

Computational Game Theory for Security and Sustainability

ALBERT XIN JIANG^{1,a)} MANISH JAIN^{1,b)} MILIND TAMBE^{1,c)}

Received: June 14, 2013, Accepted: January 6, 2014

Abstract: Security is a critical concern around the world. In many domains from counter-terrorism to sustainability, limited security resources prevent full security coverage at all times; instead, these limited resources must be scheduled, while simultaneously taking into account different target priorities, the responses of the adversaries to the security posture and potential uncertainty over adversary types. Computational game theory can help design such security schedules. Indeed, casting the problem as a Bayesian Stackelberg game, we have developed new algorithms that are now deployed over multiple years in multiple applications for security scheduling. These applications are leading to real-world use-inspired research in the emerging research area of “security games”; specifically, the research challenges posed by these applications include scaling up security games to large-scale problems, handling significant adversarial uncertainty, dealing with bounded rationality of human adversaries, and other interdisciplinary challenges.

Keywords: Stackelberg games, security, sustainability, deployed applications, column generation

1. Introduction

Security is a critical concern around the world that arises in protecting our ports, airports, transportation and other critical national infrastructure from adversaries, in protecting our wildlife and forests from poachers and smugglers, and in curtailing the illegal flow of weapons, drugs and money; and it arises in problems ranging from physical to cyber-physical systems. In all of these problems, we have limited security resources which prevent full security coverage at all times; instead, security resources must be deployed intelligently taking into account differences in priorities of targets requiring security coverage, the responses of the attackers to the security posture, and potential uncertainty over the types, capabilities, knowledge and priorities of attackers faced.

Game theory, which studies interactions among multiple self-interested agents, is well-suited to the adversarial reasoning required for security resource allocation and scheduling problems. Casting the problem as a Bayesian Stackelberg game, we have developed new algorithms for efficiently solving such games that provide randomized patrolling or inspection strategies. These algorithms have led to some initial successes in this challenging problem arena, leading to advances over previous approaches in security scheduling and allocation, e.g., by addressing key weaknesses of predictability of human schedulers. These algorithms are now deployed in multiple applications: ARMOR has been deployed at the Los Angeles International Airport (LAX) since 2007 to randomize checkpoints on the roadways entering the airport and canine patrol routes within the airport terminals [17];

IRIS, a game-theoretic scheduler for randomized deployment of the US Federal Air Marshals (FAMS) requiring significant scale-up in underlying algorithms, has been in use since 2009 [17]; PROTECT, which schedules the US Coast Guard’s randomized patrolling of ports using a new set of algorithms based on modeling bounded-rational human attackers, has been deployed in the port of Boston since April 2011 and is in use at the port of New York since February 2012 [34], and is headed for nationwide deployment; another application for deploying escort boats to protect ferries has been deployed by the US Coast Guard since April 2013 [10]; GUARDS is under evaluation for national deployment by the US Transportation Security Administration (TSA) [32], and TRUSTS [43] has been evaluated in field trials by the Los Angeles Sheriff’s Department (LASD) in the LA Metro system and a nation-wide deployment is now being evaluated at TSA. These initial successes point the way to major future applications in a wide range of security domains; with major research challenges in scaling up our game-theoretic algorithms, in addressing human adversaries’ bounded rationality and uncertainties in action execution and observation, as well as in multiagent learning.

This paper will provide an overview of the models and algorithms, key research challenges and a brief description of our successful deployments.

2. Stackelberg Security Games

Stackelberg games were first introduced to model leadership and commitment [39], and are now used to study security problems ranging from “police and robbers” scenario [11], computer network security [27], missile defense systems [5], and terrorism [33]. Models for arms inspections and border patrolling have also been modeled using inspection games [3], a related family of Stackelberg games.

¹ University of Southern California, Los Angeles, CA 90089, USA

^{a)} jiangx@usc.edu

^{b)} manish.jain@usc.edu

^{c)} tambe@usc.edu

This section provides details on this use of Stackelberg games for modeling security domains. We first give a generic description of security domains followed by *security games*, the model by which security domains are formulated in the Stackelberg game framework.

2.1 Security Domains

In a security domain, a defender must perpetually defend a set of targets using a limited number of resources, whereas the attacker is able to surveil and learn the defender's strategy and attack after careful planning. This fits precisely into the description of a Stackelberg game if we map the defender to the leader's role and the attacker to the follower's role [3], [6]. An action, or *pure strategy*, for the defender represents deploying a set of resources on patrols or checkpoints, e.g., scheduling checkpoints at the LAX airport or assigning federal air marshals to protect flight tours. The pure strategy for an attacker represents an attack at a target, e.g., a flight. The strategy for the leader is a *mixed strategy*, a probability distribution over the pure strategies of the defender. Additionally, with each target are also associated a set of payoff values that define the utilities for both the defender and the attacker in case of a successful or a failed attack. These payoffs are represented using the *security game* model, described next.

2.2 Security Games

A key assumption of security games is that the payoff of an outcome depends only on the target attacked, and whether or not it is covered by the defender [23]. The payoffs do *not* depend on the remaining aspects of the defender allocation. For example, if an adversary succeeds in attacking target t_1 , the penalty for the defender is the same whether the defender was guarding target t_2 or not.

This allows us to compactly represent the payoffs of a security game. Specifically, a set of four payoffs is associated with each target. These four payoffs are the rewards and penalties to both the defender and the attacker in case of a successful or an unsuccessful attack, and are sufficient to define the utilities for both players for all possible outcomes in the security domain. **Table 1** shows an example security game with two targets, t_1 and t_2 . In this example game, if the defender was *covering* (protecting) target t_1 and the attacker attacked t_1 , the defender would get 10 units of reward whereas the attacker would receive -1 units. We make the assumption that in a security game it is always better for the defender to cover a target as compared to leaving it uncovered, whereas it is always better for the attacker to attack an uncovered target. This assumption is consistent with the payoff trends in the real-world. A special case is *zero-sum games*, in which for each outcome the sum of utilities for the defender and attacker is zero, although in general security games are not necessarily zero-sum.

In the above example, all payoff values are exactly known. In

practice, we often have uncertainty over the payoffs and preferences of the players. Bayesian games are a well-known game-theoretic model in which such uncertainty is modeled using multiple types of players, with each associated with its own payoff values. For security games of interest, the main source of payoff uncertainty is regarding the attacker's payoffs. In the resulting *Bayesian Stackelberg game* model, there is only one leader type (e.g., only one police force), although there can be multiple follower types (e.g., multiple attacker types trying to infiltrate security) [30]. Each follower type is represented using a different payoff matrix. The leader does not know the follower's type, but knows the probability distribution over them. The goal is to find the optimal mixed strategy for the leader to commit to, given that the defender could be facing any of the follower types.

2.3 Solution Concept: Strong Stackelberg Equilibrium

The solution to a security game is a mixed strategy for the defender that maximizes the expected utility of the defender, given that the attacker learns the mixed strategy of the defender and chooses a best-response for himself. This solution concept is known as a Stackelberg equilibrium [25].

The most commonly adopted version of this concept in related literature is called Strong Stackelberg Equilibrium (SSE) [4], [9], [30], [40]. A SSE for security games is informally defined as follows (the formal definition of SSE is not introduced for brevity, and can instead be found in Ref. [23]):

Definition 1 A pair of strategies form a *Strong Stackelberg Equilibrium* (SSE) if they satisfy:

- (1) The defender plays a best-response, that is, the defender cannot get a higher payoff by choosing any other strategy.
- (2) The attacker plays a best-response, that is, given a defender strategy, the attacker cannot get a higher payoff by attacking any other target.
- (3) The attacker breaks ties in favor of the leader.

The assumption that the follower will always break ties in favor of the leader in cases of indifference is reasonable because in most cases the leader can induce the favorable strong equilibrium by selecting a strategy arbitrarily close to the equilibrium that causes the follower to strictly prefer the desired strategy [40]. Furthermore an SSE exists in all Stackelberg games, which makes it an attractive solution concept compared to versions of Stackelberg equilibrium with other tie-breaking rules. Finally, although initial applications relied on the SSE solution concept, we have since proposed new solution concepts that are more robust against various uncertainties in the model [2], [31], [42] and have used these robust solution concepts in some of the later applications.

3. Deployed and Emerging Security Applications

In this section, we describe several deployed and emerging applications of the Stackelberg game framework in different real-world domains. Besides describing successful transitions of research, our aim is to set the stage for later sections in which we discuss the research challenges that arise.

Table 1 Example of a security game with two targets.

Target	Defender		Attacker	
	Covered	Uncovered	Covered	Uncovered
t_1	10	0	-1	1
t_2	0	-10	-1	1

3.1 ARMOR for Los Angeles International Airport

Los Angeles International Airport (LAX) is the largest destination airport in the United States and serves 60–70 million passengers per year. The LAX police use diverse measures to protect the airport, which include vehicular checkpoints, police units patrolling the roads to the terminals, patrolling inside the terminals (with canines), and security screening and bag checks for passengers. The application of our game-theoretic approach is focused on two of these measures: (1) placing vehicle checkpoints on inbound roads that service the LAX terminals, including both location and timing, and (2) scheduling patrols for bomb-sniffing canine units at the different LAX terminals. The eight different terminals at LAX have very different characteristics, like physical size, passenger loads, international versus domestic flights, etc. These factors contribute to the differing risk assessments of these eight terminals. Furthermore, the numbers of available vehicle checkpoints and canine units are limited by resource constraints. Thus, it is challenging to optimally allocate these resources to improve their effectiveness while avoiding patterns in the scheduled deployments.

The ARMOR system (Assistant for Randomized Monitoring over Routes) focuses on two of the security measures at LAX (checkpoints and canine patrols) and optimizes security resource allocation using Bayesian Stackelberg games. Take the vehicle checkpoints model as an example (Fig. 1). Assuming that there are n roads, the police's strategy is placing $m < n$ checkpoints on these roads where m is the maximum number of checkpoints. ARMOR randomizes allocation of checkpoints to roads. The adversary may conduct surveillance of this mixed strategy and may potentially choose to attack through one of these roads. ARMOR models different types of attackers with different payoff functions, representing different capabilities and preferences for the attacker. ARMOR uses DOBSS (Decomposed Optimal Bayesian Stackelberg Solver) [30] to compute the defender's optimal strategy. ARMOR has been successfully deployed since August 2007 at LAX [17].

3.2 IRIS for US Federal Air Marshals Service

The US Federal Air Marshals Service (FAMS) allocates air marshals to flights originating in and departing from the United States to dissuade potential aggressors and prevent an attack should one occur. Flights are of different importance based on a variety of factors such as the numbers of passengers, the population of source and destination, and international flights from different countries. Security resource allocation in this domain is




Fig. 1 LAX checkpoints are deployed using ARMOR.

significantly more challenging than for ARMOR: a limited number of air marshals need to be scheduled to cover thousands of commercial flights each day. Furthermore, these air marshals must be scheduled on tours of flights that obey various constraints (e.g., the time required to board, fly, and disembark). Simply finding schedules for the marshals that meet all of these constraints is a computational challenge. Our task is made more difficult by the need to find a randomized policy that meets these scheduling constraints, while also accounting for the different values of each flight.


Against this background, the Iris system (Intelligent Randomization In Scheduling) has been developed and has been deployed by FAMS since October 2009 to randomize schedules of air marshals on international flights. In Iris, the targets are the set of n flights and the attacker could potentially choose to attack one of these flights. The FAMS can assign $m < n$ air marshals that may be assigned to protect these flights.

Since the number of possible schedules exponentially increases with the number of flights and resources, DOBSS is no longer applicable to the FAMS domain. Instead, Iris uses the much faster ASPEN algorithm [14] to generate the schedule for thousands of commercial flights per day.

3.3 PROTECT for US Coast Guard

The US Coast Guard's (USCG) mission includes maritime security of the US coasts, ports, and inland waterways; a security domain that faces increased risks due to threats such as terrorism and drug trafficking. Given a particular port and the variety of critical infrastructure that an attacker may attack within the port, USCG conducts patrols to protect this infrastructure; however, while the attacker has the opportunity to observe patrol patterns, limited security resources imply that USCG patrols cannot be at every location 24/7. To assist the USCG in allocating its patrolling resources, the PROTECT (Port Resilience Operational/Tactical Enforcement to Combat Terrorism) model has been designed to enhance maritime security. It has been in use at the port of Boston since April 2011, and is also in use at the port of New York since February 2012 (Fig. 2). Similar to previous applications ARMOR and IRIS, PROTECT uses an attacker-defender Stackelberg game framework, with USCG as the defender against terrorists that conduct surveillance before potentially launching an attack.

The key idea in PROTECT is also that unpredictability creates situations of uncertainty for an enemy and can be enough to deem a target less appealing. While randomizing patrol patterns is key,



(a) PROTECT is being used in Boston (b) Extending PROTECT to NY

Fig. 2 USCG boats patrolling the ports of Boston and NY.



Fig. 3 Escort boats protecting the Staten Island Ferry are using strategies generated by our system.

PROTECT also addresses the fact that the targets are of unequal value, understanding that the attacker will adapt to whatever patrol patterns USCG conducts. The output of PROTECT is a schedule of patrols which includes when the patrols are to begin, what critical infrastructure to visit for each patrol, and what activities to perform at each critical infrastructure.

While PROTECT builds on previous work, it offers key innovations. First, this system is a departure from the assumption of perfect attacker rationality noted in previous work, relying instead on a quantal response model [28] of the attacker's behavior. Second, to improve PROTECT's efficiency, a compact representation of the defender's strategies is used by exploiting equivalence and dominance. Finally, the evaluation of PROTECT for the first time provides real-world data: (i) comparison of human-generated vs PROTECT schedules, and (ii) results from an Adversarial Perspective Team's (human mock attackers) analysis. The PROTECT model has now been extended to other U.S. ports like Los Angeles/Long Beach and is moving towards nationwide deployment.

3.4 Ferry Protection for the US Coast Guard

Another problem that USCG faces is the protection of ferries, including the Staten Island Ferry in New York, from potential terrorist attacks from water. We developed a game-theoretic system for scheduling escort boat patrols to protect ferries, and this has been deployed at the Staten Island Ferry since 2013 [10] (Fig. 3). The key research challenge is the fact that the ferries are continuously moving in a continuous domain, and the attacker could attack at any moment in time. This type of moving targets domain leads to game-theoretic models with continuous strategy spaces, which presents computational challenges. Our theoretical work showed that while it is safe to discretize the defender's strategy space, discretizing the attacker's strategy space would result in loss of utility. We developed a novel algorithm that uses a compact representation for the defender's mixed strategy space while being able to exactly model the attacker's continuous strategy space. The implemented algorithm, running on a laptop, is able to generate daily schedules for escort boats with guaranteed expected utility values.

3.5 TRUSTS for Security in Transit Systems

Urban transit systems face multiple security challenges, including deterring fare evasion, suppressing crime and counter-terrorism. In particular, in some urban transit systems, including the Los Angeles Metro Rail system, passengers are legally



(a) Los Angeles Metro

(b) Barrier-free entrance to transit system

Fig. 4 TRUSTS for transit systems.

required to purchase tickets before entering but are not physically forced to do so (Fig. 4). Instead, security personnel are dynamically deployed throughout the transit system, randomly inspecting passenger tickets. This proof-of-payment fare collection method is typically chosen as a more cost-effective alternative to direct fare collection, i.e., when the revenue lost to fare evasion is believed to be less than what it would cost to directly preclude it. In the case of Los Angeles Metro, with approximately 300,000 riders daily, this revenue loss can be significant; the annual cost has been estimated at \$5.6 million [12]. The Los Angeles Sheriff's Department (LASD) deploys uniformed patrols on board trains and at stations for fare-checking (and for other purposes such as crime prevention). The LASD's current approach relies on humans for scheduling the patrols, which places a tremendous cognitive burden on the human schedulers who must take into account all of the scheduling complexities (e.g., train timings, switching time between trains, and schedule lengths).

The TRUSTS system (Tactical Randomization for Urban Security in Transit Systems) models the patrolling problem as a leader-follower Stackelberg game [43]. The leader (LASD) pre-commits to a mixed strategy patrol (a probability distribution over all pure strategies), and riders observe this mixed strategy before deciding whether to buy the ticket or not. Both ticket sales and fines issued for fare evasion translate into revenue for the government. Therefore the utility for the leader is the total revenue (total ticket sales plus penalties). The main computational challenge is the exponentially many possible patrol strategies, each subject to both the spatial and temporal constraints of travel within the transit network under consideration. To overcome this challenge, Trusts uses a compact representation of the strategy space which captures the spatiotemporal structure of the domain.

The LASD conducted field tests of this TRUSTS system in the LA Metro in 2012, and one of the feedback comments from the officers was that patrols are often interrupted due to execution uncertainty such as emergencies and arrests. Utilizing techniques from planning under uncertainty (in particular Markov Decision Processes), we proposed a general approach to dynamic patrolling games in uncertain environments, which provides patrol strategies with contingency plans [19]. This led to schedules now being loaded onto smartphones and given to officers. If interruptions occur, the schedules are then automatically updated on the smartphone app. The LASD has conducted successful field evaluations using the smartphone app, and the TSA is currently evaluating it toward nationwide deployment.

Crime presents a serious problem in transit systems like LA

Metro. Furthermore, unlike terrorists that strategically plans an attack, criminals are often opportunistic, in that their decisions are based on the available opportunities encountered. For the crime problem, we developed a new game-theoretic model that utilizes recent advances in criminology on modeling opportunistic criminals, and novel efficient algorithms that achieve speedups by exploiting the spatiotemporal structure of the domain [45].

3.6 GUARDS for US Transportation Security Agency

The United States Transportation Security Administration (TSA) is tasked with protecting the nation's over 400 airports. To aid the TSA in scheduling resources to protect airports, a new application called GUARDS (Game-theoretic Unpredictable and Randomly Deployed Security) has been developed. While GUARDS also uses Stackelberg games as ARMOR and IRIS, GUARDS faces three key challenges [32]: 1) reasoning about hundreds of heterogeneous security activities; 2) reasoning over diverse potential threats; and 3) developing a system designed for hundreds of end-users. To address those challenges, GUARDS created a new game-theoretic framework that allows for heterogeneous defender activities and compact modeling of a large number of threats and developed an efficient solution technique based on general-purpose Stackelberg game solvers. GUARDS is currently under evaluation and testing for scheduling practices at an undisclosed airport.

3.7 Emerging Applications in Networked Domains

Beyond the deployed applications above, there are a number of emerging application areas. One such area of great importance is securing urban city networks, transportation networks, computer networks and other network centric security domains. For example, after the terrorist attacks in Mumbai of 2008 [8], the Mumbai police have started setting up vehicular checkpoints on roads. We can model the problem faced by the Mumbai police as a security game between the Mumbai police and an attacker. In this urban security game, the pure strategies of the defender correspond to allocations of resources to edges in the network—for example, an allocation of police checkpoints to roads in the city. The pure strategies of the attacker correspond to paths from any *source* node to any *target* node—for example, a path from a landing spot on the coast to the airport. The strategy space of the defender grows exponentially with the number of available resources, whereas the strategy space of the attacker grows exponentially with the size of the network. In addressing this computational challenge, novel algorithms based on incremental strategy generation have been able to generate randomized defender strategies that scale up to the entire road network of Mumbai [16].

The Stackelberg game framework can also be applied to adversarial domains that exhibit ‘contagious’ actions for each player. For example, word-of-mouth advertising/viral marketing has been widely studied by marketers trying to understand why one product or video goes ‘viral’ while others go unnoticed. Counter-insurgency is the contest for the support of the local leaders in an armed conflict and can include a variety of operations such as providing security and giving medical supplies. These efforts carry a social effect beyond the action taken that can cause advantageous ripples through the neighboring population. Moreover, multiple

intelligent parties attempt to leverage the same social network to spread their message, necessitating an adversary-aware approach to strategy generation. Game-theoretic approaches can be used to generate resource allocations strategies for such large-scale, real world networks [36], [37]. This interaction can be modeled as a graph with one player attempting to spread influence while another player attempts to stop the probabilistic propagation of that influence by spreading their own influence. This ‘blocking’ problem models situations faced by governments/peacekeepers combatting the spread of terrorist radicalism and armed conflict with daily/weekly/monthly visits with local leaders to provide support and discuss grievances [13].

Game-theoretic methods are also appropriate for modeling resource allocation in cyber-security such as packet selection and inspection for detecting potential threats in large computer networks. The problem of attacks on computer systems and corporate computer networks gets more pressing each year. A number of intrusion detection and monitoring systems are being developed, e.g., deep packet inspection method that periodically selects a subset of packets in a computer network for analysis. The attacking/protecting problem can be formulated as a game between two players: the attacker (or the intruder), and the defender (the detection system). The actions of the attacker can be seen as sending malicious packets from a controlled computer to vulnerable computers. The objective of the defender is to prevent the intruder from succeeding by selecting the packets for inspection and subsequently thwarting the attack. However, packet inspections cause unwanted latency and hence the defender has to decide where and how to inspect network traffic. The computational challenge is efficiently computing the optimal defending strategies for such network scenarios [38].

3.8 Emerging Applications in Sustainability

A number of our newer applications are focused on sustainability, through suppression of environmental crime. One area is protecting forests [20], where we must protect a continuous forest area from extractors by patrols through the forest that seek to deter such extraction activity. With limited resources for performing such patrols, a patrol strategy will seek to distribute the patrols throughout the forest, in space and time, in order to minimize the resulting amount of extraction that occurs or maximize the degree of forest protection. This problem can be formulated as a Stackelberg game and the focus is on computing optimal allocations of patrol density [20].

Endangered species poaching is reaching critical levels as the populations of these species plummet to unsustainable numbers. The global tiger population, for example, has dropped over 95% from the start of the 1900s and has resulted in three out of nine species extinctions. Depending on the area and animals poached, motivations for poaching range from profit to sustenance, with the former being more common when profitable species such as tigers, elephants, and rhinos are the targets. To counter poaching efforts and to rebuild the species’ populations, countries have set up protected wildlife reserves and conservation agencies tasked with defending these large reserves. Because of the size of the reserves and the common lack of law enforcement resources, con-

servation agencies are at a significant disadvantage when it comes to deterring and capturing poachers. Agencies use patrolling as a primary method of securing the park. Due to their limited resources, however, patrol managers must carefully create patrols that account for many different variables (e.g., limited patrol units to send out, multiple locations that poachers can attack at varying distances to the outpost). Our proposed system Protection Assistant for Wildlife Security (PAWS) aims to assist conservation agencies in their critical role of patrol creation by predicting where poachers will attack and optimizing patrol routes to cover those areas.

Another emerging application domain is that of ensuring the sustainability of fish resources. Marine fisheries are acknowledged to be some of the most important food resources for countries around the world. As reported by World Wild Fund for Nature (WWF), cod are currently at risk from overfishing in the UK, Canada and most other Atlantic countries. Global cod catch has suffered a 70% drop over the last 30 years, and if this trend continues, the world's cod stocks will disappear in 15 years. Illegal, unreported, and unregulated (IUU) fishing is one of the major threats to the sustainability of ocean fish resources. As estimated by National Oceanic and Atmospheric Administration (NOAA), IUU fishing produces between 11 and 26 million tons of seafood annually, representing as much as 40 percent of the total catch in some fisheries. The driver behind IUU fishing is high economic profit and low chance of seizure. It is impossible to maintain a 24/7 presence to prevent IUU fishing everywhere due to the limited asset patrolling resources. Hence the allocation of the patrolling resources becomes a key challenge for security agencies like USCG. Utilizing data on fish locations, as well as historical data on USCG patrols and captures, we developed an algorithm that combines machine learning techniques and game-theoretic planning, and developed a prototype application that is currently under evaluation by US Coast Guard.

4. Scaling Up To Real-world Problem Sizes

The wide use of Stackelberg games has inspired theoretical and algorithmic progress leading to the development of fielded applications, as described in Section 3. For example, Dobss [30], an algorithm for solving Bayesian Stackelberg games, is central to the fielded application ARMOR in use at the Los Angeles International Airport [17]. Conitzer and Sandholm [9] gave complexity results and algorithms for computing optimal commitment strategies in Bayesian Stackelberg games, including both pure and mixed-strategy commitments.

These early works assumed that the set of pure strategies for the players are given explicitly. Many real world problems, like the FAMS and urban road networks, present billions of pure strategies to both the defender and the attacker. Such large problem instances cannot even be represented in modern computers, let alone solved using previous techniques. We have proposed models and algorithms that compute optimal defender strategies for massive real-world security domains [14], [15].

4.1 Scaling Up with Defender Pure Strategies

In this section, we describe one particular algorithm ASPEN, that

computes strong Stackelberg equilibria (SSE) in domains with a *very large* number of pure strategies (up to billions of actions) for the defender [14]. ASPEN builds on the insight that in many real-world game-theoretic problems, there exist solutions with *small support sizes*, which are mixed strategies in which only a small set of pure strategies are played with positive probability [26]. ASPEN exploits this by using a *strategy generation* approach for the defender, in which defender pure strategies are iteratively generated and added to the optimization formulation.

As an example, let us consider the problem faced by the FAMS. There are currently tens of thousands of commercial flights flying each day, and public estimates state that there are thousands of air marshals that are scheduled daily by the FAMS [22]. Air marshals must be scheduled on tours of flights that obey logistical constraints (e.g., the time required to board, fly, and disembark). An example of a schedule is an air marshal assigned to a round trip from Los Angeles to New York and back.

ASPEN [14] casts this problem as a security game, where the attacker can choose any of the flights to attack, and each air marshal can cover one schedule. Each schedule here is a feasible set of targets that can be covered together; for the FAMS, each schedule would represent a flight tour which satisfies all the logistical constraints that an air marshal could fly. A *joint schedule* then would assign every air marshal to a flight tour, and there could be exponentially many joint schedules in the domain. A pure strategy for the defender in this security game is a joint schedule. As mentioned previously, ASPEN employs strategy generation since all the defender pure strategies cannot be enumerated for such a massive problem. ASPEN decomposes the problem into a *master* problem and a *slave* problem, which are then solved iteratively. Given a number of pure strategies, the master solves for the defender and the attacker optimization constraints, while the slave is used to generate a new pure strategy for the defender in every iteration.

The iterative process is graphically depicted in Fig. 5. The master operates on the pure strategies (joint schedules) generated thus far, which are represented using the matrix \mathbf{P} . Each column of \mathbf{P} , \mathbf{J}_j , is one pure strategy (or joint schedule). An entry P_{ij} in the matrix \mathbf{P} is 1 if a target t_i is covered by joint-schedule \mathbf{J}_j , and 0 otherwise. The objective of the master problem is to compute \mathbf{x} , the optimal mixed strategy of the defender over the pure strategies in \mathbf{P} . The objective of the slave problem is to generate the best joint schedule to add to \mathbf{P} . The best joint schedule is identi-

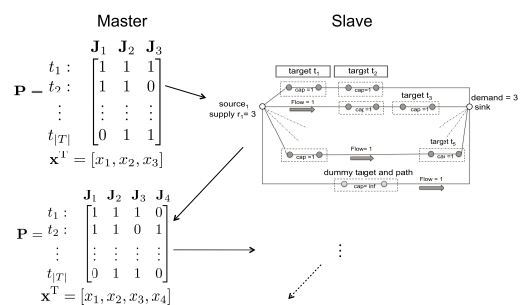


Fig. 5 Strategy generation employed in ASPEN: The schedules for a defender are generated iteratively. The *slave* problem is a novel minimum-cost integer flow formulation that computes the new pure strategy to be added to \mathbf{P} ; \mathbf{J}_4 is computed and added in this example.

fied using the concept of *reduced costs*, which measures if a pure strategy can potentially increase the defender's expected utility (the details of the approach are provided in Ref. [14]). While a naïve approach would be to iterate over all possible pure strategies to identify the pure strategy with the maximum potential, ASPEN uses a novel minimum-cost integer flow problem to efficiently identify the best pure strategy to add. ASPEN always converges on the optimal mixed strategy for the defender.

Employing strategy generation for large optimization problems is not an “out-of-the-box” approach, the problem has to be formulated in a way that allows for domain properties to be exploited. The novel contribution of ASPEN is to provide a linear formulation for the master and a minimum-cost integer flow formulation for the slave, which enables the application of strategy generation techniques. Additionally, ASPEN also provides a branch-and-bound heuristic to reason over attacker actions. This branch-and-bound heuristic provides a further order of magnitude speed-up, allowing ASPEN to handle the massive sizes of real-world problems.

4.2 Scaling Up with Defender and Attacker Pure Strategies

In domains such as the urban network security setting described in Section 3.7, the number of pure strategies of both the defender and the attacker are exponentially large. In this section, we describe the RUGGED algorithm [15], which generates pure strategies for both the defender and the attacker.

RUGGED models the domain as a zero-sum game, and computes the minimax equilibrium, since the minimax strategy is equivalent to the SSE in zero-sum games. **Figure 6** shows the working of RUGGED: at each iteration, the Minimax module generates the optimal mixed strategies $\langle \mathbf{x}, \mathbf{a} \rangle$ for the two players for the current payoff matrix, the Best Response Defender module generates a new strategy for the defender that is a best response against the attacker's current strategy \mathbf{a} , and the Best Response Attacker module generates a new strategy for the attacker that is a best response against the defender's current strategy \mathbf{x} . The rows X_i in the figure are the pure strategies for the defender, they would correspond to an allocation of checkpoints in the urban road network domain. Similarly, the columns A_j are the pure strategies for the attacker, they represent the attack paths in the urban road network domain. The values in the matrix represent the payoffs to the defender. The algorithm stops when neither of the generated best responses improve on the current minimax strategies.

The contribution of RUGGED is to provide the mixed integer for-

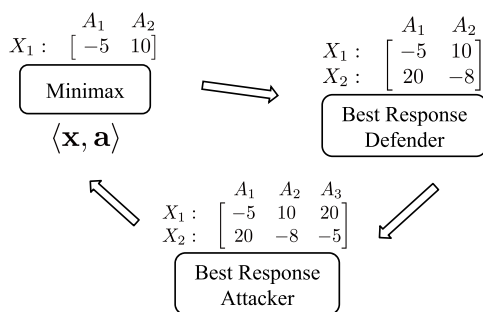


Fig. 6 Strategy Generation employed in RUGGED: The pure strategies for both the defender and the attacker are generated iteratively.

mulations for the best response modules which enable the application of such a strategy generation approach. RUGGED can compute the optimal solution for deploying up to 4 resources in real-city network with as many as 250 nodes within a reasonable time frame of 10 hours (the complexity of this problem can be estimated by observing that both the best response problems are NP-hard themselves [15]). More recent work [16] builds on RUGGED and proposes SNARES, which allows scale-up to the entire city of Mumbai, with 10–15 checkpoints.

5. Current Research

In this section we highlight several areas that we are actively doing research on, and point out some of the open research challenges.

Scalability: Driven by the growing complexity of applications, a sequence of algorithms for solving security games have been developed including DOBSS [30], ERASER [23], ASPEN [14] and RUGGED [15]. However, existing algorithms still cannot scale up to very large scale domains. While RUGGED/SNARES computes optimal solutions much faster than any of the previous approaches, much work remains to be done for it to be applicable to complex heterogeneous settings on large networks.

Besides strategy generation, another approach for dealing with an exponential number of pure strategies is to compactly represent mixed strategies as marginal probabilities of coverage on each of the targets. Because of the utility structure of security games, such marginal probabilities are sufficient to express the expected utility of the defender. Kiekintveld et al. [23] used this approach in ERASER to formulate the problem of computing SSE as a compact mixed-integer linear program. However, this approach is unable to deal with complex constraints on the defender resources [24]. Nevertheless, we have recently been able to use this approach for certain patrolling domains, including fare-enforcement patrols in urban transit systems [43] and boat patrols for protecting ferries [10]. In these domains a pure strategy is a patrol of a certain time duration over a set of locations, and the number of such pure strategies grow exponentially in the time duration. We were able to compactly represent mixed strategies as fractional flows on the *transition graph*, in which vertices are time-location pairs and arcs represent possible actions. This allowed us to formulate the optimization problems compactly which led to improved scalability. An open problem is to find other types of security domains in which the strategy space can be compactly represented. Another is to develop a hybrid approach that combines marginals and strategy generation.

Robustness: Classical game theory solution concepts often make assumptions on the knowledge, rationality, and capability (e.g., perfect recall) of players. Unfortunately, these assumptions could be wrong in real-world scenarios. Algorithms for the defender's optimal strategy have been proposed to take into account various uncertainties faced in the domain, including payoff noise [44], execution/observation error [42], and uncertain capability [2]. However, previous works assumed that the attacker knows (or with a small noise) the defender's mixed strategy. Recently An et al. [1] proposed a formal framework to model the attacker's belief update process as he observes instantiations of the

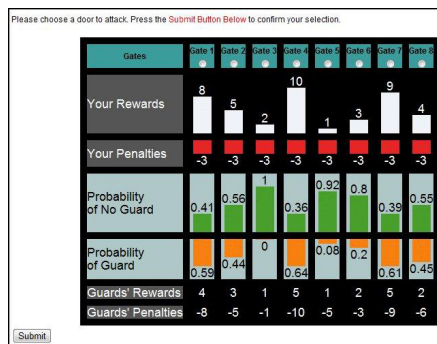


Fig. 7 Interface of the Guards and Treasures game, based on the security scenario at the LAX airport.

defender's mixed strategy. The resulting optimization problem for the defender is nonlinear and scalable computation remains an open issue.

Human adversary modeling: One required research direction is addressing bounded rationality of human adversaries. This is a fundamental problem that can affect the performance of our game theoretic solutions, since algorithms based on the assumption of the perfectly rational adversary are not robust to deal with deviations of the adversary from the optimal response. Recently, there has been some research on applying ideas from behavioral game theory (e.g., prospect theory [21] and quantal response [28]) within security game algorithms. One line of approaches is based on the quantal response model to predict the behaviors of the human adversary, and then to compute optimal defender strategies against such behavior of the adversary. These include BRQR [41] which follows the logit quantal response (QR) [28] model, and subsequent work on subjective-utility quantal response (SUQR) models [29]. The parameters of these models are estimated by experimental tuning. **Figure 7** shows the interface of an interactive game used in our human subject experiments, based on the security scenario at the LAX airport. We have made the source code of the game available at <http://teamcore.usc.edu/projects/BGT/experiment.html>.

However, in real-world security domains, we may have very limited data, or may only have some limited information on the biases displayed by adversaries. An alternative approach is based on robust optimization: instead of assuming a particular model of human decision making, try to achieve good defender expected utility against a range of possible models. One instance of this approach is MATCH [31], which guarantees a bound for the loss of the defender to be within a constant factor of the adversary loss if the adversary responds non-optimally. Another robust solution concept is monotonic maximin [18], which tries to optimize defender utility against the worst-case monotonic adversary behavior, where monotonicity is the property that actions with higher expected utility is played with higher probability. An open research challenge is to combine such robust-optimization approaches with available behavior data. Furthermore, since real-world human attackers are sometimes distributed coalitions of socially, culturally and cognitively-biased agents, we may need significant interdisciplinary research to build in social, cultural and coalitional biases into our adversary models.

Multi-Objective Optimization: In existing applications such

as ARMOR, IRIS and PROTECT, the defender is trying to maximize a single objective. However, there are domains where the defender has to consider multiple objectives simultaneously. Multi-objective security games (MOSG) have been proposed to address the challenges of domains with multiple incomparable objectives [7]. In an MOSG, the threats posed by the attacker types are treated as different objective functions which are not aggregated, thus eliminating the need for a probability distribution over attacker types. Unlike Bayesian security games which have a single optimal solution, MOSGs have a set of Pareto-optimal (non-dominated) solutions which is referred to as the Pareto frontier. By presenting the Pareto frontier to the end user, they may be able to better understand the structure of their problem as well as the trade-offs between different security strategies.

Evaluation: Evaluation in itself is a major challenge given the real-world deployment of these systems. We have conducted a number of such evaluations: simulations, human subjects in the lab, adversary perspective teams (mock attacker teams) before and after deployment, assessment by domain experts internal and external to agencies deploying these applications, and data from deployments (such as number of citations to fare-evaders) have all been used. Space precludes us from discussions of these evaluations, but they are discussed in individual publications on the applications [17], [34], [43], and all and more of these are discussed in Ref. [35].

6. Summary

Security is recognized as a world-wide challenge and game theory is an increasingly important paradigm for reasoning about complex security resource allocation. While the deployed game theoretic applications have provided a promising start, very significant amount of research remains to be done. These are large-scale interdisciplinary research challenges that call upon multi-agent researchers to work with researchers in other disciplines, be "on the ground" with domain experts, and examine real-world constraints and challenges that cannot be abstracted away.

Acknowledgments This research is supported by MURI grant W911NF-11-1-0332 and by the United States Department of Homeland Security through the Center for Risk and Economic Analysis of Terrorism Events (CREATE) under grant number 2010-ST-061-RE0001. All opinions, findings, conclusions and recommendations in this document are those of the authors and do not necessarily reflect views of the United States Department of Homeland Security.

References

- [1] An, B., Kempe, D., Kiekintveld, C., Shieh, E., Singh, S., Tambe, M. and Vorobeychik, Y.: Security Games with Limited Surveillance, *Conference on Artificial Intelligence (AAAI)* (2012).
- [2] An, B., Tambe, M., Ordonez, F., Shieh, E. and Kiekintveld, C.: Refinement of Strong Stackelberg Equilibria in Security Games, *Proc. 25th Conference on Artificial Intelligence*, pp.587–593 (2011).
- [3] Avenhaus, R., von Stengel, B. and Zamir, S.: Inspection Games, *Handbook of Game Theory*, Aumann, R.J. and Hart, S. (Eds.), Vol.3, North-Holland, Amsterdam, chapter 51, pp.1947–1987 (2002).
- [4] Breton, M., Alg, A. and Haurie, A.: Sequential Stackelberg Equilibria in Two-Person Games, *Optimization Theory and Applications*, Vol.59, No.1, pp.71–97 (1988).
- [5] Brown, G., Carlyle, M., Kline, J. and Wood, K.: A Two-Sided Optimization for Theater Ballistic Missile Defense, *Operations Research*,

- Vol.53, pp.263–275 (2005).
- [6] Brown, G., Carlyle, M., Salmeron, J. and Wood, K.: Defending Critical Infrastructure, *Interfaces*, Vol.36, No.6, pp.530–544 (2006).
 - [7] Brown, M., An, B., Kiekintveld, C., Ordonez, F. and Tambe, M.: Multi-objective optimization for security games, *Proc. 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (2012).
 - [8] Chandran, R. and Beitchman, G.: Battle for Mumbai Ends, Death Toll Rises to 195, *Times of India* (29 November 2008), available from http://articles.timesofindia.indiatimes.com/2008-11-29/india/27930171_1_taj-hotel-three-terrorists-nariman-house.
 - [9] Conitzer, V. and Sandholm, T.: Computing the Optimal Strategy to Commit to, *Proc. ACM Conference on Electronic Commerce (ACM-EC)*, pp.82–90 (2006).
 - [10] Fang, F., Jiang, A. and Tambe, M.: Optimal Patrol Strategy for Protecting Moving Targets with Multiple Mobile Resources, *AAMAS* (2013).
 - [11] Gatti, N.: Game Theoretical Insights in Strategic Patrolling: Model and Algorithm in Normal-Form, *ECAI-08*, pp.403–407 (2008).
 - [12] Hamilton, B.A.: Faregating Analysis. Report Commissioned by the LA Metro (2007), available from http://boardarchives.metro.net/Items/2007/11_November/20071115EMACItem27.pdf.
 - [13] Howard, N.J.: Finding Optimal Strategies for Influencing Social Networks in Two Player Games, Master's thesis, MIT, Sloan School of Management (2011).
 - [14] Jain, M., Kardes, E., Kiekintveld, C., Ordonez, F. and Tambe, M.: Security Games with Arbitrary Schedules: A Branch and Price Approach, *Proc. 24th AAAI Conference on Artificial Intelligence*, pp.792–797 (2010).
 - [15] Jain, M., Korzhuk, D., Vanek, O., Pechoucek, M., Conitzer, V. and Tambe, M.: A Double Oracle Algorithm for Zero-Sum Security games on Graphs, *Proc. 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (2011).
 - [16] Jain, M., Tambe, M. and Conitzer, V.: Security Scheduling for Real-world Networks, *AAMAS* (2013).
 - [17] Jain, M., Tsai, J., Pita, J., Kiekintveld, C., Rath, S., Tambe, M. and Ordonez, F.: Software Assistants for Randomized Patrol Planning for the LAX Airport Police and the Federal Air Marshal Service, *Interfaces*, Vol.40, pp.267–290 (2010).
 - [18] Jiang, A.X., Nguyen, T.H., Tambe, M. and Procaccia, A.D.: Monotonic Maximin: A Robust Stackelberg Solution Against Boundedly Rational Followers, *Conference on Decision and Game Theory for Security (GameSec)* (2013).
 - [19] Jiang, A., Yin, Z., Kraus, S., Zhang, C. and Tambe, M.: Game-theoretic Randomization for Security Patrolling with Dynamic Execution Uncertainty, *AAMAS* (2013).
 - [20] Johnson, M., Fang, F., Yang, R., Tambe, M. and Albers, H.: Patrolling to Maximize Pristine Forest Area, *Proc. AAAI Spring Symposium on Game Theory for Security, Sustainability and Health* (2012).
 - [21] Kahneman, D. and Tversky, A.: Prospect Theory: An Analysis of Decision Under Risk, *Econometrica*, Vol.47, No.2, pp.263–291 (1979).
 - [22] Keteyian, A.: TSA: Federal Air Marshals (2010), available from <http://www.cbsnews.com/stories/2010/02/01/earlyshow/main6162291.shtml>, (accessed 2011-02-01).
 - [23] Kiekintveld, C., Jain, M., Tsai, J., Pita, J., Tambe, M. and Ordonez, F.: Computing Optimal Randomized Resource Allocations for Massive Security Games, *Proc. 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp.689–696 (2009).
 - [24] Korzhuk, D., Conitzer, V. and Parr, R.: Complexity of computing optimal Stackelberg strategies in security resource allocation games, *Proc. 24th AAAI Conference on Artificial Intelligence*, pp.805–810 (2010).
 - [25] Leitmann, G.: On Generalized Stackelberg Strategies, *Optimization Theory and Applications*, Vol.26, No.4, pp.637–643 (1978).
 - [26] Lipton, R., Markakis, E. and Mehta, A.: Playing large games using simple strategies, *EC: Proc. ACM Conference on Electronic Commerce*, pp.36–41, ACM (2003).
 - [27] Lye, K. and Wing, J.M.: Game Strategies in Network Security, *International Journal of Information Security*, Vol.4, No.1–2, pp.71–86 (2005).
 - [28] McKelvey, R.D. and Palfrey, T.R.: Quantal Response Equilibria for Normal Form Games, *Games and Economic Behavior*, Vol.10, No.1, pp.6–38 (1995).
 - [29] Nguyen, T.H., Yang, R., Azaria, A., Kraus, S. and Tambe, M.: Analyzing the Effectiveness of Adversary Modeling in Security Games, *Conference on Artificial Intelligence (AAAI)* (2013).
 - [30] Paruchuri, P., Pearce, J.P., Marecki, J., Tambe, M., Ordonez, F. and Kraus, S.: Playing Games with Security: An Efficient Exact Algorithm for Bayesian Stackelberg Games, *Proc. 7th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp.895–902 (2008).
 - [31] Pita, J., John, R., Maheswaran, R., Tambe, M. and Kraus, S.: A Robust Approach to Addressing Human Adversaries in Security Games, *European Conference on Artificial Intelligence (ECAI)* (2012).
 - [32] Pita, J., Tambe, M., Kiekintveld, C., Cullen, S. and Steigerwald, E.: GUARDS - Game Theoretic Security Allocation on a National Scale, *Proc. 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (2011).
 - [33] Sandler, T. and Arce, M.D.G.: Terrorism and Game Theory, *Simulation and Gaming*, Vol.34, No.3, pp.319–337 (2003).
 - [34] Shieh, E., An, B., Yang, R., Tambe, M., Baldwin, C., DiRenzo, J., Maule, B. and Meyer, G.: PROTECT: A Deployed Game Theoretic System to Protect the Ports of the United States, *Proc. 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (2012).
 - [35] Taylor, M., Kiekintveld, C. and Tambe, M.: Evaluating Deployed Decision-Support Systems for Security: Challenges, Analysis, and Approaches, *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*, Tambe, M. (Ed.), Cambridge University Press (2011).
 - [36] Tsai, J., Qian, Y., Vorobeychik, Y., Kiekintveld, C. and Tambe, M.: Bayesian security games for controlling contagion, *Proc. ASE/IEEE International Conference on Social Computing (SocialCom)* (2013).
 - [37] Tsai, J., Nguyen, T.H. and Tambe, M.: Security Games for Controlling Contagion, *Conference on Artificial Intelligence (AAAI)* (2012).
 - [38] Vanek, O., Yin, Z., Jain, M., Bosansky, B., Tambe, M. and Pechoucek, M.: Game-Theoretic Resource Allocation for Malicious Packet Detection in Computer Networks, *Proc. 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (2012).
 - [39] von Stackelberg, H.: *Marktform und Gleichgewicht*, Springer, Vienna (1934).
 - [40] von Stengel, B. and Zamir, S.: Leadership with Commitment to Mixed Strategies, Technical Report LSE-CDAM-2004-01, CDAM Research Report (2004).
 - [41] Yang, R., Kiekintveld, C., Ordonez, F., Tambe, M. and John, R.: Improving Resource Allocation Strategy Against Human Adversaries in Security Games, *IJCAI* (2011).
 - [42] Yin, Z., Jain, M., Tambe, M. and Ordonez, F.: Risk-Averse Strategies for Security Games with Execution and Observational Uncertainty, *Proc. 25th AAAI Conference on Artificial Intelligence (AAAI)*, pp.758–763 (2011).
 - [43] Yin, Z., Jiang, A., Johnson, M., Tambe, M., Kiekintveld, C., Leyton-Brown, K., Sandholm, T. and Sullivan, J.: TRUSTS: Scheduling Randomized Patrols for Fare Inspection in Transit Systems, *Proc. 24th Conference on Innovative Applications of Artificial Intelligence (IAAI)* (2012).
 - [44] Yin, Z. and Tambe, M.: A Unified Method for Handling Discrete and Continuous Uncertainty in Bayesian Stackelberg Games, *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (2012).
 - [45] Zhang, C., Jiang, A.X., Short, M.B., Brantingham, P.J. and Tambe, M.: Modeling Crime diffusion and crime suppression on transportation networks: An initial report, *SNSC 2013: The AAAI Fall Symposium 2013 on Social Networks and Social Contagion* (2013).



Albert Xin Jiang is a postdoctoral researcher in the TEAMCORE research group at the Department of Computer Science at the University of Southern California. He received his Ph.D. from the Department of Computer Science at the University of British Columbia. Much of his research is addressing computational

problems arising in game theory, including the efficient computation of solution concepts such as Nash equilibrium, Stackelberg equilibrium and correlated equilibrium, as well as applications of game-theoretic computation to real-world domains such as large-scale infrastructure security and electronic commerce. He has published over 40 peer-reviewed technical articles. He received the Best Student Paper Award at the ACM Conference on Electronic Commerce (ACM-EC) in 2011 and was a finalist for the best paper award at AAMAS 2013. He received the 2012 CAIAC Doctoral Dissertation Award and the runner up IFAAMAS Victor Lesser Distinguished Dissertation Award. As part of the Port Resilience Operational/Tactical Enforcement to Combat Terrorism (PROTECT) team, he received a Coast Guard Meritorious Team Commendation in 2013 for creating “an innovative approach to optimize patrol schedules and actions for the Coast Guard Ports, Waterways and Coast Security missions.”



Manish Jain completed his Ph.D. from the Department of Computer Science at the University of Southern California in August 2013. His work is on the applications of game-theoretic and large-scale optimization techniques for scheduling limited resources. His work has been deployed by many agencies within the Department of Homeland Security, including the Federal Air Marshals Service and Los Angeles World Airports police. He has published numerous papers and given multiple tutorials on the topic of Security Games at many AI conferences and publication venues. His work is also the recipient of the RIST Prize by the Military Operations Research Society, and a finalist for the EURO excellence in Practice award. He has also won the best Dissertation award from the CS Department at the University of Southern California. He has also been recently nominated for the best ACM Dissertation Award in 2013. He is currently the Chief Technology Officer of Armorway Inc. Armorway is dedicated to developing products for intelligent optimized security allocations, both for law enforcement as well as private agencies, countering adversaries in domains ranging from physical infrastructure to computer networks.



Milind Tambe is Helen N. and Emmett H. Jones Professor in Engineering at the University of Southern California (USC). He is a fellow of AAAI (Association for Advancement of Artificial Intelligence) and ACM (Association for Computing Machinery), as well as recipient of the ACM/SIGART Autonomous Agents Re-

search Award, Christopher Columbus Fellowship Foundation Homeland security award, the INFORMS Wagner prize for excellence in Operations Research practice and the Rist Prize of the Military Operations Research Society. Professor Tambe has contributed several foundational papers in agents and multiagent systems; this includes areas of multiagent teamwork, distributed constraint optimization (DCOP) and security games. For this research, he has received the “influential paper award” from the International Foundation for Agents and Multiagent Systems (IFAAMAS), as well as with his research group, best paper awards at a number of premier Artificial Intelligence Conferences and workshops; these have included multiple best paper awards at the International Conference on Autonomous Agents and Multiagent Systems and International Conference on Intelligent Virtual Agents. In addition, the “security games” framework and algorithms pioneered by Professor Tambe and his research group are now deployed for real-world use by several agencies including the US Coast Guard, the US Federal Air Marshals service, the Transportation Security Administration, LAX Police and the LA Sheriff’s Department for security scheduling at a variety of US ports, airports and transportation infrastructure. This research has led to him and his students receiving the US Coast Guard Meritorious Team Commendation from the Commandant, US Coast Guard First District’s Operational Excellence Award, Certificate of Appreciation from the US Federal Air Marshals Service and special commendation given by the Los Angeles World Airports police from the city of Los Angeles. In addition, for his research, Professor Tambe has received the IBM Faculty Award, Okawa foundation faculty research award, RoboCup scientific challenge award and USC Viterbi School of Engineering use-inspired research award. Finally, for his teaching and service, Professor Tambe has received the USC Steven B. Sample Teaching and Mentoring award and the ACM recognition of service award. Recently, he co-founded ARMORWAY, a company focused on risk mitigation and security resource optimization, where he serves on the board of directors. Professor Tambe received his Ph.D. from the School of Computer Science at Carnegie Mellon University.