

## Urban Security: Game-Theoretic Resource Allocation in Networked Physical Domains

Jason Tsai, Zhengyu Yin, Jun-young Kwak, David Kempe, Christopher Kiekintveld, Milind Tambe

University of Southern California, Los Angeles, CA 90089  
{jasonetts, zhengyuy, junyoungk, dkempe, kiekintv, tambe} @usc.edu

### Abstract

Law enforcement agencies frequently must allocate limited resources to protect targets embedded in a network, such as important buildings in a city road network. Since intelligent attackers may observe and exploit patterns in the allocation, it is crucial that the allocations be randomized. We cast this problem as an attacker-defender Stackelberg game: the defender's goal is to obtain an optimal mixed strategy for allocating resources. The defender's strategy space is exponential in the number of resources, and the attacker's exponential in the network size. Existing algorithms are therefore useless for all but the smallest networks.

We present a solution approach based on two key ideas: (i) A polynomial-sized game model obtained via an approximation of the strategy space, solved efficiently using a linear program; (ii) Two efficient techniques that map solutions from the approximate game to the original, with proofs of correctness under certain assumptions. We present in-depth experimental results, including an evaluation on part of the Mumbai road network.

### Introduction

Protecting targets against potential attacks is an important problem for security forces worldwide. The general setting is as follows: An attacker assigns different values to reaching (and damaging or destroying) one of multiple targets. A defender wants to allocate resources (such as patrol cars or canine units) to capture the attacker before he reaches a target. In many of these situations, the domain has structure that is naturally modeled as a graph. For example, city maps can be modeled with intersections as nodes and roads as edges, where nodes are targets for attackers. In order to prevent attacks, security forces can schedule checkpoints on edges (e.g., roads) to detect intruders. For instance, in response to the devastating terrorist attacks in 2008 (Chandran and Beitchman 29 November 2008), Mumbai police deploy randomized checkpoints as one countermeasure to prevent future attacks (Ali 4 August 2009). The strategy for placing these checkpoints must necessarily be decided in advance of attack attempts, should account for targets of differing importance, and should anticipate an intelligent adversary who can observe the strategy prior to attacking.

Copyright © 2010, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

In light of these requirements, game-theoretic approaches have been developed to assist in generating randomized security strategies in real-world domains, including applications in use by the Los Angeles International Airport (Pita et al. 2008) and the Federal Air Marshals Service (Tsai et al. 2009). To account for the attacker's ability to observe deployment patterns, these methods model the problem as a Stackelberg game and solve for an optimal probability distribution over the possible deployments to ensure unpredictability. Novel solvers for such security games have recently been developed (Basilico, Gatti, and Amigoni 2009; Paruchuri et al. 2008; Conitzer and Sandholm 2006), but all take time at least polynomial in the number of actions of both players. In our setting, every path from an entry point to a target is an attacker action, and every set of  $r$  or fewer edges is a defender action, where  $r$  is the maximum number of checkpoints. Since the attacker's actions grow exponentially with the size of the network, and the defender's actions grow exponentially with  $r$ , existing methods quickly become too slow when applied to large real-world domains.

In this work, we develop an efficient procedure for generating checkpoint deployments based on two key ideas: (i) a polynomial-sized approximation of the strategy space solved using a linear program; (ii) two efficient sampling techniques to map solutions back to the original space. To avoid the exponential strategy space over all possible combinations of checkpoint placements (the joint distribution), our methods operate on the marginal probabilities of edges, i.e., the total probability of placing a checkpoint on an edge. Our linear program, RANGER, upper-bounds the capture probabilities along paths by the sum of marginal probabilities.

Our sampling algorithms efficiently generate joint distributions in the original problem space from RANGER's solution of marginal probabilities. We prove that under certain conditions, the actual capture probabilities of our algorithms match the upper bounds of RANGER, and thus necessarily give *optimal* payoff. *Radius Sampling* generates optimal joint distributions if certain conditions on the marginal distribution are met. *Comb Sampling* generates distributions which are optimal against an *approximating attacker* who calculates the expected value of an attack by summing the marginal probabilities in the path.

In addition to our theoretical results, we test our methods empirically. First, we evaluate the quality of RANGER

against an optimal solution technique, DOBSS, to verify the accuracy of RANGER’s approximation. Then, we evaluate the sampling procedures by testing against an *exact attacker* who plays a best response to the defender’s true joint distribution. We also apply our methods to a game model of the city of Mumbai and the targets attacked in 2008.

## Related Work

Aside from the literature on Stackelberg games for security, our approach is also based on insights from network interdiction (Washburn and Wood 1995; Israeli and Wood 2002). These are the special case of our model when there is a single target, or — equivalently — all targets have identical values. For such games, Washburn and Wood (1995) give an algorithm finding optimal strategies for both players based on Min-Cut computations. However, different target values can cause their algorithm to perform arbitrarily poorly, as we see in our experiments.

Two additional lines of work are somewhat related. Mavronicolas et al. (2008) define and analyze a network security game where each attacker can attack any node of the network, and the defender chooses a path to patrol to capture as many attackers as possible. Because the attacker is not restricted to paths, the types of results for this game are different from ours, and the focus in (Mavronicolas et al. 2008) is on understanding the impact of selfish behavior by defenders rather than optimal strategies. Hider-seeker games (Basilico, Gatti, and Amigoni 2009; Halvorson, Conitzer, and Parr 2009) are also studied on graphs, but here, the attacker’s goal is only to evade capture, not to reach any particular target.

## Problem Description

A graph-based security game models an attacker and a defender who take actions on a graph  $G = (V, E)$ , with  $n = |V|$  nodes and  $m = |E|$  edges. The attacker starts at one of the source nodes  $s \in S \subseteq V$  of his choosing and travels along a path in an attempt to reach one of the targets  $t \in T \subseteq V$ . The attacker’s pure strategies are thus all  $s$ - $t$  paths  $P$ , denoted by  $\mathcal{B}$ , from some source  $s$  to some target  $t$ . The defender tries to capture the attacker before he reaches a target, by placing up to  $r$  resources on *edges* of the graph. The defender’s pure strategies are subsets of  $r$  or fewer edges; we denote the set of all such sets by  $\mathcal{L}$ . Assuming that the defender plays  $L \in \mathcal{L}$  and the attacker  $P \in \mathcal{B}$ , the attacker is captured whenever  $P \cap L \neq \emptyset$ , and succeeds in his attack when  $P \cap L = \emptyset$ .

Unsuccessful attacks always have a payoff of  $c$  for the defender, while successful ones have a penalty of  $D(t)$ . We make the natural restriction that  $D(t) \leq c$ . We also assume that the game is zero-sum, meaning that the attacker’s payoff for a successful attack on target  $t$  is  $-D(t)$ , and  $-c$  for an unsuccessful one. We stress here that targets may have vastly different payoffs associated with them, unlike in (Washburn and Wood 1995). This distinction is crucial to model real security domains, and thus to bridge the gap between theory and practice.

Since the attacker can choose which source to enter from, for our analysis, we merge all sources into a *single source* without loss of generality. More formally, we reform the graph so that all original source-incident edges are incident to the new source. While this operation obviously changes the graph, it does so only in a way that does not impact the game: no rational attacker would ever include multiple sources in his path, and therefore, a defender will never select an edge between two sources. Those are the only edges that disappear from the problem. Thus, to simplify presentation and analysis, we will assume that the attacker always enters from a unique known source  $s$ .

In a world of increasingly sophisticated and determined attackers, a good defender strategy must take into account the fact that the attacker will observe and exploit patterns in the defender’s behavior. Thus, the game is naturally modeled as a Stackelberg game, an approach also taken (for the same reasons) in past work in security settings (Gatti 2008; Kiekintveld et al. 2009). The defender is modeled as the *leader* and moves first, by selecting a mixed strategy  $\lambda \in \Lambda$  that assigns a probability to each pure strategy  $L \in \mathcal{L}$ . The attacker is the *follower* and chooses a strategy after observing the defender’s mixed strategy. There is always a pure-strategy best response for the attacker, so we restrict the attacker to pure strategies without loss of generality. Thus, the attacker’s Stackelberg strategy is a function  $f : \lambda \mapsto P$ . For any pair of strategy profiles  $(\lambda, f)$ , the expected rewards for the defender ( $R_D$ ) and attacker ( $R_A$ ) are given by:

$$R_D(\lambda, f) = p \cdot c + (1 - p) \cdot D(t) \quad (1)$$

$$R_A(\lambda, f) = p \cdot -c + (1 - p) \cdot -D(t) \quad (2)$$

where  $t$  is the target at the end of the path specified by  $f(\lambda)$ , and  $p$  the probability that the attacker is captured on the path to  $t$  given the defender’s strategy  $\lambda$ . Although the optimal defender strategy is a Stackelberg Equilibrium, in a zero-sum this is equivalent to a Maximin strategy. Unfortunately, as  $\mathcal{L}$  has size  $\Theta(m^r)$ , and  $\mathcal{B}$  has size exponential in  $n$ , existing methods for computing such strategies do not scale to realistic problem sizes. We therefore develop a linear program and two accompanying sampling methods to efficiently solve graph-based security games.

## RANGER

We first introduce RANGER (Resource Allocation for Network-based Games with Efficient Representations), a linear program for finding an optimal set of *marginal* check-point probabilities for the defender. We denote the marginal probability associated with edge  $e$  by  $x_e$ . Formally,  $x_e = \sum_{L \in \mathcal{L}, e \in L} \lambda_L$ , where  $\lambda_L$  is the probability of the set  $L$  under  $\lambda$ . We denote the marginal distribution by  $\vec{x} = \langle x_e \rangle$ .

By reasoning over  $\vec{x}$ , we avoid the exponential size of the defender’s space. The key insight of our approach is the following simple consequence of the Union Bound: *For any path  $P$ , the capture probability under  $\lambda$  is at most  $\sum_{e \in P} x_e$ .* We use this upper bound (the sum of  $x_e$ ) as an approximation of the true capture probability in deriving RANGER. The power of our approach is that we subsequently present ways to sample joint distributions where the total capture

probability *matches* this upper bound under certain conditions; this immediately implies optimality of our procedures, and retroactively justifies the approximation.

In the RANGER linear program below,  $x_e$  is the marginal probability of placing a checkpoint on edge  $e$ . The  $d_v$  are, for each vertex  $v$ , the minimum sum of checkpoint probabilities along any path from the source,  $s$  to  $v$ <sup>1</sup>. This is enforced by the constraints (4)–(6). Constraint (7) enforces that at most  $r$  checkpoints are placed, and Constraint (3) captures the payoff for the defender.

$$\text{Maximize } R_D, \text{ s.t.:$$

$$R_D \leq (1 - d_t) \cdot D(t) + d_t \cdot c \quad (3)$$

$$d_s = 0 \quad (4)$$

$$d_v \leq \min(1, d_u + x_e) \quad \forall e = (u, v) \quad (5)$$

$$0 \leq x_e \leq 1 \quad \forall e \in E \quad (6)$$

$$\sum_{e \in E} x_e \leq r \quad (7)$$

We begin by verifying the claim that RANGER’s solution is an overestimate of the optimal solution.

**Theorem 1.** *Let  $\lambda^*$  be the optimal strategy and  $R^*$  the corresponding defender reward. Then,  $R_D^* \geq R^*$ , where  $R_D^*$  is the defender reward returned by the LP.*

*Proof.* Let  $\vec{x}^*$  be the marginal probabilities of  $\lambda^*$ . Obviously,  $0 \leq x_e^* \leq 1$ , and  $\sum_{e \in E} x_e^* = \sum_{L \in \mathcal{L}} |L| \cdot \lambda_L^* \leq r$ . For each vertex  $v$ , let  $d_v^*$  be the probability (under  $\lambda^*$ ) of capturing the intruder assuming he chooses the best path to reach  $v$ . Then  $0 \leq d_v^* \leq 1$ , and for each edge  $e = (u, v)$ ,  $d_v^* \leq d_u^* + x_e^*$ , by the Union Bound. The attacker will choose the path to maximize his own payoff  $R_A^*$ . Because the game is zero-sum,

$$R^* = -R_A^* = -\max_t \{(1 - d_t^*) \cdot -D(t) + d_t^* \cdot -c\}$$

$$= \min_t \{(1 - d_t^*) \cdot D(t) + d_t^* \cdot c\}.$$

Thus, for any target  $t$ ,  $R^* \leq (1 - d_t^*) \cdot D(t) + d_t^* \cdot c$ . Thus, the values  $R^*$ ,  $\vec{d}^*$  and  $\vec{x}^*$  are feasible for the LP; because RANGER finds the optimum feasible solution, we obtain that  $R_D^* \geq R^*$ .  $\square$

RANGER is an exponentially more compact representation of both the attacker and defender strategy spaces. This can be seen by noticing that RANGER has a polynomial number of variables with respect to  $n$  and  $m$ . Any straightforward application of prior formulations would have  $\Theta(m^r)$  variables for the defender and exponentially many ( $|\mathcal{B}|$ ) constraints for the attacker.

## Constructing Joint Distributions

To deploy security resources, we require joint schedules, drawn from a joint distribution over  $\mathcal{L}$ . We develop sampling procedures that use the  $\vec{x}$  computed by RANGER to generate a distribution over joint schedules. The principle

<sup>1</sup>Recall that we can assume a single source w.l.o.g.

behind these methods is to bring the *actual* capture probability for a target  $t$  up to the value  $d_t$ . One way to ensure this would be if no deployment ever placed two checkpoints on any  $s$ - $t$  path to any target. More generally (and informally), it is sufficient if this “checkpoint disjointness” is ensured for “critical”  $s$ - $t$  paths: those whose sum of marginal probabilities are close to the minimal ( $d_t$ ).

Notice that placing checkpoints by independently sampling from  $\vec{x}$  violates this approach with increasing frequency as  $r$  increases, and yields very suboptimal solutions. Instead, we introduce two novel sampling procedures that achieve a certain “checkpoint disjointness”, under some assumptions, and are therefore *optimal*.

## Radius Sampling

Radius Sampling (RS) interprets the marginal probabilities as “distances”, and places a security ring around the source. In this way, it avoids sampling multiple times on “critical paths”, in a way we make precise now. For any  $h \geq 0$ , we define the *ring of radius  $h$*  around  $s$  as  $R_h := \{e = (u, v) | d_u \leq h < d_v\}$ , i.e., the set of edges from a node with probability of capture at most  $h$  from  $s$  to a node with probability of capture more than  $h$  from  $s$ .

We define  $\alpha := \int_0^\infty |R_h| dh$  (a normalization constant), and the density function  $\phi(h) := \frac{|R_h|}{\alpha}$ . Notice that

$$\alpha = \sum_{e=(u,v)} (d_v - d_u) \leq \sum_e x_e \leq r. \quad (8)$$

Our algorithm works as follows: Choose a radius  $h$  from  $[0, \infty]$  according to the density function  $\phi$ . Now, choose  $r$  of the edges in  $R_h$  uniformly at random (or all edges in  $R_h$  if  $|R_h| \leq r$ ). Place checkpoints on these edges. We call the resulting set  $L_R$ . Notice that both  $h$  and  $L_R$  are random variables, and  $L_R$  is a set of at most  $r$  edges.

**Theorem 2.** *If for all  $h$ ,  $|R_h| \geq r$  or  $|R_h| = 0$ , then RS produces an optimal distribution for the defender.*

Theorem 2 follows from Lemma 3 and Theorem 1 as follows: By Lemma 3, the capture probability for any  $s$ - $t$  path is at least  $d_t$ , i.e., RANGER’s value. Therefore, the defender’s payoff is at least RANGER’s, which by Theorem 1 is at least the payoff with the optimum mixed strategy.

**Lemma 3.** *Under the assumptions of Theorem 2, let  $P$  be any  $s$ - $v$  path and  $w$  the node maximizing  $d_w$  among all nodes on  $P$ . The capture probability along  $P$  is at least  $d_w$ .*

*Proof.* We prove the lemma by induction on  $|P|$ , the number of edges on path  $P$ . In the base case  $|P| = 0$ , the only node  $v$  with a path from  $s$  is  $s$  itself, and the statement holds.

For the inductive step, let  $P$  be a path of length  $\ell + 1$  and  $e = (v', v)$  the last edge of  $P$ . Let  $P' = P \setminus \{e\}$  be the path of length  $\ell$  from  $s$  to  $v'$ , and  $w'$  the node on  $P'$  maximizing  $d_{w'}$ . By Induction Hypothesis,  $\text{Prob}[L_R \cap P' \neq \emptyset] \geq d_{w'}$ .

We distinguish two cases. If  $d_{w'} \geq d_v$ , then

$$\text{Prob}[L_R \cap P \neq \emptyset] \geq \text{Prob}[L_R \cap P' \neq \emptyset]$$

$$\geq d_{w'} \geq d_v,$$

implying the claim.

If  $d_v > d_{w'}$ , then consider the event  $\mathcal{E} = [h > d_{w'} \text{ and } e \in L_R]$ .  $\mathcal{E}$  is the event when we include  $e$  in  $L_R$

and  $h$  is sufficiently large that no edge from  $P'$  can also be sampled. The probability of  $\mathcal{E}$  is

$$\begin{aligned} & \int_{d_{w'}}^{d_v} \text{Prob}[e \in L_R \mid h = x] \phi(x) dx \\ &= \int_{d_{w'}}^{d_v} \frac{r}{|R_x|} \cdot \frac{|R_x|}{\alpha} dx = \int_{d_{w'}}^{d_v} \frac{r}{\alpha} dx \geq d_v - d_{w'}. \end{aligned}$$

Here, we substituted the definitions of the sampling process, and then used that  $\frac{r}{\alpha} \geq 1$  from Equation (8).

Whenever  $L_R$  intersects  $P'$ , by definition, we must have that  $h \leq d_{w'}$  (because no edge  $e' \in P'$  is in  $R_h$  for  $h > d_{w'}$ ). Thus, the events  $\mathcal{E}$  and  $[R_h \cap P' \neq \emptyset]$  are disjoint, and

$$\begin{aligned} \text{Prob}[R_h \cap P \neq \emptyset] &\geq \text{Prob}[(R_h \cap P' \neq \emptyset) \cup \mathcal{E}] \\ &= \text{Prob}[R_h \cap P' \neq \emptyset] + \text{Prob}[\mathcal{E}] \\ &\geq d_{w'} + (d_v - d_{w'}) \\ &= d_v. \end{aligned}$$

The penultimate step used the induction hypothesis as well as the inequality  $\text{Prob}[\mathcal{E}] \geq d_v - d_{w'}$  derived above.  $\square$

To implement RS, we need to find  $d_v$  values to each node, determine the edges and weights for each  $R_h$ , and sample according to the above procedure. Each step takes time polynomial in  $n$ . Thus, we have a polynomial-time procedure that *optimally* solves a graph-based security game under the conditions of Theorem 2. Previously known techniques either required time exponential in the graph size or can not provide quality guarantees. However, since we cannot guarantee performance when Radius Sampling's condition is not met, we also explore another sampling algorithm.

## Comb Sampling

Now, we consider a somewhat simpler case for the defender: the attacker only observes marginal distributions and approximates the capture probability on any path by adding the probabilities. This may occur because observing the full joint probability distribution is much more time- and resource-intensive than observing the marginals. When only able to observe marginals, adding probabilities is a reasonable and conservative approximation for the attacker.

Comb Sampling (CS) is based on two ideas: (1) If the marginals of the joint distribution match RANGER's  $x_e$  values, then an attacker summing probabilities will choose the target  $t$  and path  $P$  predicted by RANGER. (2) If the edges on  $P$  are chosen mutually exclusively, then the capture probability on  $P$  matches that of RANGER.

Let  $e_1, \dots, e_{|P|}$  be the edges on the path  $P$  (in arbitrary order), and  $e_{|P|+1}, \dots, e_m$  the remaining edges, in arbitrary order. For each  $1 \leq j \leq m$ , let  $X_j = \sum_{i < j} x_i$ , and define the interval  $I_j = [X_j, X_j + x_j)$ . Because  $\sum_i x_i = r$  (w.l.o.g.), the  $I_j$  form a disjoint cover of the interval  $[0, r)$ . We now generate a deployment,  $L_C$ , as follows: Pick a number  $y \in [0, 1)$  uniformly at random, and include in  $L_C$  all edges  $e_j$  such that  $y + k \in I_j$  for some integer  $k$ . In other words, include exactly the edges which "own" the intervals containing the points  $y, y + 1, y + 2, \dots, y + r - 1$ . This samples exactly  $r$  edges.

**Lemma 4.** *Given a marginal distribution  $\vec{x}$ , CS will exactly meet all marginal probabilities,  $x_e$ .*

*Proof.* Consider any edge  $e_j$ , and two cases. If  $I_j \subseteq [k, k + 1)$  for some  $k$  (i.e.,  $I_j$  contains no integer point), then  $e_j$  is included if and only if  $k + y \in I_j$ , which happens with probability  $|I_j| = x_j$ . On the other hand, if  $I_j = [X_j, k) \cup [k, X_j + x_j)$ , then  $e_j$  is included if and only if  $y + k - 1 \in [X_j, k)$  or  $y + k \in [k, X_j + x_j)$ ; because  $x_k \leq 1$ , this happens with probability  $(k - X_j) + (X_j + x_j - k) = x_j$ .  $\square$

Lemma 4 ensures that the attacker will follow the path predicted by RANGER. Now consider  $P$ . If  $\sum_{e \in P} x_e \geq 1$ , then for any  $y$ , some edge  $e \in P$  will be included, so the attacker is always captured. Otherwise, an edge from  $P$  is included if and only if  $y < \sum_{e \in P} x_e$ , which happens with probability  $\sum_{e \in P} x_e$ , i.e., the sum of marginals on  $P$ . Combined with Lemma 4, this guarantees that RANGER's reward is achievable using CS if the defender faces an approximating attacker. The sampling time is clearly polynomial in the graph size. Thus, we have a polynomial-time procedure to optimally defend against an approximating attacker.

## Experiments

### Quality Comparison

Our first evaluation studies the solution quality generated by Radius and Comb Sampling in the general case, against both exact and approximating attackers, as defined in the Introduction. Neither method is guaranteed to achieve the optimal value against an exact attacker, so we evaluate their approximation quality. We compare them against DOBSS, which computes the optimal solution against an exact attacker. DOBSS may not be optimal against an approximating attacker, so we also report the quality of DOBSS against approximating attackers, labeled DOBSS Marginal. As a benchmark, we include a simple Independent Sampling (IS) strategy, wherein for each checkpoint, edge  $e$  is selected independently with probability  $\frac{x_e}{r}$ .

We generate random graphs that are representative of the domains of interest. First, we test on random geometric graphs to estimate performance for road network domains. Then we test on scale-free graphs as an analogy for subway networks. For each graph type, we generate 500 instances each of 4, 5, and 6-node, zero-sum games and report results in Table 1. Every graph has one source, between 1 and  $|V| - 1$  targets with values from -1 to -10 for the defender when successfully attacked, and 1 to 3 checkpoints. The payoffs for capture are all 0. *Graphs were kept simple so DOBSS could solve them within a reasonable time limit.*

For each graph, we run each sampling method on the marginals produced by RANGER and calculate the *actual* capture probabilities for the joint distribution generated. Using the true capture probabilities, we select an optimal target for the attacker and compute expected payoffs based on this choice. We compare these rewards against DOBSS to evaluate the quality of our methodology.

Table 1 shows the number of cases where a difference in quality of more than 0.001 existed between methods. Empirically, RANGER computes very good estimates of the optimal reward value, never differing from DOBSS for more than 5% of cases. Unsurprisingly, Independent Sampling

frequently results in suboptimal distributions (44%–78%). Remarkably, CS attains the optimal expected reward in every single one of the 3,000 graphs tested. RS also performs very well, never differing from DOBSS for more than 11% of the games. DOBSS Marginal never differs from DOBSS or RANGER, indicating that DOBSS remains optimal against an approximating attacker (not shown). However, as runtime experiments will show, DOBSS is completely incapable of solving reasonable game sizes.

	Random Geo.			Scale-Free		
	4	5	6	4	5	6
RG > DOBSS	12	8	5	0	4	22
IS < DOBSS	220	283	280	389	347	247
CS < DOBSS	0	0	0	0	0	0
RS < DOBSS	0	0	3	0	29	53

Table 1: Results by number of cases (RG - RANGER).

## Mumbai

As a real-world trial, we use our algorithms to create security policies for the southern tip of Mumbai, shown in Figure 1, which was an area of heavy terrorist activity in 2008. The region is modeled as a graph with 35 nodes and 58 edges. Attackers can potentially enter from *any* entry node, chosen based on historical and likely entry points. A total of four target nodes are chosen based on historical attacks, marked with black circles in Figure 1. These are held constant throughout testing. Payoffs are decided as in the Quality Comparison experiments.

Figure 2a shows the averaged defender rewards obtained across eight configurations, each with their own setup of target values and sources, with each configuration being run with checkpoints varying from 2 to 10. Figure 2b shows results averaged across a different set of eight configurations, each with their own setup of target values and 4 checkpoints, with each configuration being run with the number of sources increasing from 1 to 7.






Figure 1: Target layout for southern Mumbai.



(a) Avg reward by checkpoints. (b) Avg reward by sources.

Figure 2: Rewards in Mumbai domain.



(a) No. edges used. (b) Standard dev. of Prob.

Figure 3: RANGER strategies in Mumbai domain.

Since DOBSS is unable to solve even the simplest case within the 20-minute limit, we include only CS, RS, Minimum Cut (Mincut), as well as three natural defense strategies: *Uniform Random* (Uniform), *Entry-Adjacent* (Entry-Adj) and *Weighted-Target-Adjacent* (Target-Adj). MC, as introduced by (Washburn and Wood 1995), contracts all sources into a super-source and all targets into a super-target and finds the minimum cut on the resulting graph, uniformly randomizing resources across it. This effectively ignores target value variation, but is extremely efficient. Uniform randomizes checkpoints uniformly across all edges in the graph. Entry-Adj places checkpoints on edges incident to entry nodes with equal probability. Target-Adj places checkpoints on edges incident to target nodes, weighted according to their payoff. The  $y$ -axis shows the expected reward in both graphs of Figure 2. The  $x$ -axis of Figure 2a shows the number of checkpoints allowed and the number of sources in Figure 2b. RANGER ran in 0.2 seconds in all trials.

In all cases, the two sampling methods developed here outperformed all others, with RS performing worse than CS when its optimality condition was violated more severely. MC, which does not prioritize targets, performs worse than the Uniform strategy in some situations, which sometimes places more coverage on higher value targets simply because there are more roads to them.

RANGER actually exploits resources differently and to better effect, which we explore in Figure 3. Data was used from the games in Figure 2b. Figure 3a shows the average number of edges with non-zero marginal probability in Entry-Adj and RANGER strategies as the number of sources increases. Clearly, RANGER uses *fewer* edges than Entry-Adj as the number of entry points increases, but uses them




Figure 4: Runtimes for RANGER, ERASER, and DOBSS.

to better effect.

Figure 3b shows the standard deviation of marginal probabilities of edges. While Entry-Incident remains constant at 0, RANGER’s results vary from 0 (with only one source) to 0.2, which is a standard deviation of 20% in marginal probability on edges. Although we cannot provide guarantees on the performance of Radius or Comb Sampling against an exact attacker in general, the techniques yield non-trivial mixed strategies of high quality in practice that outperform the sensible alternatives explored here.

### Runtime Comparison

We have discussed the exponentially smaller solution space RANGER, RS, and CS operate in; we now show runtime comparisons to verify our conclusions empirically. Specifically, we evaluate the runtime performance of RANGER against the fastest-known exact algorithm for solving general Bayesian Stackelberg games, DOBSS (Paruchuri et al. 2008), as well as a faster solver for security games, ERASER (Kiekintveld et al. 2009). DOBSS serves as a benchmark, since it provides the optimal solution against exact attackers. ERASER exploits structural properties that exist in many security domains, but requires that defender actions be independent from each other, which is violated in our domain since placing two checkpoints on one path will violate this assumption. Thus, ERASER can be used as an efficient, approximate alternative to DOBSS, so we compare RANGER’s runtime against it.

Figure 4 shows the scaling of each method’s runtime with respect to  $n$ . In these experiments, we use complete graphs (3 to 8 vertices) and random geometric graphs (4 to 12 vertices), each with one source, one target (value -1 to -10 for defender), and two checkpoints. The  $x$ -axis shows the number of vertices in the graph, and the  $y$ -axis shows runtime in seconds. Each result was averaged over 10 trials. Unsurprisingly, DOBSS and ERASER scale poorly and are only able to solve the problem with up to 7 vertices for complete graphs and 8 (DOBSS) and 11 (ERASER) vertices on random geometric graphs. RANGER is capable of solving games with 400 vertices within the time limit (not shown).

### Conclusions

In this work, we provide three primary contributions. First, we develop a linear program, RANGER, to efficiently cre-

ate optimal marginal distributions for placing checkpoints and prove that the reward found is an overestimate of the true optimal reward. Second, we introduce Radius Sampling, which we show produces optimal joint distributions under specific conditions. Third, we develop Comb Sampling, which we prove guarantees optimality against an approximating attacker. We complete the discussion by providing experimental verification of the high quality of our techniques on random graphs and a real-world domain.

### Acknowledgement

This research was supported by the United States Department of Homeland Security (DHS) through the Center for Risk and Economic Analysis of Terrorism Events (CRE-ATE) under grant number 2007-ST-061-000001. Opinions, conclusions or recommendations herein are solely those of the authors and do not necessarily reflect views of DHS.

### References

- Ali, S. A. 4 August 2009. Rs 18L seized in nakabandi at Vile Parle. *Times of India*.
- Basilico, N.; Gatti, N.; and Amigoni, F. 2009. Leader-follower strategies for robotic patrolling in environments with arbitrary topologies. In *AAMAS-09*.
- Chandran, R., and Beitchman, G. 29 November 2008. Battle for Mumbai ends, death toll rises to 195. *Times of India*.
- Conitzer, V., and Sandholm, T. 2006. Computing the optimal strategy to commit to. In *EC-06*.
- Gatti, N. 2008. Game theoretical insights in strategic patrolling: Model and algorithm in normal-form. In *ECAI-08*.
- Halvorson, E.; Conitzer, V.; and Parr, R. 2009. Multi-step Multi-sensor Hider-Seeker Games. In *IJCAI-09*.
- Israeli, E., and Wood, R. K. 2002. Shortest-path network interdiction. *Networks* 40(2):97–111.
- Kiekintveld, C.; Jain, M.; Tsai, J.; Pita, J.; Tambe, M.; and Ordóñez, F. 2009. Computing Optimal Randomized Resource Allocations for Massive Security Games. In *AAMAS-09*.
- Mavronicolas, M.; Papadopoulou, V.; Philippou, A.; and Spirakis, P. 2008. A network game with attackers and a defender. *Algorithmica* 51(3):315–341.
- Paruchuri, P.; Pearce, J. P.; Marecki, J.; Tambe, M.; Ordóñez, F.; and Kraus, S. 2008. Playing games with security: An efficient exact algorithm for Bayesian Stackelberg games. In *AAMAS-08*.
- Pita, J.; Jain, M.; Western, C.; Portway, C.; Tambe, M.; Ordóñez, F.; Kraus, S.; and Parachuri, P. 2008. Deployed ARMOR protection: The application of a game-theoretic model for security at the Los Angeles International Airport. In *AAMAS-08 (Industry Track)*.
- Tsai, J.; Rathi, S.; Kiekintveld, C.; Ordóñez, F.; and Tambe, M. 2009. IRIS - A Tool for Strategic Security Allocation in Transportation Networks. In *AAMAS-09 (Industry Track)*.
- Washburn, A., and Wood, K. 1995. Two-person zero-sum games for network interdiction. *Operations Research* 43(2):243–251.