A Polynomial Time Algorithm for Spatio-Temporal Security Games

SOHEIL BEHNEZHAD, University of Maryland

MAHSA DERAKSHAN, University of Maryland

MOHAMMADTAGHI HAJIAGHAYI, University of Maryland

ALEKSANDRS SLIVKINS, Microsoft Research

An ever-important issue is protecting infrastructure and other valuable targets from a range of threats from vandalism to theft to piracy to terrorism. The “defender” can rarely afford the needed resources for a 100% protection. Thus, the key question is, *how to provide the best protection using the limited available resources*.

We study a practically important class of security games that is played out in space and time, with targets and “patrols” moving on a real line. A central open question here is whether the Nash equilibrium (i.e., the minimax strategy of the defender) can be computed in polynomial time. We resolve this question in the affirmative. Our algorithm runs in time polynomial in the input size, and only polylogarithmic in the number of possible patrol locations (M). Further, we provide a *continuous* extension in which patrol locations can take arbitrary real values. Prior work obtained polynomial-time algorithms only under a substantial assumption, e.g., a constant number of rounds. Further, all these algorithms have running times polynomial in M , which can be very large.

1 INTRODUCTION

Protecting infrastructure and other valuable targets from a range of threats from vandalism to theft to piracy to terrorism is an ever-important issue around the world, aggravated recently by increased threats of piracy and terrorism. Providing 100% protection usually requires more money or other resources than the “defender” can commit. Thus, the key question is, *how to provide the best protection using the limited resources that are available*.

A successful recent approach casts this issue in game-theoretic terms, modeling it as a *security game*: a zero-sum game between the *defender* who has some targets to protect, and the *attacker* who strives to inflict damage on these targets. Usually the defender needs to commit to a particular allocation of resources, such as the schedule of patrols, whereas the attacker can strike at will; this corresponds to a classic game-theoretic model called a *Stackelberg game*. The defender can (and should) randomize, e.g. so as to prevent the attacker

Supported in part by NSF CAREER award CCF-1053605, NSF BIGDATA grant IIS-1546108, NSF AF:Medium grant CCF-1161365, DARPA GRAPHS/AFOSR grant FA9550-12-1-0423, and another DARPA SIMPLEX grant. Authors’ emails: {soheil,mahsaa,hajiagha}@cs.umd.edu, slivkins@microsoft.com.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC'17, June 26–30, 2017, Cambridge, Massachusetts, USA. ACM ISBN 978-1-4503-4527-9/17/06...\$15.00.

© 2017 Copyright is held by the owner/author(s). Publication rights licensed to ACM.

DOI: <http://dx.doi.org/10.1145/3033274.3085133>

from exploiting a particular gap in the patrol schedule. The attacker can be strategic and optimize his attack according to his beliefs about the defender's strategy. The literature has mostly adopted a pessimistic view, in which the attack is the exact best response to the defender's actual strategy. Thus, the defender's goal is to use an optimal (minimax) strategy.

This approach has resulted in a flurry of research activity, including several awards and nominations. Further, it has been adopted in a number of real-world deployments, ranging from patrol boats to airport checkpoints to US air marshals to an urban transit system to wildlife protection. (Many of them have been recognized with commendations and awards.) Other potential applications include protecting aid convoys in unstable regions, and protecting ships from piracy.

Most applications of security games are *spatio-temporal* in the sense that the patrols move from one location to another with a limited speed, and only protect targets that are sufficiently close. Then the defender's strategy is a rather complicated object: a pure strategy should specify a trajectory for every patrol, possibly choosing from a very large number of possible locations. Further, in many applications targets have their own trajectories that need to be taken into consideration.

The first-order question in spatio-temporal security games is computing the equilibrium, i.e., the minimax strategy for the defender. More specifically, we focus on exact equilibrium computation in polynomial time. The prior work (e.g., [5, 8, 20]), as well as the present paper, considers a one-dimensional space (that is, patrol and target locations are on the real line), and discretizes it uniformly into the possible patrol locations. The relevant parameters are: the number of patrols (K), the number of targets (A), the number of rounds of scheduling (T), and the number of possible patrol locations (M). The input specifies trajectories of targets and their values; the trajectories may be arbitrary, and the values may change over time. Thus, the input size is $O(T \cdot A + \log(K \cdot M))$. What makes the problem particularly challenging is that the number of pure strategies — tuples of patrol trajectories — is as large as $(M)^{KT}$.

A central open question here is whether the Nash equilibrium (i.e., the minimax strategy of the defender) can be computed in polynomial time. We resolve this question in the affirmative: our main result is an algorithm that computes the exact equilibrium in time polynomial in the input size. In particular, the running time scales only *polylogarithmically* in the number of possible patrol locations. Moreover, we provide a *continuous* extension in which patrol locations can take arbitrary real values, under a mild technical assumption that the target locations are rational. The dependence on the number of patrols is argued away: while a pure strategy of the defender must specify a trajectory for each patrol, we prove that $\text{poly}(TA)$ patrols suffices to protect all targets. The output is a distribution over $\text{poly}(TA)$ -many pure strategies.

We improve over the state-of-art prior work [8, 20] in several ways. First, the running time in [8] is exponential in the number of patrols (K) and becomes impractical even for $K = 3$ [20]. Second, [20] achieves a polynomial running time only under a substantial assumption: either a constant number of rounds, or that all targets have a unit value at all times, or that the “protection ranges” of the patrols are so small that they cannot overlap for any two adjacent patrol locations. Third, the running times in [8, 20] depend polynomially on the the number of patrol locations (M). Finally, the polynomial running time in [20] relies on the Ellipsoid Algorithm for solving linear programs, which is notoriously slow in practice.

Our techniques. Our main algorithm works on the *discretized version* of the problem, in which the possible locations for patrols are integers from 1 to M (there is no such restriction

on the location of targets). The algorithm consists of three parts: partitioning the spatial domain, formulating patrol placements in a single time point, and combining them to find the optimal strategy for all time points. Below we describe these three parts one by one.

First, we partition the spatial domain into a relatively small number of intervals so that the patrol locations inside each interval are “equivalent” to one another as far as our problem is concerned. Then we use these intervals as “atomic” patrol locations, thereby replacing the dependence on M with the dependence on the number of intervals. The partitioning algorithm starts from the last round T and goes backwards in time: for each round t it constructs a collection of intervals based on the target locations at time t and the intervals constructed for time $t + 1$, so as to ensure the desired “equivalence” property. We bound the number of intervals by $\mathcal{O}(T^3A)$.

Second, for each time point t we construct a graph G_t which models any possible *snapshot* of the patrol placements at this time. More specifically, every patrol placement at time t can be mapped to a specific path in G_t , and every *randomized* patrol placement can be mapped to a specific unit flow in G_t . Furthermore, we define the cost of a path/flow in G_t such that it equals the maximum utility of the attacker under the corresponding (randomized) patrol placement at time t , and can be computed via a linear program.

Third, we create a linear program that “unifies” the graphs G_t , and use this LP to construct the minimax strategy for the defender. The LP ensures that the (randomized) patrol placements computed in each G_t are consistent with one another, in the sense that there is a valid transition from one round to the next, without violating the speed restriction. To accomplish this, the LP finds min-cost flows in each G_t , and includes additional linear constraints that guarantee consistency. We post-process the solution of this LP and remove the crossing edges in the flows. Finally, we incrementally construct a mixed strategy of the defender based on the post-processed solution and prove that it is indeed the optimal strategy.

In the *continuous* version of the problem, patrol locations can take arbitrary real values, and the target locations are rational. We first re-scale all target locations to integers, and prove that this re-scaled problem instance admits a discrete solution. Then we use the algorithm from the discretized version. It is essential that the running time of the latter is polylogarithmic in M .

Related work. Security games have been studied extensively in the past decade, see the book [18] as well as more recent work, e.g. [5, 8, 12, 19, 20]. The research concerned both theoretical foundations as well as applications. Publicized real-world deployments include: US Coast Guard patrol boats [8], canine-patrol and vehicle-checkpoints scheduling in Los Angeles airport (LAX) [16], scheduling flights for air marshals by US Federal Air Marshal Service [10], airport passenger screening by US Transportation Security Administration [6], fare inspection in Los Angeles transit system [21], and wildlife protection in Malaysia [9].

Most relevant to the present work are papers on computing minimax strategy in zero-sum spatio-temporal security games. While the initial work assumed static targets [18], some of the later work addressed moving targets [5, 8, 20] (as discussed above). On a related note, if the patrols are allowed to *accelerate*, with an upper bound on the acceleration, then computing the defender’s minimax strategy becomes NP-hard [19, 20]. Other work concerned solving security games that are not (necessarily) spatio-temporal or zero-sum, e.g. [7, 11, 20]. A notable line of work in security games assumes that the defender does not fully know attackers’ values for the targets, but can learn more about them over multiple rounds

of interaction with the said attackers (see [15] for a recent survey of a subset of this work, as well as [2, 4, 13, 14]).

In a broader game-theoretic context, our work is related to Stackelberg games and equilibrium computation. Originally introduced to model competing firms, Stackelberg games is a classic concept in game theory which appears in many textbooks and countless papers.

Computing Nash Equilibria is a central problem in algorithmic economics. While this problem is known to be PPAD-hard in general, polynomial-time algorithms exist for many natural classes of games, particularly for zero-sum games (for background, see a survey [17] and references therein). Yet, these algorithmic results are insufficient for games in which the number of pure strategies can be exponential in the input size (see [1, 3] for examples of such games).

Further directions. Many ideas in this paper may be useful for solving other spatio-temporal security games. In particular, the overall algorithmic framework of locally solving each “time layer” under some compatibility constraints and then merging the “time layers” to compute the global optimum solution appears broadly applicable.

We believe our techniques can be extended to achieve a polynomial time algorithm for several extensions of the model. In particular, we can incorporate additional constraints on the patrols, such as obstacles that the patrols cannot cross over, or speed limits that depend on a particular location. We can also handle scenarios when the spatial domain or the timeline are not evenly discretized. (For ease of presentation, we do not include these extensions in the present paper).

A general way to model such extensions is to assume that the range of valid movements for each location is given in the input. Whenever we still have a property that the patrols do not need to cross each other in an optimal solution (which indeed is a very natural property for homogeneous patrols), our techniques achieve a polynomial-time algorithm in the input size. However, the input size for this extended model gets large, and no longer scales polylogarithmically in M .

That said, some important special cases allow for succinct input. For example, a small number of obstacles can be specified directly, rather than given implicitly via the ranges of valid movements. Designing a polynomial-time algorithm for such cases requires a problem-specific pre-processing step for partitioning the locations (which could potentially be very different from the partitioning step in this paper). However, the rest of the algorithm could be essentially the same.

It is very tempting to extend our model to a two-dimensional space. We believe some of our techniques can be useful for this extension, most importantly the compatibility constraint technique from Section 3.3. The main challenge in extending our approach is an appropriate generalization of the “day graphs”.

2 PRELIMINARIES

Our goal is to find an optimal patrol scheduling strategy for the defender to protect a set of A mobile targets in a one-dimensional space. Figure 1 illustrates the problem in a 2-D diagram; the x-axis denotes the evenly discretized temporal domain containing $T + 1$ time points, and the y-axis denotes the one-dimensional space of length M . We say a spatial position i is above j , if $i > j$ and similarly define a spatial position i to be under j , if $i < j$.

The defender has K homogeneous patrols to protect a set of moving targets from a potential attack. Patrols have a maximum speed of Δ . This means, a move from a position

Fig. 1. An illustration of how targets and patrols move in discrete (a) and continuous (b) models of the problem.

m_t at time t to m_{t+1} at time $t+1$ is invalid if $|m_{t+1} - m_t| > \Delta$. We consider two models of the problem: *discretized model*, denoted by $\mathcal{DSG}$ (Figure 1-a), and *continuous model*, denoted by $\mathcal{CSG}$ (Figure 1-b). In the discretized model, the position of any patrol at any time point is an integer between 0 and M , but in the continuous model, the patrol locations are not restricted to be integers (note that the temporal domain is still discretized). Furthermore, for any target a , $h_{a,t}$ and $w_{a,t}$ respectively denote its position and weight at time t . Note that in both models, there is no restriction on the position and the speed of targets and the weight of the targets could change from time to time (e.g., ferries may not carry the same number of people at different times). The patrols protect any target within their protection radius R (a fixed number for all patrols). That is, a patrol k at position m_t at time t , protects a target a , if $|h_{a,t} - m_t| \leq R$. In Figure 1, the grey ranges around patrols denote the area they protect. We denote the set of targets, the set of spatial positions, the set of patrols and the set of all time points by $[A]$, $[M]$, $[K]$ and $[T]$ respectively.

A *patrol path*, is a sequence of T positions $(m_1, m_2, \dots, m_T)$, such that for any $t \in [T]$, a move from m_t to m_{t+1} does not violate the speed limit (the black paths in Figure 1 denote patrol paths). A pure strategy of the defender is a set of K patrol paths denoted by $\{v_k\}_{k \in [K]}$. A mixed strategy of the defender, is a probability distribution over her pure strategies. A pure strategy of the attacker is a single target-time pair (a, t) which means the attacker attacks target a at time t . Let $\{v_k\}_{k \in [K]}$ be the pure strategy of the defender and (a, t) be the pure strategy of the attacker, attacker's utility is 0 if target a is protected by at least one patrol at time t and it is $w_{a,t}$ if it is not protected by any patrols. We assume the game is zero-sum and find minmax strategies.

Without loss of generality, we can assume $K \leq TA$. This observation comes from the fact that with only TA patrols, the defender can provide a 100% protection without needing any more patrols. To do this, for any target-time pair (a, t) , the defender can put a still patrol at the location of target a at time t .

3 DISCRETE MODEL

The main goal of this section is to prove the following theorem:

THEOREM 3.1. *There is a polynomial time (in input size) algorithm to solve $\mathcal{DSG}$.*

Algorithm 1 Partitions the given positions into intervals

```

1: function GETINTERVALS
2:   for  $t = T$  to 1 do
3:     if  $t < T$  then  $P_t \leftarrow \text{GETINTERVALPOINTS}(t, P_{t+1})$ 
4:     else  $P_t \leftarrow \text{GETINTERVALPOINTS}(t, \emptyset)$ 
5:      $\mathcal{I}_t \leftarrow \text{Array}()$ 
6:     for any two consecutive items  $p_i$  and  $p_{i+1}$  in  $P_t$  do
7:        $\mathcal{I}_t.\text{INSERT}([p_i, p_{i+1}])$ 
8:   return  $\langle \mathcal{I}_1, \dots, \mathcal{I}_T \rangle$ 
9: function GETINTERVALPOINTS( $t, P_{t+1}$ )
10:   $P_t \leftarrow \text{SortedArray}()$ 
11:   $P_t.\text{INSERT}(0), P_t.\text{INSERT}(M)$ 
12:   $\epsilon \leftarrow$  a sufficiently small number in  $\mathbb{R}^+$ 
13:  for each target  $a \in [A_t]$  do
14:    if  $h_{a,t} - R > 0$  then  $P_t.\text{INSERT}(h_{a,t} - R)$ 
15:    if  $h_{a,t} + R < M$  then  $P_t.\text{INSERT}(h_{a,t} + R + \epsilon)$ 
16:  for  $p$  in  $P_{t+1}$  do
17:    if  $p - \Delta > 0$  then  $P_t.\text{INSERT}(p - \Delta)$ 
18:    if  $p + \Delta < M$  then  $P_t.\text{INSERT}(p + \Delta)$ 
19:  for any item  $p_i$  in  $P_t$  do
20:    if interval  $[p_i, p_{i+1})$  does not contain any patrol position then
21:       $P_t.\text{REMOVE}(p_i)$ 
22:  return  $P_t$ 

```

3.1 Partitioning The Positions

The number of pure strategies in a single time point, even with only one patrol, is not polynomial in the input size, since the number of possible locations, M , could be exponentially larger than the input size. To overcome this difficulty, we partition the spatial positions into polynomially many sets of consecutive positions which we call *intervals* and we only keep track of these intervals instead of maintaining the exact position of a patrol within the intervals. For example, assume there is only one target, one patrol and $T = 1$, then it only matters whether the patrol's protection range contains the location of the target or not and the exact position of the patrol does not matter.

We use Algorithm 1 to partition the positions into meaningful intervals. For any given time point t , the function GETINTERVALPOINTS, generates a sorted array $P_t = \langle p_1, \dots, p_{n_t} \rangle$ of numbers that we call *interval points* and GETINTERVALS uses these generated interval points to partition the spatial positions of any time point t to intervals:

$$\mathcal{I}_t = \langle [p_1, p_2), [p_2, p_3), \dots, [p_{n_t-1}, p_{n_t}) \rangle.$$

The intervals are assumed to be left-closed and right-open to simplify the calculations. We use $\mathcal{I}_t[i]$ to denote the i -th interval in $\mathcal{I}_t$ and use $\mathcal{I}_t[i : j]$ to denote the set of consecutive intervals $\{\mathcal{I}_t[i], \mathcal{I}_t[i+1], \dots, \mathcal{I}_t[j]\}$.

The following lemma proves the total number of intervals is polynomial in the input size and as a corollary of that, Algorithm 1 runs in polynomial time in the input size.

Fig. 2. Interval points are added to the partitioning sets in two ways: (a) around any target there are two interval points, and (b) any interval point propagates to the previous time points. A combination of both is shown in (c).

LEMMA 3.2. *The total number of intervals created by Algorithm 1 is $\mathcal{O}(T^3A)$.*

Proof. To prove this, we charge any interval point to a target/time pair and show no target/time pair will be charged more than $\mathcal{O}(T^2)$ times and since there are at most $\mathcal{O}(AT)$ target/time pairs, there will not be more than $\mathcal{O}(T^3A)$ interval points. Note that there are two ways for interval points to be added to a partitioning set P_t (Figure 2):

- (1) For any target a at time point t , two interval points $h_{a,t} - R$ and $h_{a,t} + R + \epsilon$ are added to P_t . We charge these two interval points to (a, t) .
- (2) For any interval point p in P_{t+1} two interval points $p - \Delta$ and $p + \Delta$ are added to P_t . We recursively charge these two intervals to the target/time pair that the interval point p is charged to.

This means an interval point p could be charged to a target/time pair (a, t) , only if $h_{a,t} = p \pm i\Delta + R + \epsilon$ or $h_{a,t} = p \pm i\Delta - R$ for any integer i where $i \leq t$. Although this condition is only necessary and not sufficient, but it implies at most $\mathcal{O}(T^2)$ interval positions could be charged to an arbitrary target/time pair and therefore the total number of partition points is $\mathcal{O}(T^3A)$. $\square$

Note that since by Lemma 3.2, the total number of interval points at any time point t is polynomial in the input size, function $\text{GETINTERVALPOINTS}(t, P_{t+1})$, which is simply a loop over P_{t+1} and the targets, runs in polynomial time.

COROLLARY 3.3. *Algorithm 1 halts in polynomial time in the input size.*

The following lemmas prove two important properties of the partitions generated by Algorithm 1. These properties basically imply all patrol locations within the same interval are equivalent as far as the problem is concerned.

LEMMA 3.4. *Let k and k' be two patrols in the same interval at any time t . The set of targets that k and k' protect at time t are equal.*

LEMMA 3.5. *Let $[s_i, f_i]$ and $[s_j, f_j]$ be two arbitrary intervals in $\mathcal{I}_t$ and $\mathcal{I}_{t+1}$ respectively. If there exists a feasible move from an arbitrary position in $[s_i, f_i]$ to a position in $[s_j, f_j]$, for any position in $[s_i, f_i]$, there exists a feasible move to a position in $[s_j, f_j]$.*

We can now define the feasible set of a set of consecutive intervals:

Fig. 3. The white intervals at time $t + 1$, are in the feasible set of the white intervals at time t

Definition 3.6 (feasible sets). We define the feasible set of $\mathcal{I}_t[i : j]$, denoted by $\text{feas}_t(\mathcal{I}_t[i : j])$, to be a subset of $\mathcal{I}_{t+1}$, containing an interval $\mathcal{I}_{t+1}[i']$ iff there exists a feasible move from a position in some interval in $\mathcal{I}_t[i : j]$ to $\mathcal{I}_{t+1}[i']$. We may occasionally abuse this notation and use the simpler form of $\text{feas}_t(\mathcal{I}_t[i])$ instead of $\text{feas}_t(\mathcal{I}_t[i : i])$ (Figure 3).

It is easy to see the following corollary of Definition 3.6:

COROLLARY 3.7. *For any $\mathcal{I}_t[i : j]$ there exists a consecutive interval set $\mathcal{I}_{t+1}[i' : j']$ such that*

$$\text{feas}_t(\mathcal{I}_t[i : j]) = \mathcal{I}_{t+1}[i' : j'].$$

Definition 3.8 (interval path). We define an *interval path* to be a sequence of T intervals $\langle \mathcal{I}_i[x_i] \rangle_T$, such that for any time point $t \in [T]$, $\mathcal{I}_{t+1}[x_{t+1}] \in \text{feas}_t(\mathcal{I}_t[x_t])$. Moreover, for any interval path $\xi = \langle \mathcal{I}_i[x_i] \rangle_T$, we define $S(\xi)$ to be the set of all patrol paths that are within ξ . More formally, a patrol path $\langle m_i \rangle_T$ is within $S(\xi)$ if and only if for any time point $t \in [T]$, m_t is in interval $\mathcal{I}_t[x_t]$.

Note that any patrol path is within exactly one interval path, since intervals do not overlap and they cover all locations. It could also be obtained from the following lemma that $S(\xi)$ is never empty for an interval path ξ . The proof is to choose any position in $\mathcal{I}_1[x_1]$ and following the valid movements until we reach a position in $\mathcal{I}_T[x_T]$.

LEMMA 3.9. *For any interval path $\xi = \langle \mathcal{I}_i[x_i] \rangle_T$, there is at least one patrol path in $S(\xi)$.*

Note that by Lemma 3.4, two patrols that are in the same interval, protect the same set of targets at that specific time. This implies that the patrol paths that are within the same interval path, protect the same set of targets at all times and could be replaced with one another in any strategy, without changing the utilities. This means the amount of information encoded in an interval path is sufficient to describe the important characteristics of strategies and find the optimal one.

3.2 Strategies In a Single Time Point

In this section we explain how to locally find the best strategy for a single time point ignoring the speed limitations. Note that although we proved the number of intervals is polynomial in the input size, there are still exponentially many different ways to place our K patrols in them. This section describes how we can resolve this problem and find the best strategy.

We use the term *snapshot* to denote a patrol placement at a single time point and formally define it as follows:

Fig. 4. Figure (a) shows the vertices of a day graph and figure (b) shows a sample day graph. The highlighted path in figure (b) shows a canonical path and its equivalent patrol placement is shown in the vertical bar next to it, each grey piece denotes an intervals and the black dots denote the location of patrols in them.

Definition 3.10 (snapshots). A pure snapshot at time point t , is an assignment of patrols to intervals of time t . We denote it by a sorted sequence of K intervals $\langle \mathcal{I}_t[y_i] \rangle_K$ such that for any $i \in [K]$, $y_i \leq y_{i+1}$. A mixed snapshot, denoted by $\{(d_i, p_i)\}_n$ is a probability distribution over n pure snapshots where p_i denotes the probability of choosing pure snapshot d_i and $\sum_{i=1}^n p_i = 1$.

For any time point t , we construct a weighted directed graph G_t (called a *day graph*) and give a one-to-one mapping between pure snapshots at time t and paths from S_t (source vertex) to S'_t (sink vertex), where S_t and S'_t are two specific vertices of G_t . Moreover, we map any mixed snapshot at time point t , to a network flow of 1 unit from S_t to S'_t . This mapping is not necessarily one-to-one and many mixed snapshots may be mapped to the same network flow; however, the maximum utility of the attacker in all such mixed snapshots, will be the same.

In Definition 3.11 we formally explain how G_t is constructed. An informal explanation of it is as follows: the vertex set of G_t , as shown in Figure 4-a, includes a vertex S_t , a vertex S'_t and a grid of $K \times n_t$ vertices (recall that n_t is the number of intervals at time t) each denoted by $V_t[x, y]$. There is an edge from S_t to any vertex in the first column of the grid $V_t[1, y]$, and there is an edge from any vertex in the last column of the grid $V_t[K, y]$ to S'_t . Also for any x, y , and y' , there is an edge from $V_t[x, y]$ to $V_t[x + 1, y']$ if $y \leq y'$. Furthermore, we define a *canonical path* to be any path from S_t to S'_t and define a *canonical flow* to be any flow of unit 1 from S_t to S'_t (Definition 3.12). We give a one-to-one mapping between canonical paths in G_t and pure snapshots at time point t in Definition 3.13 and map any mixed snapshot to a canonical flow in Definition 3.14. Figure 4-b shows a sample day graph, a canonical path in it and its equivalent pure snapshot. Moreover, we assign weights to the edges of G_t such that the maximum payoff of the attacker for a pure (mixed) snapshot equals the cost of its corresponding canonical path (flow). The cost of a canonical path and a canonical flow is defined in Definition 3.12.

For any target a and any intervals i and i' at time t , we define the binary value $\mathcal{C}_t(i, i'; a)$ to be 1 if the following two conditions hold: (1) target a is located in a position between i and i' (non-inclusive) at time t , and, (2) a could not be protected by any patrol at any

arbitrary position in i or i' ; otherwise we set $\mathcal{C}_t(i, i'; a)$ to be 0. We similarly define two binary variables $\mathcal{C}_t(\emptyset, i; a)$ and $\mathcal{C}_t(i, \emptyset; a)$ for the border cases. We set $\mathcal{C}_t(\emptyset, i; a)$ to be 1 iff target a is in a position below i where no patrol in i can protect it and set $\mathcal{C}_t(i, \emptyset; a)$ to be 1 iff target a is in a position above i where no patrol in i can protect it. Assume a patrol placement p does not protect a target a at time t . Let i and i' be the closest intervals to target a , that contain at least one patrol and are below and above a respectively (set them to be $\emptyset$ if no such interval exists), then by definition $\mathcal{C}_t(i, i'; a) = 1$. Using this definition, we can now formally define a day graph and canonical path/flow.

Definition 3.11 (day graph). Given a time point t , we construct graph G_t as follows:

- (1) Graph G_t contains a vertex S_t (source), a vertex S'_t (sink) and $K \times n_t$ other vertices, each denoted by $V_t[x, y]$ for $1 \leq x \leq K$ and $1 \leq y \leq n_t$.
- (2) For any y such that $1 \leq y \leq n_t$, there is an edge from S_t to $V_t[1, y]$. If e denotes an edge of this kind, for any target a , we define $c_{e,a}$ to be $\mathcal{C}_t(\emptyset, \mathcal{I}_t[y]; a)$.
- (3) For any y such that $1 \leq y \leq n_t$, there is an edge from $V_t[K, y]$ to S'_t . If e denotes an edge of this kind, for any target a , we define $c_{e,a}$ to be $\mathcal{C}_t(\mathcal{I}_t[y], \emptyset; a)$.
- (4) For any two vertices $V_t[x, y]$ and $V_t[x+1, y']$, if $y \leq y'$, there is an edge from $V_t[x, y]$ to $V_t[x+1, y']$. If e denotes this edge, for any target a , we define $c_{e,a} = \mathcal{C}_t(\mathcal{I}_t[y], \mathcal{I}_t[y']; a)$.

Definition 3.12 (canonical path/flow). In a day graph G_t any path from S_t to S'_t is a canonical path. Let $E = \{e_1, e_2, \dots, e_{K+1}\}$ be the set of edges in a canonical path, and a be an arbitrary target. We define the cost of this canonical path for target a to be:

$$\sum_{e \in E} c_{e,a} \cdot w_{t,a} \quad (1)$$

Also, any flow of unit 1 from S_t to S'_t is a canonical flow. We denote any canonical flow with a function $f : E(G_t) \rightarrow [0, 1]$, where $f(e)$ denotes the flow passing through an edge e . The cost of an arbitrary canonical flow f , for day graph G_t is:

$$\max \sum_{e \in E(G_t)} f(e) \cdot c_{e,a} \cdot w_{t,a} \quad \forall a; \text{ where } a \in [A.] \quad (2)$$

Intuitively speaking, $c_{e,a}$ is 1 if and only if having edge e in a canonical path p implies that in the “equivalent” pure snapshot of p , target a is not covered by any patrol. The formal mapping of canonical paths and flows to snapshots is as follows.

Definition 3.13 (pure snapshot mapping). Let $s_t = \langle \mathcal{I}_t[y_i] \rangle_K$ be a pure snapshot (recall that $y_i \leq y_{i+1}$ for any $i \in [K]$). We map s_t to the following canonical path:

$$p_t = \langle S_t, V_t[1, y_1], V_t[2, y_2], \dots, V_t[K, y_K], S'_t \rangle$$

and similarly map p_t to s_t and say s_t and p_t are equivalent.

Definition 3.14 (mixed snapshot mapping). Let $m = \{(d_i, p_i)\}_n$ be a mixed snapshot at time t . Also let f_i denote a flow of unit p_i from S_t to S'_t through the edges of the equivalent canonical path of d_i . We construct flow f as follows: for any edge e of G_t , $f(e) = \sum_{i=1}^n f_i(e)$. Note that since by definition of a mixed snapshot, $\sum_{i=1}^n p_i = 1$, f is a flow of unit 1 from S_t to S'_t , and hence is a canonical flow. We map m to f .

In Lemma 3.15 we prove that the payoff of the attacker if he attacks target a at time t while the placement of patrols is represented by the pure strategy s , equals to the cost of the target a in the canonical path equivalent to s . Then in Lemma 3.16 we prove that the maximum payoff of the attacker at time t while the strategy of defender is represented

by the mixed snapshot m is equal to the cost of canonical flow equivalent to m . These two lemmas can be directly obtained by the given definitions, however, for space limitations, their formal proofs are left to the appendix.

LEMMA 3.15. *Let s be a pure snapshot at time t , and a be an arbitrary target. The payoff of the attacker with respect to s , if he attacks the target a at time t , equals the cost of the target a in the canonical path equivalent (Definition 3.13) to s .*

LEMMA 3.16. *Let r be a mixed snapshot at time t and let f denote the canonical flow that r is mapped to. The maximum expected payoff of the attacker at time t with respect to r , equals the cost of f .*

3.3 Best Strategy For All Time Points

Lemma 3.16 implies if our goal is to minimize the maximum payoff of the attacker at a single time point t , it suffices to find a canonical flow in G_t with minimum cost. Although this works for the special case when $T = 1$, but it does not consider the movement of patrols and their speed limits. More precisely, a pure strategy for the defender could be shown as a sequence of pure snapshots $\langle s_1, s_2, \dots, s_T \rangle$. However there is one important condition: for any $i \in [T]$, there must be a feasible transition from s_i to s_{i+1} . This is also the case for mixed snapshots and two consecutive ones may not be necessarily compatible. In this section we resolve this issue and prove Theorem 3.1.

Our algorithm to find the optimal strategy of the defender consists of three main steps. In the first step, which is explained in more details in Section 3.3.1, we run an LP that returns a canonical flow for each day graph $G_1, \dots, G_T$. Apart from the constraints to ensure we get valid canonical flows with minimum overall cost, our LP contains an extra constraint for compatibility of these canonical flows. In the second step (Section 3.3.2), while keeping the overall characteristics of these canonical flows unchanged, we adjust them in a way to make sure no two crossing edges in any of the day graphs have a positive flow. Finally, in the third step (Section 3.3.3), we construct a mixed strategy for the defender based on the adjusted canonical flows.

Let s_t and s_{t+1} denote two pure snapshots representing the placement of patrols in a valid pure strategy p at two consecutive times. In the following lemma we prove that there exists a feasible move from i -th interval of s_t to the i -th interval of s_{t+1} (recall that intervals in pure snapshots are sorted based on their position). We prove this lemma by induction on the number of patrols. At each step we prove that there exists a feasible move from the top most interval in s_t to the top most interval in s_{t+1} and we prove if we match these two together and remove them, we can construct another pure strategy that contains the remaining intervals.

LEMMA 3.17. *If $\langle \mathcal{I}_t[y_i^t] \rangle_K$ and $\langle \mathcal{I}_{t+1}[y_i^{t+1}] \rangle_K$ are two pure snapshots at time t and $t + 1$ in at least one valid pure strategy p , then for any $j \in [K]$ we have $\mathcal{I}_{t+1}[y_j^{t+1}] \in \text{feas}_t(\mathcal{I}_t[y_j^t])$.*

In the following definition, we define what it means for a patrol path to be *intervally above*, *below* or *equal* to another patrol path:

Definition 3.18. Let $v = \langle m_0, m_1, \dots, m_T \rangle$ and $v' = \langle m'_0, m'_1, \dots, m'_T \rangle$ be two patrol paths. We say v and v' are *intervally equal* if for any $t \in [T]$, m_t and m'_t are in the same interval. We also define v to be *intervally under* v' , if for any $t \in [T]$, either $m_t < m'_t$ or m_t and m'_t are in the same interval. Similarly, we define v to be *intervally above* v' , if v' is *intervally under* v .

Next, in Lemma 3.19 we prove that there exists an optimal strategy of the defender that for any pure strategy p in its support, there is an ordering of interval paths in p such that, i -th interval path is always intervally under j -th interval path if $1 \leq i < j \leq K$. To prove this we use Lemma 3.17 that indicates there exists a possible move from the k -th interval in pure snapshot s_1 to the k -th interval in pure snapshot s_2 if s_1 and s_2 represent the patrols' placement of a pure strategy in two consecutive times and $1 \leq k \leq K$. For any patrol $k \in [K]$, we construct an interval path that contains the k -th interval of all the pure snapshots in p , and we assign patrol k to this interval path. It is easy to see that if we order the patrols from 1 to K the interval path assigned to patrol k_1 is under the interval path assigned to patrol k_2 if $1 \leq k_1 < k_2 \leq K$. This lemma is very similar to Lemma 3 of [20] but adopted to intervals paths.

LEMMA 3.19. *There exists an optimal mixed strategy of the defender, such that for every pure strategy p in its support there is an ordering of interval paths $\langle \xi_1, \dots, \xi_K \rangle$ such that the following condition holds for this ordering: for any two interval paths ξ_i and ξ_j , in the pure strategy p , ξ_i is intervally under ξ_j if $i < j$.*

Again, for space limitations, the full proof is left to the appendix. However, intuitively, starting from any given optimal solution one can swap the remaining path of any two patrols that cross each other without losing anything. This eventually resolves all crosses and gives a desired optimal solution.

Let s denote an optimal strategy of the defender that satisfies the condition mentioned in Lemma 3.19, and let $\langle \xi_{p,1}, \xi_{p,2}, \dots, \xi_{p,K} \rangle$ denote the ordering of interval paths in pure strategy p in support of s such that $\xi_{p,i}$ is intervally under $\xi_{p,j}$ if $1 \leq i < j \leq K$. Without loss of generality we assume for any $i \in [K]$ the same patrol is assigned to interval path $\xi_{p,i}$ for all p in support of s , and it is denoted by k_i . Therefore, if $\langle \mathcal{I}_t[y_1], \mathcal{I}_t[y_2], \dots, \mathcal{I}_t[y_K] \rangle$ denotes a pure snapshot that represents the patrols' placement in an arbitrary time point $t \in [T]$ in pure strategy p in support of s , $\mathcal{I}_t[y_i]$ is the position of patrol k_i in pure strategy p at this time. Moreover, let m denote the mixed snapshot of strategy s at time t . The flow passing through the vertex $V_t[i, j]$, denotes the probability with which patrol k_i is placed in the j -th interval at time t . So, all the data related to position of patrol k_i at time t is in the column i of the day graph of time t . We use this later in the paper.

3.3.1 Linear Programming. In this section we explain how the first step of our algorithm, the LP, works.

Note that a flow of 1 unit in G_t with minimum weight, minimizes the attacker's payoff at time t . To minimize the attacker's payoff at all time points, we need to minimize the cost of the canonical flow with the maximum cost. To do this, for any edge e in any day graph G_t we define an LP variable $f_t(e)$ which specifies the amount of flow passing through e . Moreover for any time point t , we include the following constraints in our LP:

- (1) For each vertex v of G_t (except for the source vertex S_t and the sink vertex S'_t) the amount of ingoing flow to v is equal to the amount of outgoing flow from v .
- (2) The amount of outgoing flow from S_t is 1.
- (3) The amount of ingoing flow to S'_t is 1.
- (4) The amount of flow passing through any edge e is not negative.
- (5) The cost of flow through any edge e and for any target a in G_t , specified by $w_{e,a} \times f_t(e)$ is not more than u .

And we set the objective function of our LP to minimize u , which is the overall cost of canonical flows. However, as we said earlier the canonical flows we find must be compatible;

$$\min \quad u \quad (3)$$

$$f_t^-(v) = f_t^+(v) \quad \forall t, v : t \in [T], v \in V(G_t) - \{S_t, S'_t\} \quad (4)$$

$$f_t^+(S_t) = 1 \quad \forall t : t \in [T] \quad (5)$$

$$f_t^-(S'_t) = 1 \quad \forall t : t \in [T] \quad (6)$$

$$f(e) \geq 0 \quad \forall e, t : t \in [T], e \in G_t \quad (7)$$

$$\sum_{e \in E(G_t)} f(e) \cdot c_{e,a} \cdot w_{t,a} \leq u \quad \forall t, a : t \in [T], a \in [A_t] \quad (8)$$

$$\sum_{v \in V_t[k;i:j]} f_t^+(v) \leq \sum_{v' \in \text{feas}_t(V_t[k;i:j])} f_{t+1}^+(v') \quad \forall t, k, i, j : t \in [T], k \in [K], 1 \leq i \leq j \leq n_t \quad (9)$$

Linear Program 1. Variable $f_t(e)$, which is defined for any edge e in the day graph G_t where t could be any time point in $[T]$, denotes the amount of flow passing through e . By $f_t^-(v)$ we mean the total flow coming into vertex v (i.e., $f_t^-(v) = \sum f_t(e)$ where e is any edge ending at v). Also $f_t^+(v)$ denotes the total flow coming out of vertex v (i.e., $f_t^+(v) = \sum f_t(e)$ where e is any edge starting from v).

therefore apart from the aforementioned constraints, we define a compatibility constraint. Recall that $\text{feas}_t(\mathcal{I}_t[i : j])$ denotes a collection of intervals at time $t + 1$ and contains an interval i' , iff there is a valid move from an interval in $\mathcal{I}_t[i : j]$ to i' . We define a very similar concept for day graphs:

Definition 3.20. Let $V_t[x; i : j]$ denote the set of consecutive grid vertices $\{V_t[x, i], V_t[x, i + 1], \dots, V_t[x, j]\}$ in G_t . Recall that by definition of $\mathcal{I}_t[i : j]$, any vertices in $V_t[x; i : j]$ is equivalent to an interval in $\mathcal{I}_t[i : j]$. We define $\text{feas}_t(V_t[x; i : j])$ as follows: $\text{feas}_t(V_t[x; i : j])$ is a subset of grid vertices of G_{t+1} containing a vertex $V_{t+1}[x, i']$ if and only if $V_{t+1}[x, i']$ is equivalent to an interval in $\text{feas}_t(\mathcal{I}_t[i : j])$.

The compatibility constraint we use in our LP is as follows: for any set of consecutive vertices $V_t[x; i : j]$, the amount of flow passing through the vertices in $V_t[x; i : j]$ is not more than the amount of flow passing through the vertices in $\text{feas}_t(V_t[x; i : j])$. Intuitively, this constraint indicates that for any set of consecutive intervals $\mathcal{I}_t[i : j]$, the probability that there exists a patrol in it, should not be more than the probability of having a patrol in its feasible set ($\text{feas}_t(\mathcal{I}_t[i : j])$) in the next time point. Note that by definition, $\text{feas}_t(\mathcal{I}_t[i : j])$ contains all of the valid intervals that a patrol in $\mathcal{I}_t[i : j]$ can move to; therefore it is obvious why this constraint is necessary. The sufficiency of this constraint to prove compatibility of snapshots, however, comes later when we explain how we construct an optimal strategy based on the adjusted LP solution. The formal definition of the LP is given in Linear Program 1.

By the end of Section 3.3, we prove the solution of LP 1 is equal to the utility of the attacker if both players play their optimal strategies. Lemma 3.21 proves a weaker claim:

LEMMA 3.21. *The solution of Linear Program 1 gives a lower bound for the utility of the attacker when both players play their optimal strategies.*

To prove Lemma 3.21, we start from an optimal strategy of the defender satisfying the condition of Lemma 3.19 that the interval paths do not cross each other. Based on this strategy, we construct a feasible solution for LP 1 in which the value of u is equal to the maximum possible utility of the attacker. Note that this only proves LP 1 gives a lower bound

Fig. 5. The whole idea of how to locally resolve the issue of having a crossing flow. The details are formally explained in the text.

for the utility of the attacker when both players play their minimax strategies since the feasible solution we considered is not necessarily the optimum solution of the LP (although as we said before, we will later prove that they are exactly the same).

3.3.2 Adjusting The LP Solution. In this section we give an algorithm that adjusts any optimal solution of LP 1 to resolve their “crossing flows”. We later use the adjusted solution to construct the defender’s optimal strategy.

We start by defining what we mean by *crossing edges* and *crossing flows*:

Definition 3.22 (crossing edges and crossing flow). Let $e = (V_t[x, y_1], V_t[x + 1, y_2])$ and $e' = (V_t[x, y'_1], V_t[x + 1, y'_2])$ be two arbitrary edges of day graph G_t where $y_1 < y'_1$. We say e and e' cross, if and only if $y_2 > y'_2$. Moreover, if f is a canonical flow of G_t , the crossing edge pair (e, e') is a crossing flow in f , if $f(e) > 0$ and $f(e') > 0$.

In the following definition, we give a total ordering on the crossing flows in a canonical flow, which we later use in the algorithm we provide to resolve them.

Definition 3.23 (crossing flows’ ordering). Let (e_1, e_2) and (e_3, e_4) be two crossing flows of a canonical flow. Also let $e_1 = (V_t[x, y_1], V_t[x + 1, y'_1])$, $e_2 = (V_t[x, y_2], V_t[x + 1, y'_2])$, $e_3 = (V_t[x', y_3], V_t[x' + 1, y'_3])$, and $e_4 = (V_t[x', y_4], V_t[x' + 1, y'_4])$ be the vertices of these edges. We say $(e_1, e_2) < (e_3, e_4)$ if and only if one of the following conditions hold:

- (1) $x < x'$
- (2) $x = x'$ and $y_1 < y_3$.
- (3) $x = x'$ and $y_1 = y_3$ and $y'_1 < y'_3$.
- (4) $x = x'$ and $y_1 = y_3$ and $y'_1 = y'_3$ and $y'_2 < y'_4$.
- (5) $x = x'$ and $y_1 = y_3$ and $y'_1 = y'_3$ and $y'_2 = y'_4$ and $y_2 < y_4$.

Algorithm 2 is the formal pseudo-code of how we resolve all crossing flows. At each step, the algorithm resolves the minimum crossing flow (minimum based on the total ordering defined in Definition 3.23) and continues this process until there is no other one. Figure 5 illustrates how a single crossing flow is resolved.

LEMMA 3.24. *Function RESOLVECROSSES of Algorithm 2 does not increase the total cost of the input canonical flow.*

To prove Lemma 3.24, we consider all different locations of targets for which changing the flows might affect the utility of players and prove in none of these cases the total cost is increased. The complete proof is left to the appendix.

Algorithm 2 Resolves crossing flows

```

1: function RESOLVEXCROSSES( $f_1, G_1, f_2, G_2, \dots, f_T, G_T$ )
2:   for each timepoint  $t \in [T]$  do
3:     while  $f_t$  contains any crossing flow do
4:        $\{(V_t[x, y_1], V_t[x+1, y_2]), (V_t[x, y'_1], V_t[x+1, y'_2])\} = \text{FINDNEXTCROSS}(G_t, f_t)$ 
5:        $e \leftarrow (V_t[x, y_1], V_t[x+1, y_2])$ 
6:        $e' \leftarrow (V_t[x, y'_1], V_t[x+1, y'_2])$ 
7:       RESOLVEXCROSS( $e, e'$ )
8:   function FINDNEXTCROSS( $G_t, f_t$ )
9:      $(e, e') \leftarrow$  find the minimum crossing flow in  $f_t$ 
10:    return  $(e, e')$ 
11:  function RESOLVEXCROSS( $e, e'$ )
12:     $f_m \leftarrow \min(f_t(e), f_t(e'))$ 
13:     $f_t(e') \leftarrow f_t(e') - f_m$ 
14:     $f_t(e) \leftarrow f_t(e) - f_m$ 
15:     $(V_t[x, y_1], V_t[x+1, y_2]) \leftarrow e$ 
16:     $(V_t[x, y'_1], V_t[x+1, y'_2]) \leftarrow e'$ 
17:     $f_t((V_t[x, y_1], V_t[x+1, y'_2])) \leftarrow f_m$ 
18:     $f_t((V_t[x, y'_1], V_t[x+1, y_2])) \leftarrow f_m$ 

```

LEMMA 3.25. *The running time of the function RESOLVEXCROSSES in the Algorithm 2 is polynomial.*

The proof scheme of Lemma 3.25 is to show the minimum crossing flow at step i is strictly less than the minimum crossing flow at step $i+1$ (after the previously minimum crossing flow is resolved). Consequently, since the total number of possible crossing edges of a day graph is polynomial, the number of steps until the algorithm halts is polynomial. Again, we left the formal proof of this lemma to the appendix for space limitations.

Note that another property of Algorithm 2 is that the flow passing through a vertex will not change and it is only the amount of flow passing through the edges that changes (Figure 5). This proves most of the constraints of LP 1 will still hold. The only two constraints that consider the flow passing through the edges, and not the vertices, are number 7 and number 8. The former will be true since the process does not produce any negative flow and the latter is true since by Lemma 3.24 the total cost does not change.

Consequently, the following statement is true since we can first solve LP 1 by any polynomial time LP-solver and then run Algorithm 2 on its solution.

COROLLARY 3.26. *There exists a polynomial time algorithm that finds $f \langle f_1, \dots, f_T \rangle$, a collection of canonical flows, that is a solution of LP 1, and for any $t \in [T]$, f_t does not contain any crossing flow.*

3.3.3 Constructing A Strategy. Assuming f is a non-crossing solution of LP 1, this section gives an algorithm to find a mixed strategy of the defender that equivalent to the set of canonical flows in f and finally proves Theorem 3.1. We first define what we mean by the top most flow path of a non-crossing canonical flow:

Definition 3.27 (Top-Most Flow Path). Let f_t be a canonical flow of G_t without any crossing flows. We say a canonical path $p = \langle e_1, e_2, \dots, e_{K+1} \rangle$ of G_t is a flow path of f_t if for

any k ($1 \leq k \leq K+1$), $f_t(e_k) > 0$. The top-most flow path of f_t is the flow path of f_t that is above all other flow paths of f_t (that is well-defined because f_t does not have any crossing flow). The size of the flow path p of f_t , denoted by $|p|$, is m if for any k , $f_t(e_k) \geq m$ and there exists an edge e_i of p such that $f_t(e_i) = m$.

LEMMA 3.28. *Let $f = \langle f_1, \dots, f_T \rangle$ be a collection of non-crossing canonical flows that satisfies the compatibility constraint (constraint number 9) of LP 1. For any t that $\{t, t+1\} \subset [T]$, the top-most flow paths of f_t and f_{t+1} are compatible.*

Proof. Let $\langle S_t, V_t[1, y_1], \dots, V_t[K, y_K], S'_t \rangle$ and $\langle S_{t+1}, V_{t+1}[1, y'_1], \dots, V_{t+1}[K, y'_K], S'_{t+1} \rangle$ respectively denote the top-most flow paths of f_t and f_{t+1} . It suffices to prove for any $k \in [K]$, there is a valid movement from the corresponding interval of $V_t[k, y_k]$ to the corresponding interval of $V_{t+1}[k, y'_k]$. To do so, we assume this is not the case and obtain a contradiction. Let $V_{t+1}[k, y'_k] \notin \text{feas}_t(V_t[k, y_k])$ for some $k \in [K]$, then one of the following conditions should hold:

- (1) $V_{t+1}[k, y'_k]$ is below the feasible range $\text{feas}_t(V_t[k, y_k])$.
- (2) $V_{t+1}[k, y'_k]$ is above the feasible range $\text{feas}_t(V_t[k, y_k])$.

If the first condition is true, the contradiction is that $V_{t+1}[k, y'_k]$ cannot be in the top-most flow path of f_{t+1} . To see this, note that we know by constraint 9 of LP 1 that the total flow passing through the vertices of $\text{feas}_t(V_t[k, y_k])$ is not less than the flow passing through $V_t[k, y_k]$, therefore there is a vertex in $\text{feas}_t(V_t[k, y_k])$ (and above $V_{t+1}[k, y'_k]$) with a non-negative flow and thus $V_{t+1}[k, y'_k]$ cannot be in the top-most flow path of f_{t+1} .

If the second condition is true, the contradiction is that constraint 9 cannot be satisfied. To see this, note that since $V_t[k, y_k]$ is in the top-most flow path of f_t , no flow passes through the vertices above it and therefore $\sum_{v \in V_t[k; 1:y]} f_t^+(v) = 1$. However, since $V_{t+1}[k, y'_k]$ is above the feasible range $\text{feas}_t(V_t[k, y_k])$, $\sum_{v' \in \text{feas}_t(V_t[k; 1:y])} f_{t+1}^+(v') < 1$ which means constraint 9 that indicates the value of the latter summation should not be less than the former one, cannot be satisfied. $\square$

THEOREM 3.29. *Let $f = \langle f_1, \dots, f_T \rangle$ be a solution of LP 1 without any crossing flows. There exists a polynomial time algorithm to find a mixed strategy of the defender that is equivalent to f .*

To prove Theorem 3.29, we show the following iterative algorithm constructs the desired mixed strategy in polynomial time:

- (1) Find the top-most flow paths $p_1, \dots, p_T$ of $f_1, \dots, f_T$.
- (2) Construct the pure strategy p , corresponding to $p_1, \dots, p_T$.
- (3) Add p to s with probability $q = \min |p_i|$.
- (4) For any edge e of any p_i , decrease $f_i(e)$ to $f_i(e) - q$.
- (5) If there is any flow left in $f_1, \dots, f_T$, repeat all the steps.

Note that by Lemma 3.28, if the compatibility constraint of LP 1 is satisfied, the top-most flow paths are compatible. Since at the first round of the algorithm, we have an actual solution of the LP, the compatibility constraint is obviously satisfied. To completely prove the correctness of this algorithm, we also need to show after each iteration, changing the flows does not violate the compatibility constraint. For space limitations, we left this part of the proof to the appendix. Furthermore, the running time of this algorithm is polynomial in the input size since in each iteration, the flow passing through at least one edge decreases to zero and the total number of edges is polynomial.

We can now prove the main theorem of Section 3.

THEOREM (3.1). *There is a polynomial time (in input size) algorithm to solve $\mathcal{DSG}$.*

Proof. Recall that by Lemma 3.21, the optimal solution of LP 1 is a lower bound for the utility of the attacker when both players play their optimal (minimax) strategies. Also note that by Lemma 3.26 and Theorem 3.29, we can construct a mixed strategy s of the defender that is equivalent to the optimal solution of LP 1 in polynomial time. This means the maximum utility of the attacker when the defender plays s , is equal to its lower bound and therefore s minimizes the maximum expected utility of the attacker: i.e., it is a minimax strategy of the defender. $\square$

4 CONTINUOUS MODEL

In this section we prove the following theorem for the continuous model:

THEOREM 4.1. *There exists a polynomial time algorithm to find an optimal solution for $\mathcal{CSG}$.*

The given proof is based on a technical assumption that all numbers in the input are rational.

Proof of Theorem 4.1: The main idea of this proof is to reduce any instance of $\mathcal{CSG}$ to an instance of $\mathcal{DSG}$, for which we know there exists a polynomial time algorithm.

Recall that any rational number can be represented by a fraction $\frac{a}{b}$ such that both a and b are integers. Let B be the set of denominators in this fractional representation of all numbers in the input. We define m to be the product of all numbers in B . Note that the number of digits needed to represent m is polynomial in the input size since every number in B appears in the input. To create an instance of $\mathcal{DSG}$, we multiply all target positions, Δ , R and M , given in the instance of $\mathcal{CSG}$ to m .

To use the algorithm for $\mathcal{DSG}$, it suffices to prove in the scaled solutions, there exists an optimal solution that places patrols only in the integer locations. To do this, we prove that for any given patrol path $p_1 = \langle m_i \rangle_T$ in the scaled version, there exists a patrol path $p_2 = \langle m'_i \rangle_T$ that covers the same set of targets and for any $t \in [T]$, m'_t is an integer position. It suffices to set m'_t to be $\lfloor m_t \rfloor$. Note that since the position of all targets and the protecting ranges of the patrols in the scaled version are all integers, this patrol path protects exactly the same set of targets. Now we can use the algorithm of Theorem 3.1 to find the optimal solution of this scaled input and then scale it back to the original size by dividing the patrols' locations by m . $\square$

REFERENCES

- [1] AmirMahdi Ahmadinejad, Sina Dehghani, MohammadTaghi Hajiaghayi, Brendan Lucier, Hamid Mahini, and Saeed Seddighin. 2016. From duels to battefields: Computing equilibria of blotto and other games. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*.
- [2] Maria-Florina Balcan, Avrim Blum, Nika Haghtalab, and Ariel D. Procaccia. 2015. Commitment Without Regrets: Online Learning in Stackelberg Security Games. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC*. 61–78.
- [3] Soheil Behnezhad, Sina Dehghani, Mahsa Derakhshan, MohammadTaghi HajiAghayi, and Saeed Seddighin. 2017. Faster and Simpler Algorithm for Optimal Strategies of Blotto Game. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*.
- [4] Avrim Blum, Nika Haghtalab, and Ariel D. Procaccia. 2014. Learning Optimal Commitment to Overcome Insecurity. In *Annual Conference on Neural Information Processing Systems (NIPS)*. 1826–1834.

- [5] Branislav Bošanský, Viliam Lisý, Michal Jakob, and Michal Pěchouček. 2011. Computing time-dependent policies for patrolling games with mobile targets. In *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 3*. International Foundation for Autonomous Agents and Multiagent Systems, 989–996.
- [6] Matthew Brown, Arunesh Sinha, Aaron Schlenker, and Milind Tambe. 2016. One Size Does Not Fit All: A Game-Theoretic Approach for Dynamically and Effectively Screening for Threats. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*. 425–431.
- [7] Vincent Conitzer and Tuomas Sandholm. 2006. Computing the optimal strategy to commit to. In *Proceedings 7th ACM Conference on Electronic Commerce (EC-2006)*. 82–90.
- [8] Fei Fang, Albert Xin Jiang, and Milind Tambe. 2013. Optimal patrol strategy for protecting moving targets with multiple mobile resources. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*. International Foundation for Autonomous Agents and Multiagent Systems, 957–964.
- [9] Fei Fang, Thanh Hong Nguyen, Rob Pickles, Wai Y. Lam, Gopalasamy R. Clements, Bo An, Amandeep Singh, Milind Tambe, and Andrew Lemieux. 2016. Deploying PAWS: Field Optimization of the Protection Assistant for Wildlife Security. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*. 3966–3973.
- [10] Christopher Kiekintveld, Manish Jain, Jason Tsai, James Pita, Fernando Ordóñez, and Milind Tambe. 2009. Computing optimal randomized resource allocations for massive security games. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, 689–696.
- [11] Dmytro Korzhuk, Vincent Conitzer, and Ronald Parr. 2010. Complexity of Computing Optimal Stackelberg Strategies in Security Resource Allocation Games.. In *AAAI*.
- [12] Joshua Letchford and Vincent Conitzer. 2013. Solving Security Games on Graphs via Marginal Probabilities. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence*.
- [13] Joshua Letchford, Vincent Conitzer, and Kamesh Munagala. 2009. Learning and Approximating the Optimal Strategy to Commit To. In *Algorithmic Game Theory, Second International Symposium, SAGT*. 250–262.
- [14] Janusz Marecki, Gerald Tesauro, and Richard Segal. 2012. Playing repeated Stackelberg games with unknown opponents. In *International Conference on Autonomous Agents and Multiagent Systems, AAMAS*. 821–828.
- [15] Giuseppe De Nittis and Francesco Trovò. 2016. Machine Learning Techniques for Stackelberg Security Games: a Survey. *Technical report on arxiv.org* abs/1609.09341 (2016).
- [16] James Pita, Manish Jain, Janusz Marecki, Fernando Ordóñez, Christopher Portway, Milind Tambe, Craig Western, Praveen Paruchuri, and Sarit Kraus. 2008. Deployed ARMOR protection: the application of a game theoretic model for security at the Los Angeles International Airport. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems: industrial track*. International Foundation for Autonomous Agents and Multiagent Systems, 125–132.
- [17] Tim Roughgarden. 2010. Computing Equilibria: A Computational Complexity Perspective. *Economic Theory* 42, 1 (2010), 193–236.
- [18] Milind Tambe. 2011. *Security and game theory: algorithms, deployed systems, lessons learned*. Cambridge University Press.
- [19] Haifeng Xu. 2016. The Mysteries of Security Games: Equilibrium Computation Becomes Combinatorial Algorithm Design. In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC*. 497–514.
- [20] Haifeng Xu, Fei Fang, Albert Xin Jiang, Vincent Conitzer, Shaddin Dughmi, and Milind Tambe. 2014. Solving Zero-Sum Security Games in Discretized Spatio-Temporal Domains.. In *AAAI*. Citeseer, 1500–1506.
- [21] Zhengyu Yin, Albert Xin Jiang, Milind Tambe, Christopher Kiekintveld, Kevin Leyton-Brown, Tuomas Sandholm, and John P Sullivan. 2012. TRUSTS: Scheduling randomized patrols for fare inspection in transit systems using game theory. *AI Magazine* 33, 4 (2012), 59.