the beginning of a typical iteration of the WHILE loop (Steps 4–9), the value of $\boldsymbol{\delta}(t)$ equals the *remaining coverage* required for target t before we can attain its threshold λ_t . If we increment a column (θ, \mathbf{p}) by 1, then the cost of our solution will increase by $c(\theta)$, and at the same time, the coverage of target t will increase (Step 5) by the amount $\Delta_{Cov}(\mathbf{p}, \boldsymbol{\delta}, t)$. Hence, the increase in total coverage of all the targets (Step 6) will be equal to $\Delta_{Cov}(\mathbf{p}, \boldsymbol{\delta})$. Let us term this quantity $\Delta_{Cov}(\mathbf{p}, \boldsymbol{\delta})$ as *marginal coverage*. The algorithm myopically selects the column that has the maximum marginal coverage to cost ratio (Step 7), and increments that column by 1 (Step 8). The remaining coverage required for all the targets are adjusted accordingly (Step 9). The WHILE loop terminates (Step 4) when $\boldsymbol{\delta} = 0$, that is, when all the targets in \mathcal{T}_{small} have been covered up to their corresponding thresholds.

Note that IP (9) contains exponentially many variables $x(\theta, \mathbf{p})$. This follows from the fact that the set \mathcal{P} of possible assignment vectors is exponential in size. Hence, we have to prove that Dobson’s algorithm can be implemented in polynomial time. We also need to establish a bound on the approximation ratio.

Lemma 4. *Dobson’s algorithm (Figure 3) can be used to solve the Integer Program (9). Although IP (9) contains exponentially many variables, an approximate version of Dobson’s algorithm can be implemented in polynomial time. It returns a feasible solution to IP (9), where each variable $x(\theta, \mathbf{p})$ is assigned a nonnegative integral value $X(\theta, \mathbf{p})$. The solution satisfies two properties.*

1. *The objective value of the solution is at most $O(\log |\mathcal{T}| + \log |\mathcal{S}|)$ times the optimal objective value of IP (9).*
2. *The number of variables taking nonzero values are polynomially bounded.*

Proof (Sketch). The approximation ratio of Dobson’s algorithm [4] grows logarithmically with the maximum column sum of the coefficient matrix. Recall that (Fact 2) each $\eta(\mathbf{p}, s)$ is an integer lying between 0 and $|\mathcal{S}|^2$, and the number of constraints in IP (9) is at most $|\mathcal{T}|$. Thus, in case of IP (9), the maximum column sum is upper bounded by $|\mathcal{S}|^2 \times |\mathcal{T}|$. Hence the approximation ratio is given by $O(\log(|\mathcal{S}|^2 |\mathcal{T}|)) = O(\log |\mathcal{S}| + \log |\mathcal{T}|)$. Next, we will show that an approximate version of Dobson’s algorithm can be implemented in polynomial time, and asymptotically, it gives the same approximation ratio of $O(\log |\mathcal{S}| + \log |\mathcal{T}|)$.

It is sufficient to discuss the implementations of Step 3, the WHILE loop (Steps 4–9) and Step 10. We implement Step 3 by implicitly assuming that all the $X(\theta, \mathbf{p})$ values have been initialized to zero. During the course of the algorithm, we keep track of only those $X(\theta, \mathbf{p})$ values that have been incremented at least once. Since each λ_t is an integer lying between 0 and $|\mathcal{S}|^2$ (Fact 2), the total coverage required of all targets is at most $|\mathcal{T}| \times |\mathcal{S}|^2$. Every iteration of the WHILE loop (Steps 4–9) contributes at least 1 towards this total coverage, and it increments exactly one $X(\theta, \mathbf{p})$. Therefore, at the termination of the WHILE loop, the number of nonzero $X(\theta, \mathbf{p})$ values is upper bounded by the polynomial $|\mathcal{T}| \times |\mathcal{S}|^2$. In Step 10, the algorithm returns only these nonzero $X(\theta, \mathbf{p})$ values, and all other variables are implicitly set to zero.

Regarding the WHILE loop (Steps 4–9), note that the marginal coverage $\Delta_{Cov}(\mathbf{p}, \boldsymbol{\delta})$, as a function of the assignment vector \mathbf{p} , is monotone and submodular. To be more precise, fix some resource type θ . Next, take any two assignment vectors $\mathbf{p}, \mathbf{p}' \in \mathcal{P}_\theta$ such \mathbf{p} is dominated by \mathbf{p}' , that is, $\mathbf{p}(s) \leq \mathbf{p}'(s)$ for all $s \in \mathcal{S}$. Furthermore, suppose that $\sum_{s \in \mathcal{S}} \mathbf{p}'(s) < 1$. Fix some schedule $s^* \in \mathcal{S}(\theta)$ and consider two new assignment vectors $\mathbf{p}_1, \mathbf{p}'_1 \in \mathcal{P}_\theta$ with the following properties. For all $s \in \mathcal{S} \setminus \{s^*\}$, we have $\mathbf{p}_1(s) = \mathbf{p}(s)$, and $\mathbf{p}'_1(s) = \mathbf{p}'(s)$. We also have $\mathbf{p}_1(s^*) = \mathbf{p}(s^*) + 1/|\mathcal{S}|^2$, and $\mathbf{p}'_1(s^*) = \mathbf{p}'(s^*) + 1/|\mathcal{S}|^2$. Now, submodularity of marginal coverage means that the next inequality will always be satisfied.

$$\Delta_{Cov}(\mathbf{p}'_1, \boldsymbol{\delta}) - \Delta_{Cov}(\mathbf{p}', \boldsymbol{\delta}) \leq \Delta_{Cov}(\mathbf{p}_1, \boldsymbol{\delta}) - \Delta_{Cov}(\mathbf{p}, \boldsymbol{\delta})$$

We can exploit this submodularity condition as follows. While implementing Step 7, suppose we have correctly guessed the resource type θ that maximizes the marginal coverage to cost ratio.³ All we need to do is to find an assignment vector $\mathbf{p} \in \mathcal{P}_\theta$ with maximum marginal coverage, subject to the constraints that each component of the vector \mathbf{p} is an integral multiple of $1/|\mathcal{S}|^2$, and the sum of all the components is at most *one* (Lemma 2). This is equivalent to maximizing a monotone submodular function subject to a budget constraint [10]. A simple greedy algorithm is known to give a $(1 - 1/e)$ approximation for this problem.

To summarize, in polynomial time we can obtain a column (θ, \mathbf{p}) that gives a constant approximation to the optimal ratio of marginal coverage to cost. Going through Dobson’s proof [4], it is easy to verify the following statement. Even if we implement Step 7 in this approximate fashion, the algorithm will have the same asymptotic approximation guarantee. This concludes the proof. \square

³ Clearly, in $O(|\theta|)$ time we can iterate over all possible resource types.

We are now ready to present our algorithm for the SG_{small} problem (Figure 4). First, we solve IP (9) according to Lemma 4. Let B^* denote the set of columns of IP (9) where the corresponding variable is set to some nonzero value, that is, $B^* = \{(\theta, \mathbf{p}) : X(\theta, \mathbf{p}) \neq 0\}$. For each $(\theta, \mathbf{p}) \in B^*$, we purchase $X(\theta, \mathbf{p})$ resources of type θ and tag them with the assignment vector \mathbf{p} . Finally, each purchased resource is randomly assigned to some schedule according to its assignment vector; and this process occurs *independently* of all other resources.

Solve IP (9) according to Lemma 4.
 Define $\mathcal{P}_\theta^* = \{\mathbf{p} \in \mathcal{P}_\theta : X(\theta, \mathbf{p}) \neq 0\}$ for all $\theta \in \Theta$.
 Let \mathcal{R}_θ be the set of resources of type θ that will be purchased.
 Let $\mathcal{R} = \bigcup_{\theta \in \Theta} \mathcal{R}_\theta$ be the set of *all* resources that will be purchased.
 Let $\mathcal{R}_{\theta, \mathbf{p}} \subseteq \mathcal{R}_\theta$ denote the resources in \mathcal{R}_θ whose assignment probabilities are specified by $\mathbf{p} \in \mathcal{P}_\theta^*$.
 For all resource-types $\theta \in \Theta$
 $|\mathcal{R}_\theta| = \sum_{\mathbf{p} \in \mathcal{P}_\theta^*} X(\theta, \mathbf{p})$.
 For all $\mathbf{p} \in \mathcal{P}_\theta^*$: $|\mathcal{R}_{\theta, \mathbf{p}}| = X(\theta, \mathbf{p})$.
 Randomly assign each resource $r \in \mathcal{R}_{\theta, \mathbf{p}}$ to schedules, according to the assignment probabilities specified by \mathbf{p} , independently of all other resources.

Fig. 4. Approximation Algorithm for the SG_{small} Problem

Theorem 2. *The algorithm described in Figure 4 gives an $O(\log |\mathcal{T}| + \log |\mathcal{S}|)$ approximation to the SG_{small} problem.*

Proof. If we purchase the resources according to Figure 4, then the total cost is equal to the objective value of the IP solution returned by Lemma 4. Now Lemma 2, Lemma 3 and Lemma 4 imply that this cost is at most $O(\log |\mathcal{T}| + \log |\mathcal{S}|)$ times the cost incurred by OPT. It remains to show that the solution defends all the targets in \mathcal{T}_{small} according to their threshold requirements. Fix some target $t \in \mathcal{T}_{small}$ for the rest of this proof. Given any purchased resource $r \in \mathcal{R}$, let \mathbf{p}_r be its assignment vector according to Figure 4. Since the $X(\theta, \mathbf{p})$ values constitute a feasible solution to IP (9), we have that $\sum_{r \in \mathcal{R}} \sum_{s: t \in T(s)} \eta(\mathbf{p}_r, s) \geq \lambda_t$. Recall that $\eta(\mathbf{p}_r, s) = \mathbf{p}_r(s) |\mathcal{S}|^2$ and $\lambda_t = \lceil e \times q_t \times |\mathcal{S}|^2 \rceil \geq eq_t |\mathcal{S}|^2$. For all $r \in \mathcal{R}$, define $\phi_r(t) = \sum_{s: t \in T(s)} \mathbf{p}_r(s)$. Therefore, we get

$$\sum_{r \in \mathcal{R}} \phi_r(t) = \sum_{r \in \mathcal{R}} \sum_{s: t \in T(s)} \mathbf{p}_r(s) \geq eq_t \quad (14)$$

The probability that resource r does *not* defend target t is given by the expression $1 - \sum_{s: t \in T(s)} \mathbf{p}_r(s) = 1 - \phi_r(t)$. Since the event of assigning a resource to some (random) schedule occurs independently of other resources, the probability that *no resource* defends the target t is equal to $\prod_{r \in \mathcal{R}} (1 - \phi_r(t))$. Since $q_t \leq 1/e$,

$$\prod_{r \in \mathcal{R}} (1 - \phi_r(t)) \leq \prod_{r \in \mathcal{R}} \exp(-\phi_r(t)) = \exp\left(-\sum_{r \in \mathcal{R}} \phi_r(t)\right) \leq \exp(-eq_t)$$

Thus, the probability of defending target t is at least $1 - \exp(-eq_t) \geq qt$. \square

Remark. We note that it is possible to devise a polynomial time algorithm that gives an $O(\log |\mathcal{T}|)$ approximation to the SG_{small} problem. We have to consider an exponential sized Linear Program that is similar to IP (9), solve it approximately using an approximate separation oracle for its dual, and then directly employ randomized rounding. However, the running time of such an algorithm might become prohibitive.

4 Conclusion

We investigated the security game problem when there is an unlimited supply of resources that can be purchased at a cost. We designed an algorithm for (deterministically) purchasing some resources at minimum cost, and then randomly assigning them to schedules so that each target is defended with at least a certain probability. The algorithm is efficient and gives a logarithmic approximation.

Since this problem is a generalization of SET-COVER, we cannot get a sub-logarithmic approximation ratio. However, if each target has at most two schedules that are capable of defending it (a generalization of the VERTEX-COVER problem) and resources are homogeneous, then we can get a constant factor approximation algorithm. We leave open the questions of exploring other settings with better approximation guarantees, and investigating the fixed parameter tractability of the problem.

Acknowledgements. The authors thank Dmytro Korzhyk and Ronald Parr for several helpful discussions. This research was supported by NSF under award numbers IIS-0812113, IIS-0953756, CCF-1101659, CCF-0745761, CCF-1008065, a gift from Cisco, by Conitzer’s Alfred P. Sloan Research Fellowship, and by Munagala’s Alfred P. Sloan Research Fellowship.

References

1. B. An, J. Pita, E. Shieh, M. Tambe, C. Kiekintveld, and J. Marecki. GUARDS and PROTECT: next generation applications of security games. *ACM SIGecom Exchanges*, 10(1):31–34, 2011.
2. X. Chen and X. Deng. Settling the complexity of two-player Nash equilibrium. In *FOCS*, pages 261–272, 2006.
3. C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a Nash equilibrium. In *STOC*, pages 71–78, 2006.
4. G. Dobson. Worst-case analysis of greedy heuristics for integer programming with nonnegative data. *Mathematics of Operations Research*, 7(4):515–531, 1982.
5. K. Etessami and M. Yannakakis. On the complexity of Nash equilibria and other fixed points. *SIAM J. Comput.*, 39(6):2531–2597, 2010.
6. U. Feige. A threshold of $\ln n$ for approximating set-cover. *Journal of the ACM*, 45(4):634–652, 1998.

7. D. Hochbaum. *Approximation Algorithms for NP-hard Problems*. PWS Publishing Company, 1997.
8. C. Kiekintveld, M. Jain, J. Tsai, J. Pita, F. Ordóñez, and M. Tambe. Computing optimal randomized resource allocations for massive security games. In *AAMAS*, pages 689–696, 2009.
9. D. Korzhyk, V. Conitzer, and R. Parr. Complexity of computing optimal Stackelberg strategies in security resource allocation games. In *AAAI*, pages 805–810, 2010.
10. G. Nemhauser, L. Wolsey, and M. Fisher. An analysis of approximations for maximizing submodular set functions. *Mathematical Programming*, 14(1):265–294, 1978.
11. J. Pita, M. Jain, F. Ordóñez, C. Portway, M. Tambe, and C. Western. Using game theory for Los Angeles Airport security. *AI Magazine*, 30(1):43–57, 2009.
12. J. Pita, M. Jain, C. Western, C. Portway, M. Tambe, F. Ordóñez, S. Kraus, and P. Parachuri. Deployed ARMOR protection: The application of a game-theoretic model for security at the Los Angeles International Airport. In *AAMAS - Industry and Applications Track*, pages 125–132, 2008.
13. J. Tsai, S. Rath, C. Kiekintveld, F. Ordóñez, and M. Tambe. IRIS - a tool for strategic security allocation in transportation. In *AAMAS - Industry Track*, pages 37–44, 2009.

A Appendix

A.1 Equivalence between the two problems

The Setting. We are given a set of targets \mathcal{T} and a set of schedules \mathcal{S} . In Section 2, we assumed that the resources are available in *unlimited supply*. In contrast, here we assume that the set of resources is *fixed*. Consequently, we no longer have to deal with the notion of “resource type”. Let \mathcal{R} be the fixed set of resources available. Overloading the notation introduced in Section 2, a resource $r \in \mathcal{R}$ can be assigned to at most one schedule in the subset $S(r) \subseteq \mathcal{S}$. Whenever some resource is assigned to schedule $s \in \mathcal{S}$, it defends all targets in the subset $T(s) \subseteq \mathcal{T}$.

The Utility Model. This setting can be viewed as a strategic game between the attacker and the defender. Every pure strategy of the defender corresponds to an assignment of the available set of resources \mathcal{R} to the schedules in \mathcal{S} . In general, there can be exponentially many pure strategies (assignments of resources to schedules) for the defender, and the game may not be succinctly representable in normal form. The pure strategy of an attacker corresponds to the target where the attack is going to take place.

If attack takes place at target $t \in \mathcal{T}$ and it is being defended, then the utility of the attacker (resp. the defender) is given by $U_a^1(t)$ (resp. $U_d^1(t)$). On the other hand, if the attack takes place at target $t \in \mathcal{T}$ and it is *not* defended, then the utility of the attacker (resp. the defender) is given by $U_a^2(t)$ (resp. $U_d^2(t)$). It is assumed that $U_a^1(t) \leq U_a^2(t)$, and $U_d^1(t) \geq U_d^2(t)$ for all targets $t \in \mathcal{T}$.

Problem 1 (SG-Utility) *An instance of the SG-UTILITY problem specifies the utility functions of the two players. First, the defender commits to a mixed*

strategy. The attacker observes this commitment, and then responds by attacking the target that maximizes the attacker's expected utility. The objective is to find the optimal (utility maximizing) mixed strategy for the defender to commit to.

Problem 2 (SG-Decision) *An instance of the SG-DECISION problem specifies a threshold requirement $q_t \in [0, 1]$ for every target $t \in \mathcal{T}$. The objective is to output "YES" (resp. "NO") if it is (resp. is not) possible to randomly assign the resources to the schedules in such a way that each target $t \in \mathcal{T}$ is defended with probability at least q_t .*

Lemma 5. *If we can solve the SG-UTILITY problem in polynomial time, then we can also solve the SG-DECISION problem in polynomial time.*

Proof (Sketch). Suppose that we are given an instance of the SG-DECISION problem with the thresholds requirements $q_t \in [0, 1]$ for every target $t \in \mathcal{T}$. Create a new (dummy) target t^* . Set the utilities of the attacker and the defender in the following way. The defender's utility is 1 whenever target t^* is attacked (regardless of whether it is defended), and 0 whenever any other target is attacked (regardless of whether it is defended). The attacker's utility for attacking target t^* is 1 (regardless of whether it is defended). Finally, the attacker's utility for attacking any target $t \in \mathcal{T}$ is $1 + q_t$ if it is *not* being defended, and is q_t if it is being defended. In terms of previously introduced notations, we have

$$\begin{aligned} U_d^1(t^*) &= U_d^2(t^*) = 1 \\ \forall t \in \mathcal{T}, \quad U_d^1(t) &= U_d^2(t) = 0 \\ U_a^1(t^*) &= U_a^2(t^*) = 1 \\ \forall t \in \mathcal{T}, \quad U_a^1(t) &= q_t, \quad U_a^2(t) = 1 + q_t \end{aligned}$$

The lemma now follows from Fact 3 and Fact 4.

Fact 3 *Apply the algorithm for the SG-UTILITY problem on the set of targets $\mathcal{T} \cup \{t^*\}$. It will return the optimal mixed strategy for the defender to commit to. If the defender's expected utility from this commitment is 1, then the output of the SG-DECISION problem (for the set of targets \mathcal{T}) is "YES".*

Proof. Since the defender's commitment guarantees her an utility of 1, the attack has to take place at target t^* . Fix any target $t \in \mathcal{T}$. The attacker's expected utility for attacking target t has to be *at most* her utility for attacking target t^* . Let γ_t (resp. γ_{t^*}) be the probability that target t (resp. target t^*) is defended under the commitment strategy. Therefore, we have

$$\begin{aligned} \gamma_t \times U_a^1(t) + (1 - \gamma_t) \times U_a^2(t) &\leq \gamma_{t^*} \times U_a^1(t^*) + (1 - \gamma_{t^*}) \times U_a^2(t^*) \\ \implies \gamma_t \times q_t + (1 - \gamma_t) \times (1 + q_t) &\leq \gamma_{t^*} \times 1 + (1 - \gamma_{t^*}) \times 1 \\ \implies q_t &\leq \gamma_t \end{aligned}$$

In other words, the defender's commitment strategy defends each target $t \in \mathcal{T}$ with probability at least q_t . \square

Fact 4 *Apply the algorithm for the SG-UTILITY problem on the set of targets $\mathcal{T} \cup \{t^*\}$. It will return the optimal mixed strategy for the defender to commit to. If the defender’s expected utility from this commitment is strictly less than 1, then the output of the SG-DECISION problem (for the set of targets \mathcal{T}) is “NO”.*

Proof. We will prove the contrapositive of the claim. Suppose that the output of the SG-DECISION problem, for the set of targets \mathcal{T} , is “YES”. Hence, there is a random assignment \mathcal{D} of the resources to schedules that defends every target $t \in \mathcal{T}$ with probability at least q_t . Suppose that we use the same random assignment \mathcal{D} as the defender’s commitment strategy. Following the proof of Fact 3, it is easy to see (assuming the attacker breaks ties in defender’s favor) that the attack will take place at target t^* . Consequently, the defender will get a utility of 1. \square

We now make the following standard assumption [9, ?].

Assumption 1 *For every schedule $s \in \mathcal{S}$ and every subset of targets $F \subseteq T(s)$, there exists some schedule $s' \in \mathcal{S}$ such that $T(s') = F$. In other words, the schedules satisfy the downward closure property.*

Assumption 1 implies that there are exponentially many schedules. However, they can be represented by only considering the *maximal* subsets of targets that constitute valid schedules. The approximation ratio and the running time of our algorithm in Section 3 remain the same under this compact representation.

Lemma 6. *If we can solve the SG-DECISION problem in polynomial time and Assumption 1 holds, then we can also solve the SG-UTILITY problem in polynomial time.*

Proof (Sketch). We are given an instance of the SG-UTILITY problem that specifies the utility functions of the attacker and the defender. Consider the following question.

Question 1 *Take any target $t \in \mathcal{T}$ and any value $\gamma_t \in [0, 1]$. Is there a commitment strategy for the defender that (a) ensures that the attack will take place at target t , and (b) defends target t with probability exactly equal to γ_t ?*

We will show (Fact 5) how to efficiently resolve Question 1 by using the algorithm for the SG-DECISION problem as a black box. Given Fact 5, we can execute the following procedure. For all $t \in \mathcal{T}$, we perform a binary search on the maximum probability with which target t can be defended, ensuring that the attack takes place at this target. This maximum probability uniquely determines the maximum expected utility the defender can obtain, subject to the constraint that the attack takes place at target t . We execute the binary search by resolving Question 1 with a new pair of inputs (t, γ_t) on every occasion. Iterating over all possible targets, we can identify the optimal commitment strategy for the defender, and solve the SG-UTILITY problem.

Fact 5 *If Assumption 1 holds, then we can resolve Question 1 in polynomial time by using the algorithm for the SG-DECISION problem as a black box.*

Proof. Suppose that there is a commitment strategy \mathcal{D} that resolves Question 1 in the affirmative. Such a strategy \mathcal{D} can exist if and only if the attacker's expected utility for attacking target t is *not* less than her expected utility for attacking any other target $t' \in \mathcal{T} \setminus t$. Let $w_{t'}$ be the probability that strategy \mathcal{D} defends target t' . Therefore, we have

$$\gamma_t \times U_a^1(t) + (1 - \gamma_t) \times U_a^2(t) \geq w_{t'} \times U_a^1(t') + (1 - w_{t'}) \times U_a^2(t')$$

Rearranging the terms in the above inequality, we get

$$w_{t'} \geq \frac{U_a^2(t') - \gamma_t \times U_a^1(t) - (1 - \gamma_t) \times U_a^2(t)}{U_a^2(t') - U_a^1(t')} = q_{t'} \text{ (say)}$$

Hence, we call the algorithm for the SG-DECISION problem to verify if there exists a commitment strategy that defends every target $t' \in \mathcal{T} \setminus \{t\}$ with probability at least $q_{t'}$, while defending target t with probability at least γ_t .

Remark. Note that we actually have to defend target t with probability *exactly equal* to γ_t . We point out how to satisfy this constraint. Suppose that the solution returned by the algorithm for the SG-DECISION problem defends target t with a probability strictly greater than γ_t . Using Assumption 1, we can easily reduce the probability of defending target t until it becomes exactly equal to γ_t , without affecting the the probabilities of defending the other targets. \square

A.2 Proof of Lemma 1

Let R^* be the set of resources purchased by OPT_{big} . Note that the set R^* may contain multiple resources of same type. For all resources $r \in R^*$ and schedules $s \in S(\theta_r)$, let $v_{r,s}^*$ denote the probability with which OPT_{big} assigns resource r to schedule s . Next, we define

$$w^*(\theta, s) = \sum_{r \in R^* : \theta_r = \theta} v_{r,s}^* \quad \forall \theta \in \Theta, s \in S(\theta)$$

In other words, if we treat the probabilities $v_{r,s}^*$ as fractional assignments of resources to schedules, then the quantity $w^*(\theta, s)$ gives the total fractional assignment to schedule s due to the resources of type θ that were purchased.

The $w^*(\theta, s)$ values satisfy Constraint (4), since OPT_{big} defends every target $t \in \mathcal{T}_{big}$ with probability at least q_t and by the union bound, the left hand side of Constraint (4) is an overestimate of the probability with which target t is defended. Furthermore, the $w^*(\theta, s)$ values satisfy Constraint (5) since the probabilities $v_{r,s}^*$ are nonnegative.

Finally, we note that OPT_{big} purchases at least $\sum_{s \in S(\theta)} w^*(\theta, s)$ resources of type θ , and this contributes at least the amount $\sum_{s \in S(\theta)} c(\theta) w^*(\theta, s)$ to the total cost. Summing over all resource-types, we see that the optimal objective value of LP (3) lower bounds the cost incurred by OPT_{big} . \square

A.3 Proof of Lemma 2

Let R^* denote the set of resources purchased by OPT. We will construct a feasible solution DISCRETE-OPT satisfying the two properties mentioned in the theorem. For every resource $r \in R^*$, we purchase r and *an additional* resource \tilde{r} of the same type. Thus, the total cost incurred is 2 times the cost of OPT, ensuring that Property 1 is satisfied.

Let \mathcal{D}^* be the distribution according to which OPT randomly assigns the resources in R^* to the schedules. For DISCRETE-OPT, we start with a distribution \mathcal{D} , where the resources in R^* are assigned according to distribution \mathcal{D}^* , and the additional resources $\{\tilde{r} : r \in R^*\}$ are *not* assigned to any schedule at all. Clearly, if DISCRETE-OPT follows distribution \mathcal{D} , then it will defend all the targets in \mathcal{T}_{small} according to their threshold requirements, and hence it will be a feasible solution to the SG_{small} problem. However, it may violate Property 2.

Fix some resource $r \in R^*$. We will show how to transform \mathcal{D} into a new distribution $\mathcal{D}^{(r)}$ for DISCRETE-OPT such that the solution remains feasible, the assignment probabilities of both the resources r and \tilde{r} become integral multiples of $1/|\mathcal{S}|^2$, and furthermore, the assignment probabilities of all other resources remain unchanged. Repeatedly applying this transformation for every resource in R^* , we can satisfy Property 2.

We need to introduce some further notation. Let y_{rs} be the probability that resource $r \in R^*$ is assigned to schedule $s \in S(\theta_r)$ in distribution \mathcal{D} . Define

$$\alpha_{rs} = \frac{|\mathcal{S}|^2 \times y_{rs} - \lfloor |\mathcal{S}|^2 \times y_{rs} \rfloor}{|\mathcal{S}|^2}, \quad A_r = \sum_{s \in S(\theta_r)} \alpha_{rs}, \quad \beta_{rs} = y_{rs} - \alpha_{rs} \quad (15)$$

The transformation from \mathcal{D} to $\mathcal{D}^{(r)}$ is illustrated in Figure 5. Step (2) guarantees that $\{r, \tilde{r}\}$ are the *only* resources whose assignment probabilities are changed. Steps (1), (2), (3), (4) ensure that all the targets in \mathcal{T}_{small} are defended according to their threshold requirements. It is also easy to see that the resource r gets assigned to any schedule s with probability $y_{rs} \times (\beta_{rs}/y_{rs}) = \beta_{rs}$, which is an integral multiple of $1/|\mathcal{S}|^2$. It remains to verify that the assignment probabilities of the additional resource \tilde{r} are multiples of $1/|\mathcal{S}|^2$; this is the purpose of Step (5).

The probability of resource \tilde{r} getting assigned to *some* schedule during Step (4) is given by the expression

$$\sum_{s \in S(\theta_r)} y_{rs} \times \left(1 - \frac{\beta_{rs}}{y_{rs}}\right) = A_r \quad (16)$$

By definition, each $\alpha_{rs} \leq 1/|\mathcal{S}|^2$. Therefore,

$$(1 - A_r) = 1 - \sum_{s \in S(\theta_r)} \alpha_{rs} \geq 1 - 1/|\mathcal{S}| \geq \sum_{s \in S(\theta_r)} (1/|\mathcal{S}|^2 - \alpha_{rs}) \quad (17)$$

The above equations (16), (17) have two important consequences. First, during Step (4), resource \tilde{r} is *not* assigned to any schedule with probability $(1 - A_r)$. Second, it is feasible to design a random process as described in Step (5).

- | |
|--|
| <ol style="list-style-type: none"> 1. According to distribution \mathcal{D}
Draw an assignment of the resources to the schedules. 2. Assign every resource other than $\{r, \tilde{r}\}$ according to outcome of the draw. 3. Suppose the outcome says resource r should be assigned to schedule s. 4. Perform a random coin toss where $\Pr[\text{HEADS} = \beta_{rs}/y_{rs}]$.
IF the outcome of the coin toss is HEADS, THEN
Assign the resource r to schedule s.
ELSE IF the outcome of the coin toss is TAILS, THEN
Assign the additional resource \tilde{r} to schedule s. 5. IF resource \tilde{r} is <i>not</i> assigned to any schedule during Steps (1)–(4), THEN
Randomly assign \tilde{r} to one of the schedules in $S(\theta_r)$; each schedule $s \in S(\theta_r)$ is selected with probability $(1/ \mathcal{S} ^2 - \alpha_{rs})/(1 - A_r)$. |
|--|

Fig. 5. Constructing the distribution $\mathcal{D}^{(r)}$.

Fix some schedule $s \in S(\theta_r)$. Resource \tilde{r} may get assigned to schedule s either in Step (4) or in Step (5), and these two events are mutually exclusive. The former event occurs with probability $y_{rs} \times (1 - \beta_{rs}/y_{rs}) = \alpha_{rs}$, whereas the latter event occurs with probability

$$(1 - A_r) \times \frac{(1/|\mathcal{S}|^2 - \alpha_{rs})}{(1 - A_r)} = 1/|\mathcal{S}|^2 - \alpha_{rs}$$

Adding up these two values, we see that resource \tilde{r} is assigned to every schedule $s \in S(\theta_r)$ with the same probability $1/|\mathcal{S}|^2$. This concludes the proof of the lemma. □

A.4 Proof of Lemma 3

In the solution DISCRETE-OPT, let $x^*(\theta, \mathbf{p})$ denote the number of purchased resources that satisfy *both* of the following conditions.

1. The resource is of type $\theta \in \Theta$, and
2. For all $s \in \mathcal{S}$, the probability that the resource is assigned to schedule s is given by $\mathbf{p}(s)$.

Every $x^*(\theta, \mathbf{p})$ is a non-negative integer, and consequently, these variables satisfy Constraint (11). Recall that each resource of type θ costs $c(\theta)$. Therefore, the total cost of DISCRETE-OPT is given by the expression $\sum_{\theta \in \Theta, \mathbf{p} \in \mathcal{P}_\theta} c(\theta)x^*(\theta, \mathbf{p})$, which corresponds to the objective of IP (9).

We proceed towards verifying Constraint (10). Here, we consider two cases separately, depending on the value of q_t .

Case 1. Fix some target $t \in \mathcal{T}_{small}$ such that $q_t \geq 1/|\mathcal{S}|^2$. In this case,

$$\lambda_t = \lceil e q_t |\mathcal{S}|^2 \rceil \leq 4q_t |\mathcal{S}|^2$$

First, note that the $x^*(\theta, \mathbf{p})$ values satisfy the following constraint.

$$\sum_{\theta \in \Theta} \sum_{\mathbf{p} \in \mathcal{P}_\theta} \sum_{s: t \in T(s)} \mathbf{p}(s) \times x^*(\theta, \mathbf{p}) \geq q_t$$

The above inequality holds since DISCRETE-OPT defends target t with probability at least q_t , and by the union bound, the left hand side is an overestimate of the probability of defending target t . Multiplying both sides by $|\mathcal{S}|^2$, we get

$$\sum_{\theta \in \Theta} \sum_{\mathbf{p} \in \mathcal{P}_\theta} \sum_{s: t \in T(s)} \eta(\mathbf{p}, s) x^*(\theta, \mathbf{p}) \geq q_t |\mathcal{S}|^2$$

Recall that $4q_t |\mathcal{S}|^2 \geq \lambda_t$. Hence, if we set $\hat{x}(\theta, \mathbf{p}) = 4 \times x^*(\theta, \mathbf{p})$, then the $\hat{x}(\theta, \mathbf{p})$ values will satisfy Constraint (10) of the integer program.

Case 2. Fix some target $t \in \mathcal{T}_{small}$ such that $q_t < 1/|\mathcal{S}|^2$. In this case,

$$\lambda_t = \lceil e q_t |\mathcal{S}|^2 \rceil \leq 4$$

By Corollary 1, the solution DISCRETE-OPT defends target t with probability at least $1/|\mathcal{S}|^2$. Mimicking the argument for Case 1, we see that the $x^*(\theta, \mathbf{p})$ values satisfy the following constraint.

$$\sum_{\theta \in \Theta} \sum_{\mathbf{p} \in \mathcal{P}_\theta} \sum_{s: t \in T(s)} \eta(\mathbf{p}, s) x^*(\theta, \mathbf{p}) \geq 1$$

Recall that $\lambda_t \leq 4$. Hence, if we set $\hat{x}(\theta, \mathbf{p}) = 4 \times x^*(\theta, \mathbf{p})$, then the $\hat{x}(\theta, \mathbf{p})$ values will satisfy Constraint (10) of the integer program.

To summarize, in both cases, the $\hat{x}(\theta, \mathbf{p})$ values are integral, and constitute a feasible solution to the Integer Program (9) with an objective that is 4 times the cost of DISCRETE-OPT. This concludes the proof of the lemma. \square