

An Overview of Recent Application Trends at the AAMAS Conference: Security, Sustainability, and Safety

Manish Jain, Bo An, Milind Tambe

■ *A key feature of the AAMAS conference is its emphasis on ties to real-world applications. The focus of this article is to provide a broad overview of application-focused papers published at the AAMAS 2010 and 2011 conferences. More specifically, recent applications at AAMAS could be broadly categorized as belonging to research areas of security, sustainability, and safety. We outline the domains of applications, key research thrusts underlying each such application area, and emerging trends.*

A key feature of the Autonomous Agents and Multiagent Systems (AAMAS) conference is its emphasis on ties to real-world applications. This emphasis of trying to marry theory and practice at AAMAS goes all the way back to the origins of its predecessor conferences, such as the first International Conference on Autonomous Agents (Johnson 1997). However, the effort to tie research into practical applications got a significant boost with the establishment of the industry track at AAMAS, which was more recently renamed as the innovative applications track.

Over the past few years, within this industry and innovative applications track at AAMAS and other related tracks at sister conferences including Innovative Applications of AI (IAAI), there have been presentations of several successful transitions of key component technologies of agents and multiagent systems. On the one hand, individual agents integrate multiple components and capabilities, for example, planning, learning, reactivity, goal orientedness, and they act autonomously while being situated in their environment — thus facilitating their application in real-world settings. On the other hand, multiagent systems and techniques focused on reasoning about multiple agents reflect the fact that there exist many autonomous agents (cooperative or self-interested) in the real world, and capturing their interaction establishes higher veracity of the model. This

appropriateness of agent and multiagent systems to model complex real-world problems has led to successful transitions of practically applied technologies ranging from belief desire intention (BDI) frameworks, to game-theoretic approaches, to auction frameworks, to biologically inspired approaches. These previously successful applications have been reviewed in the literature and textbooks on multiagent systems (Wooldridge 2009, Shoham and Leyton-Brown 2008).

This article focuses on the more recent efforts to marry research with practical applications that have been reported at AAMAS over the past two years. Specifically, focusing on papers from AAMAS 2010 and AAMAS 2011, we will discuss the three broad areas that have been the focus of the transition from research into practice: security, sustainability, and safety. With respect to security, research at AAMAS has emphasized the use of game-theoretic techniques to schedule limited security resources to protect targets of economic and political importance. For example, ARMOR (Pita et al. 2008; Jain et al. 2010b) schedules checkpoints and canine patrols at the Los Angeles International Airport whereas IRIS (Tsai et al. 2009; Jain et al. 2010b) schedules federal air marshals on board international flights of U.S. air carriers. More game-theoretic scheduling assistants are being designed for other security agencies as well, such as GUARDS (Pita et al. 2011) for scheduling activities conducted by the Transportation Security Administration. GUARDS is being evaluated at an undisclosed airport for potential nationwide deployment. Finally, PROTECT (An et al. 2011a) is in use for scheduling the patrols of the U.S. Coast Guard in the port of Boston and beyond.

Multiagent systems have also been applied to research on the sustainable use of energy resources (Chalkiadakis et al. 2011; Kamboj, Kempton, and Decker 2011; Kok 2010). Sustainable production, delivery, and use of energy is an important challenge of today. One of the ways this can be done is by developing intelligent systems, like smart grids (Vytelingum et al. 2010b; Ramchurn et al. 2011), that can efficiently predict the use of energy and dynamically optimize its delivery. The distributed nature of the energy grid and the individual interests of users make multiagent modeling an appropriate approach for this problem. Multiagent research in this area has primarily focused on developing techniques based on game-theoretic approaches (including coalitional game theory) and auctions that help reduce the usage and wastage of energy (Vytelingum et al. 2010a; Gerdling et al. 2011; Lamparter, Becher, and Fischer 2010; Vandael et al. 2011).

With respect to the final area of application, safety, multiagent systems have been applied for disaster response simulations, air-traffic manage-

ment, evacuation simulations, and related applications (Ramchurn et al. 2010; Dos Santos and Bazzan 2011; Schurr, Picciano, and Marecki 2010). The key advantage is the rich models of individual agents that can be brought to bear in such applications. For instance, from large-scale citywide evacuations to small-scale evacuations of buildings, emergency evacuations are unfortunately a perpetual fixture in society. While commercial evacuation simulation tools have begun to explore agent-based simulations,¹ researchers at AAMAS have brought to bear richer models of agents in such simulations, allowing us to gauge the impact of different environmental, emotional, and informational conditions (Tsai et al. 2011).

Security Applications

The last five years have witnessed the successful application of multiagent systems in reasoning about complex security problems (Basilico, Gatti, and Amigoni 2009; Korzhyk, Conitzer, and Parr 2010; Jain et al. 2010b; Pita et al. 2011; An et al. 2011a). The framework of game theory is becoming very popular in the arena of security, in part due to the increasing need to address the challenges posed by terrorism, drugs, and crime. Yet, limited security resources cannot be everywhere all the time, raising a crucial question of how to best utilize our limited security resources.

Game theory provides a sound mathematical approach for deploying limited security resources to maximize their effectiveness. As mentioned previously, there have been a wide range of actual deployed applications of game theory for security, such as ARMOR and IRIS. This set of applications and associated algorithms has added to the already significant interest in developing multiagent systems applying game theory for security. We now describe the application of game theory to these security domains, partitioned into four key subsections: (1) problem domains; (2) game-theoretic solution approaches; (3) deployments; and (4) research challenges.

Problem Domains

Security scenarios addressed in previous work exhibit the following important characteristics: there is a leader/follower dynamic between the security forces and terrorist adversaries, since the police commit to a security policy first while the adversaries conduct surveillance to exploit any weaknesses or patterns in the security strategies (Tambe 2011). A security policy here refers to some schedule to patrol, check, or monitor the area under protection. There are limited security resources available to protect a very large space of possible targets, so it is not possible to provide complete coverage at all times. Moreover, the tar-

gets in the real world clearly have different values and vulnerabilities in each domain. Additionally, there is uncertainty over many adversary types. For example, the security forces may not know whether they would face a well-funded terrorist or a local gang member or some other threat. Typically, the security forces are interested in a randomized schedule, so that surveillance does not yield predictable patterns; yet they wish to ensure that more important targets have a higher protection and that they guard against an intelligent adversary's adaptive response to their randomized schedule. We now describe some security domains where game-theoretic applications have been successfully deployed.

Los Angeles International Airport (LAX) is the fifth busiest airport in the United States, the largest destination airport in the United States, and serves 60 to 70 million passengers per year (Stevens et. al. 2006). 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 game-theoretic approach focuses 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, foot traffic, or international versus domestic flights. These factors contribute to the differing risk assessments of these eight terminals. The numbers of available vehicle checkpoints and canine units are limited by resource constraints, so the key challenge is to apply game-theoretic algorithms to intelligently allocate these resources — typically in a randomized fashion — to improve their effectiveness while avoiding patterns in the scheduled deployments.

The United States Federal Air Marshal Service (FAMS) places undercover law enforcement personnel aboard flights of U.S. air carriers originating in and departing the United States to dissuade potential aggressors and prevent an attack should one occur.² The exact methods used to evaluate the risks posed by individual flights is not made public by the service, and many factors might influence such an evaluation. For example, flights have different numbers of passengers, and some fly over densely populated areas while others do not. International flights also serve different countries, which may pose different risks. Special events can also change the risks for particular flights at certain times. The scale of the domain is massive. There are currently tens of thousands of commercial

flights scheduled each day, and public estimates state that there are thousands of air marshals (Griffin, Johnston, and Schwarzschild 2008). Air marshals must be scheduled on tours of flights that obey various constraints (for example, the time required to board, fly, and disembark). Simply finding schedules for the marshals that meet all of these constraints is a computational challenge. The 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.

The United States Transportation Security Agency (TSA) is tasked with protecting the nation's transportation systems.³ One set of systems in particular is the more than 400 airports that service approximately 28,000 commercial flights and up to approximately 87,000 total flights⁴ per day. To protect this large transportation network, the TSA employs approximately 48,000 transportation security officers, who are responsible for implementing security activities at each individual airport. While many people are aware of common security activities, such as individual passenger screening, this is just one of many security layers TSA personnel implement to help prevent potential threats. These layers can involve hundreds of heterogeneous security activities executed by limited TSA personnel leading to a complex resource allocation challenge. While activities like passenger screening are performed for every passenger, the TSA cannot possibly run every security activity all the time. Thus, while the resources required for passenger screening are always allocated by the TSA, it must also decide how appropriately to allocate its remaining security officers among the layers of security to protect against a number of potential threats, while facing challenges such as surveillance and an adaptive adversary as mentioned before.

The United States Coast Guard patrols harbors to safeguard the maritime and security interests of the country. Figure 1 shows an example of the types of boats used in patrols conducted by the Coast Guard in Boston. The Coast Guard continues to face a challenging future with an evolving asymmetric threat within the maritime environment both within the Maritime Global Commons but also within the ports and waterways that make up the United States Maritime Transportation System (MTS). The Coast Guard can cover any subset of patrol areas in any patrol schedule. It can also perform many security activities at each patrol area. The challenge for the Coast Guard again is to design a randomized patrolling strategy given that it needs to protect a diverse set of targets along the harbor and the attacker conducts surveillance and is adaptive.



Figure 1. U.S. Coast Guard Conducting a Patrol at the Port of Boston.

Game-Theoretic Solution Approaches

ARMOR, IRIS, GUARDS, and PROTECT, deployed for the security domains mentioned above, build on the game-theoretic foundations to reason about two types of players — the security force and the adversary — to provide a randomized security policy. The algorithms used in these applications build on several years of research reported in the AAMAS conference main track and workshops (Paruchuri et al. 2005; 2006; 2007; Jain, Kiekintveld, and Tambe 2011; Jain et al. 2011). Although the security systems use the newest algorithms from this line of research, we first provide an introduction to key game-theoretic concepts and then describe the solution approaches.

Stackelberg Game

A generic Stackelberg game has two players: a leader, and a follower. These players need not represent individuals but could also be groups that cooperate to execute a joint strategy, such as a police force or a terrorist organization. Each player has a set of possible pure strategies, or the actions that the player can execute. A mixed strategy

allows a player to play a probability distribution over pure strategies. Payoffs for each player are defined over all possible pure-strategy outcomes for both the players. The payoff functions are extended to mixed strategies by taking the expectation over pure-strategy outcomes. The follower can observe the leader's strategy and then act in a way to optimize its own payoffs. Thus, the attacker's strategy in a Stackelberg game is a best response to the leader's strategy.

The most common solution concept in game theory is a Nash equilibrium, which is a profile of strategies for each player in which no player can gain by unilaterally changing to another strategy (Osbourne and Rubinstein 1994). Strong Stackelberg equilibrium is a refinement of Nash equilibrium; it is a form of equilibrium where the leader commits to a strategy first, and the follower provides a best response while breaking ties in favor of the leader.⁵ This Strong Stackelberg equilibrium is the solution concept adopted in security applications (Osbourne and Rubinstein 1994; von Stengel and Zamir 2004; Conitzer and Sandholm 2006; Paruchuri et al. 2008).

	Covered	Uncovered
Defender	5	-20
Attacker	-10	30

Table 1. Example Payoffs in a Security Game for an Attack on One Specific Target.

The Bayesian extension to the Stackelberg game allows for multiple types of players, with each type associated with its own payoff values (Harsanyi and Selten 1972; Paruchuri et al. 2007; 2008). For real-world security domains, we assume that there is only one leader type (for example, only one police force), although there are multiple follower types (for example, multiple groups of adversaries are trying to infiltrate security). Each follower type is represented by a different payoff matrix. The leader does not know the follower's type. The goal is to find the optimal mixed strategy for the leader to commit to, given that each follower type will know the mixed strategy of the leader when choosing its own strategy.

Security Domain Representation

In a security domain, a defender must perpetually defend the site in question, whereas the attacker is able to observe the defender's strategy and attack when success seems most likely. This is appropriately modeled as a Stackelberg game if we map the attacker to the follower's role and the defender to the leader's role (Avenhaus, von Stengel, and Zamir 2002; Brown et al. 2006; Tambe 2011). The actions for the security forces represent the action of scheduling a patrol or checkpoint, for example, a checkpoint at the LAX airport or a federal air marshal scheduled to a flight. The actions for an adversary represent an attack at a target, for example, a terminal at LAX or a flight. The strategy for the leader is a mixed strategy spanning the various possible actions.

We now introduce a further specialization of Stackelberg games prominently used in security applications so far, called security games (Kiekintveld et al. 2009). In a security game, associated with each target are four payoffs defining the possible outcomes for an attack on the target, as shown in table 1. Thus, in this example, if the attacker attacked this target and it was being "covered" by the defender, then the attacker would be unsuccessful and would receive a payoff of -10. On the other hand, the defender would receive a payoff of 5 units in this particular situation. Thus, the payoffs in a security game depend only on the target attacked and whether or not it is covered by the defender. They do not depend on the remaining

aspects of the schedule, such as which set of unattacked targets are covered or which specific defense resource provides coverage.

Algorithms

Over the years, significant research has focused on continually improving the set of algorithms used to solve, or find, the optimal mixed strategy in Bayesian Stackelberg games. These algorithms have been the basis of the deployed applications. This section provides a quick tour of the algorithms that have been used in the deployed applications. While the initial algorithm (Conitzer and Sandholm 2006) provided a linear programming approach, it did not address multiple adversary types, which were important in the first application ARMOR deployed at the LAX airport. Instead, ARMOR relied on DOBSS (Paruchuri et al. 2008), which was designed to scale up for many adversary types. The ERASER algorithm (Kiekintveld et al. 2009) developed next was used in the first version of IRIS. It was capable of scaling up to a large number of defender actions, a requirement for the FAMS domain given the large number of flights the federal air marshals could fly. However, ERASER was not capable of generating schedules over flight tours with more than two flights, thereby motivating the development of ASPEN (Jain et al. 2010a). ASPEN can compute optimal solution over arbitrary tour sizes and scheduling constraints and is the algorithm of choice in the second version of IRIS. GUARDS uses DOBSS again with a novel domain representation (Pita et al. 2011), whereas PROTECT uses further research advances (An et al. 2011a). These new algorithms use mixed-integer linear programming formulations to compute the strong Stackelberg equilibrium. We now describe the ERASER algorithm to give the readers an understanding of the underlying mixed-integer program.

ERASER was the first algorithm that took as input a security game and solved for the optimal coverage vector corresponding to a strong Stackelberg equilibrium strategy for the defender. A coverage vector here implies a probability distribution, which defines the defender's probability of protecting each target. ERASER computes the coverage vector C that maximizes the defender's payoff, subject to the constraints that (1) the attacker will be able to learn this coverage vector C and best-respond to it, and (2) the sum total of coverage across all targets is limited to the number of available resources. The mixed-integer linear program of ERASER is presented in equations 1 through 7. Equations 2 and 3 force the attacker to choose an attack vector $A = (a_i)$ in a way to attack a single target with probability 1. Equation 4 restricts the coverage vector $C = (c_i)$ to probabilities in the range $[0, 1]$, and equation 5 constrains the coverage by the number of available resources.

Equations 6 and 7 compute the defender's pay-

off d and the attacker's payoff k . Here, $U_{\square}(t, C)$ represents the expected utility to the defender when the attacker attacks target t and the defender executes the coverage strategy C . Similarly, $U_{\Psi}(t, C)$ represents the expected utility to the attacker. Z is a large positive constant relative to maximum payoff value. In this way, Equation 7 forces the attacker to compute the optimal strategy to the defender strategy C . Similarly, Equation 6 computes the defender payoff d given the defender's and the attacker's strategy. Taken together, the objective and equations 6 and 7 imply that C and A are mutual best responses in any optimal solution.

$$\max d \quad (1)$$

$$a_t \in \{0,1\} \forall t \in T \quad (2)$$

$$\sum_{t \in T} a_t = 1 \quad (3)$$

$$c_t \in [0,1] \forall t \in T \quad (4)$$

$$\sum_{t \in T} c_t \leq m \quad (5)$$

$$d - U_{\square}(t, C) \leq (1 - a_t) \cdot Z \forall t \in T \quad (6)$$

$$0 \leq k - U_{\Psi}(t, C) \leq (1 - a_t) \cdot Z \forall t \in T \quad (7)$$

Deployments and Results

Having described the foundations of game-theoretic algorithms, we now briefly discuss the game-theoretical models for the applications discussed above. We then evaluate their performance in the real world.

Constructing a Game Model

Instantiating a real-world security domain in a specific Stackelberg game model involves specifying details of three aspects: (i) the possible targets that could be attacked, for example the terminals at LAX; (ii) the defense resources and constraints on how they may or may not be scheduled, for example the number of available canines; and (iii) the payoffs that describe the outcomes of attacks on each target for both the defender and the attacker.

The payoffs provided to the game model define the outcome for both the defender and the attacker in case the attack on a particular target was successful or unsuccessful. These payoffs are provided by domain experts. The payoffs for a security domain, and the exact methods used by the domain experts to arrive at these values are sensitive information. Risk analysts use a detailed set of questions to arrive at the exact payoff values; some of the considerations for payoff values are outlined (Pita et al. 2008; Tsai et al. 2009; Jain et al. 2010b; Tambe 2011).

Evaluation

While ARMOR and IRIS have been successfully deployed for a number of years, evaluating their impact in the real world is not easy. There are also security concerns in making evaluations of security policies publicly available and ethical concerns in not providing the best security possible to a control group. It is important to understand that there is no 100 percent security; all that these game-theoretic algorithms are trying to do is to increase adversary cost and uncertainty. We use at least five types of evaluation in answering the evaluation question: (1) models and simulations in the laboratory; (2) experiments with human subjects; (3) evaluations by domain experts; (4) comparison of game-theoretic strategies with previous deployment strategies; and (5) impersonation of an adversary using teams of security officers to test a security strategy. Finally, researchers continue to look for additional evidence and data that would provide additional evaluation and potential pointers to improvements in the deployment of game-theoretic algorithms.

The key conclusions from comparison against previous deployment techniques are as follows. When compared to human schedulers, we find that the game-theoretic approaches provide more unpredictability. Human schedulers tend to generate predictable patterns, and this weakness of human schedulers was noted in the case of LAX police schedules as well as for FAMS schedules (Murr 2007; U.S. Government Accountability Office 2009). Indeed, human inability to generate random patterns is well studied (Wagenaar 1972). It would seem that the task of scheduling a tour for an air marshal is already quite complex; requiring further that the tours be unpredictable just creates a significant cognitive burden for a human. When compared to a uniform random schedule, game-theoretic schedules perform better since they can account for differing weights of different targets (Jain et al. 2010b). Similarly, game-theoretic scheduling outperforms simple weighted random schedules since game theory explicitly accounts for an adaptive adversary (Jain et al. 2010b).

While these conclusions are supported by our different evaluation techniques, we present one example result from IRIS in terms of our simulations. The results are shown in figure 2. Here, the x -axis shows the number of schedules or flight tours that the federal air marshals could fly and the y -axis shows the expected utility for the defender where higher expected utility is better for the defender. In these experiments, each schedule was a tour composed of one departure flight and one arrival flight. The number of air marshals available to do these flight tours was kept fixed to one in all the experiments. We compare the expected utility from the IRIS strategy with the expected utilities

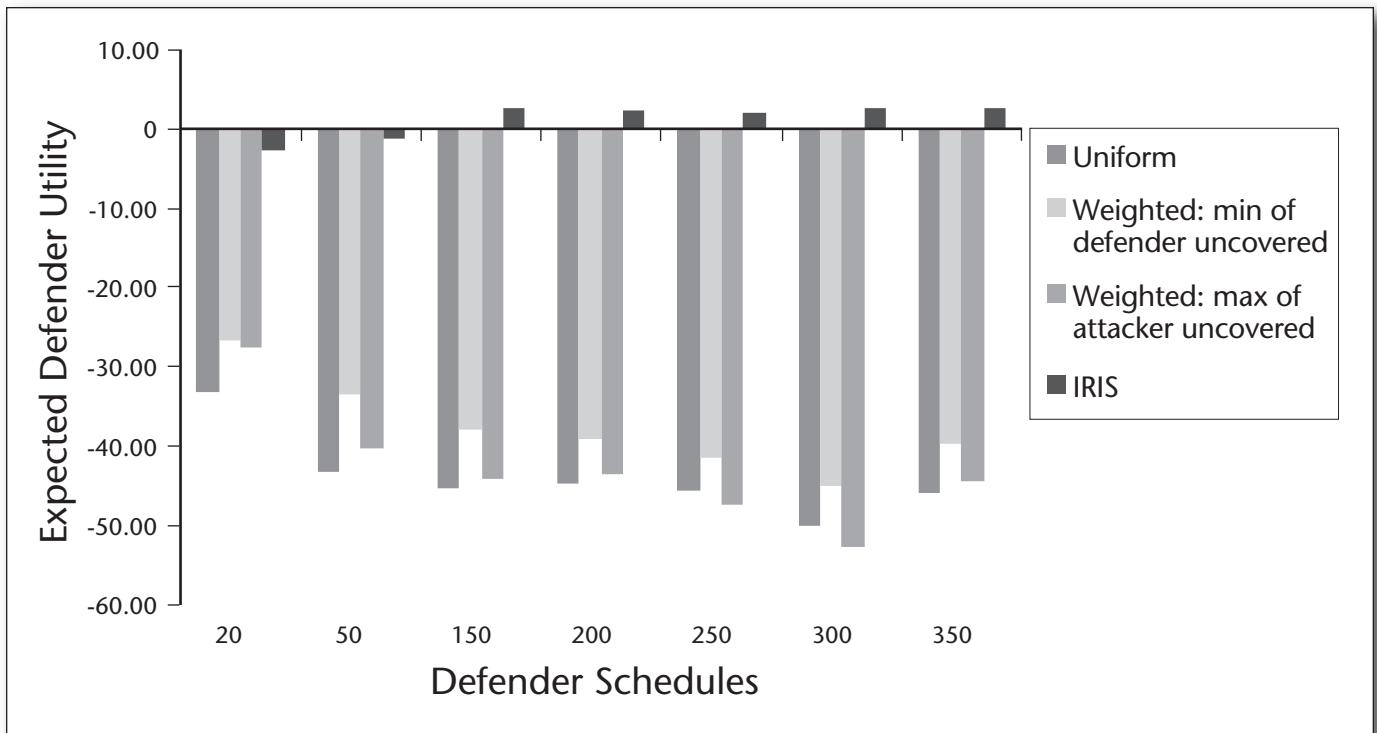


Figure 2. Comparison of IRIS Schedules with Uniform Random and Naïve Weighted Random Techniques.

from uniform and simple weighted random strategies. The results show that the IRIS strategy gives a higher expected utility to the defender in all the settings. Experiments comparing game-theoretic schedules with other weighted randomization techniques, as well as with previously used scheduling practices, also showed that game-theoretic schedules performed better (Jain et al. 2010b). More details of the evaluation can be found in Tambe (2011).

Further evidence of the merit of these software assistants is the adoption and continued use of these tools. ARMOR has been deployed by the LAX police since August 2007 and IRIS began to be used by FAMS in October 2009 after undergoing a non-public internal evaluation. PROTECT is deployed in Boston by the Coast Guard since April 2011 and it is being considered for further deployments. Finally, GUARDS is being evaluated by the TSA at an undisclosed airport.

Research Challenges

While the deployed applications have advanced the state of the art, significant future research remains to be done. At least two of the primary challenges relate to scaleup and robustness. In the following, we highlight key research thrusts in both.

First, with respect to scaleup, algorithms for

security games must address increases in both the number of defender strategies as well as the number of attacker strategies. A key motivating domain for such a scaleup is when defending cities against potential attackers. For example, police in the city of Mumbai have started scheduling a limited number of checkpoints on roads in response to the Mumbai attacks of 2008 (Ali 2009).⁶ Security game algorithms could potentially be used to schedule randomized checkpoints in such settings. In such domains, the strategy space of the defender grows exponentially with the number of available resources, and the strategy space of the attacker grows exponentially with the size of the road network considered. The latest technique to schedule such checkpoints is based on a “double oracle approach,” which does not require the enumeration of the entire strategy space for either of the players (Jain et al. 2011). However, significant further scaleup is required to handle a city of the size of Mumbai.

Second with respect to robustness, our solution algorithms must be robust to the significant uncertainty faced in the domain. For example, while the Stackelberg formulation assumes that the adversary conducts careful surveillance and thus has perfect knowledge of the defender’s mixed strategy, in reality, the adversary’s surveillance may be limited or error prone, requiring security game algo-

gorithms to be robust to such an occurrence (Yin et al. 2011). Similarly, these algorithms must handle the significant uncertainty of the defender's model of the adversary's payoffs (Kiekintveld, Marecki, and Tambe 2011) and uncertainty over the capability of the attacker as well (An et al. 2011b). While there are many such uncertainties, we will briefly highlight work that focuses on the adversary's bounded rationality, which introduces uncertainty in the adversary's decision procedure. In addition to computational game theory, this research also focuses on addressing human biases and cognitive limitations when computing solutions to the game-theoretic models. Thus, this work has led to a new area of research combining behavioral or experimental game theory (Camerer, Ho, and Chong 2004) with security game algorithms. It marries concepts like anchoring bias (Fox and Rottenstreich 2003), prospect theory (Kahneman and Tversky 1979; Hastie and Dawes 2001), and quantal response (McKelvey and Palfrey 1995) with computational game theory, resulting in a novel approach to model real-world players.

As an example of this style of research, an Internet-based computer game inspired by the security situation at LAX was designed to test game-theoretic schedules against human opponents (Yang et al. 2011). Figure 3 shows a screenshot of this game. In this game, the doors represent the terminals that need to be protected. The values shown for the door define the payoffs for the players of the game. The defender was simulated by a pirate, who happened to guard a few of the doors using a pre-specified scheduling strategy. The human subject, analogous to the attacker, was able to observe the pirate for a few observations, and then made a choice as to which door to attack. The outcome of the game was dependent on whether or not a pirate guarded the door chosen by the human subject. The students were given a bankroll at the start of the game; they added to the bank for every success and money was deducted for every failure. The net results of tests with human subjects showed that the standard game-theoretic strategies perform better compared to uniform and naïve weighted random strategies against humans as well (Pita et al. 2010). It also showed that strategies that exploited human biases performed even better than these standard game-theoretic strategies.

Other Applications

This section summarizes some other recent multi-agent applications based on papers from AAMAS 2010 and AAMAS 2011, including those in sustainable energy and safety, which include traffic management, disaster management, air-traffic management, and health. We begin with applications in sustainable energy.

Smart Grid Management/Coordination

Providing sustainable energy is a critical grand challenge facing the world today and it affects all aspects of development. One way to mitigate the challenge is using renewable energy sources, such as hydroelectricity, solar energy, wind energy, wave power, geothermal energy, bioenergy, and tidal power. The other approach is saving energy during its distribution and consumption. In order to efficiently deliver and use energy, energy systems (such as the smart grid) should be able to predict and intelligently respond to the behavior and actions of all electric power users connected to it. In addition, such systems should allow dynamic optimization of system operations and resources. The distributed nature and autonomous behavior of these systems lend themselves to a multiagent methodology.

The function of an electrical grid, as shown in figure 4, is to aggregate multiple networks and power generation companies. Smart grids increase the connectivity, automation, and coordination between suppliers, consumers, and networks that perform either long-distance transmission or local distribution tasks. Given the existence of multiple entities in the smart grid, smart grid management/coordination is crucial for the creation of a robust, intelligent electricity supply network. However, smart grid management is challenging due to the dynamic nature of the grid and the self-interested nature of all the entities participating in the grid.

There are two lines of work related to regulating the energy supply and consumption. One is using different types of storage devices with appropriate (dis)charging strategies, and the other focuses on using different market mechanisms to match supply and consumption. Most smart grid technologies are trying to balance demand and supply in order better to integrate distributed intermittent renewable energy sources. Renewable energy often depends on environmental conditions (for example, wind speeds) that can vary significantly over a short time. Therefore, it is difficult (even impossible) for supply to continuously follow the vagaries of consumer demand. In recent years, efficient low-cost energy storage devices have been widely used to support sustainable energy provisioning and balance demand and supply. While energy usage can be potentially improved by using such devices, it is possible that individual homes charge at the same time according to their own needs. This will cause a higher peak in demand in the electricity market, and in the worst case, it could cause blackouts and infrastructure damage if the total demand were to exceed network capacity. Vytelingum et al (2010b) provide a game-theoretic framework for modeling storage devices in large-scale systems where each storage device is owned

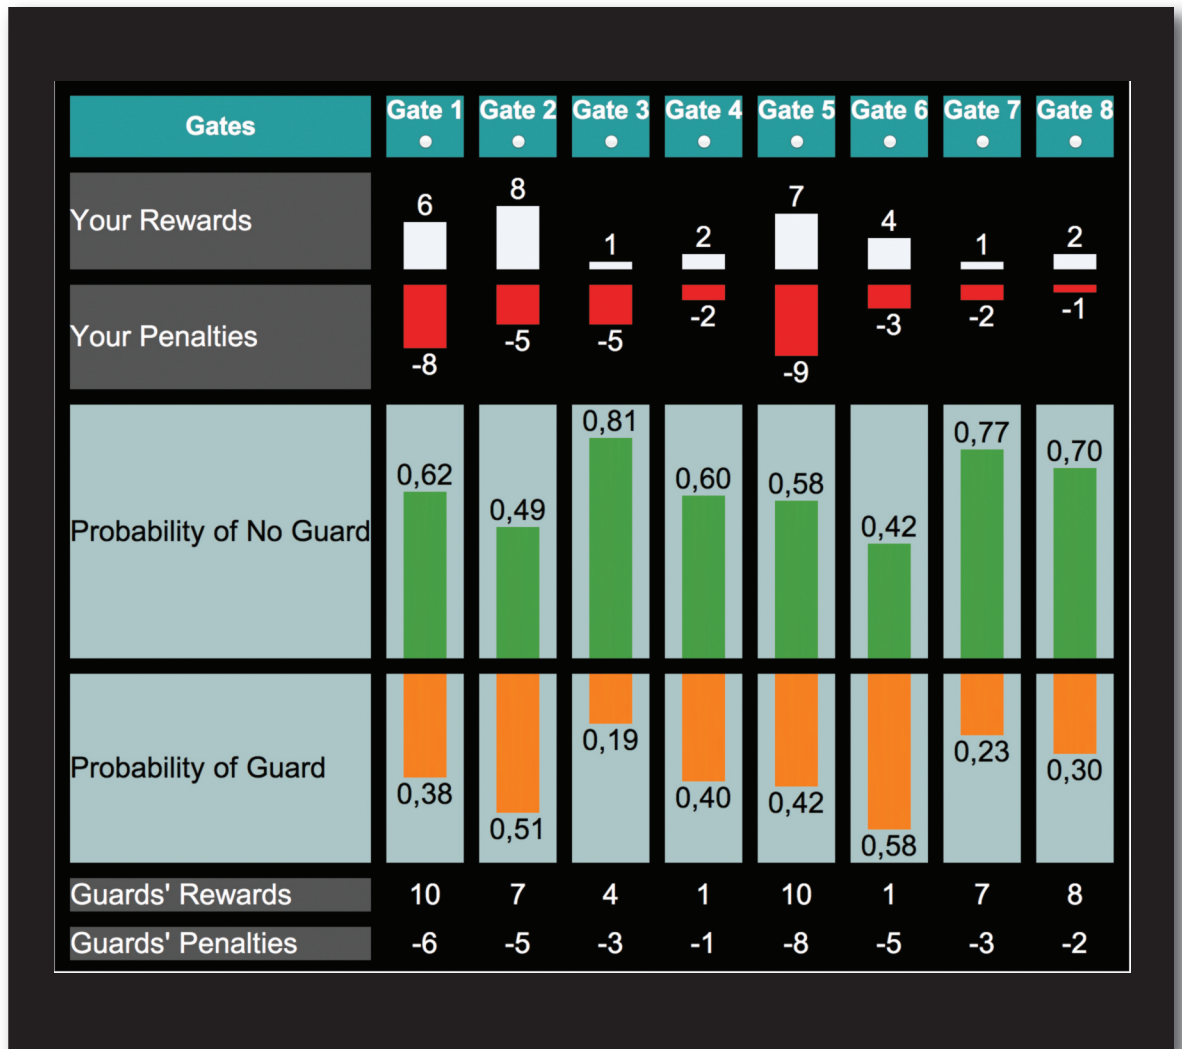


Figure 3. Internet-Based Computer Game Simulating Checkpoints at LAX Airport.

by a self-interested agent that aims to maximize its monetary profit. Under certain assumptions, the proposed agent-based microstorage management strategy allows all storage devices in the system to converge to profitable and efficient behavior. Ramchurn et al (2011) consider more complex deferrable loads and managing the comfort in the home. In addition to microstorage devices at homes, plug-in electric drive vehicles (EDVs), that is, vehicles that use electricity to power at least part of their drive trains, can be integrated into the smart grid and can provide power storage services to the smart grid. Although individual EDVs control too little power to sell in the market at an individual level, a large group of EDVs may form an aggregate or coalition that controls enough power to meaningfully sell in the various electricity markets. A prototype system has been deployed in the

real world and it is shown that a vehicle has an incentive to participate in coalitions (Kamboj, Kempton, and Decker 2011).

In addition to using small storage devices, different market mechanisms have been proposed to regulate the energy consumption and supply in different scenarios at the smart grid. The most widely used mechanism is auction (Vytelingum et al. 2010a; Gerding et al. 2011; Lamparter, Becher, and Fischer 2010). For example, a continuous double auction (CDA) mechanism with agents' trading strategies is used to balance different traders in the market (Vytelingum et al. 2010a). An online allocation mechanism is proposed for electric vehicle owners to bid for power and the time window for charging (Gerding et al. 2011). Since the smart grid market is a complex dynamic market, those mechanisms often can guarantee truthfulness and effi-

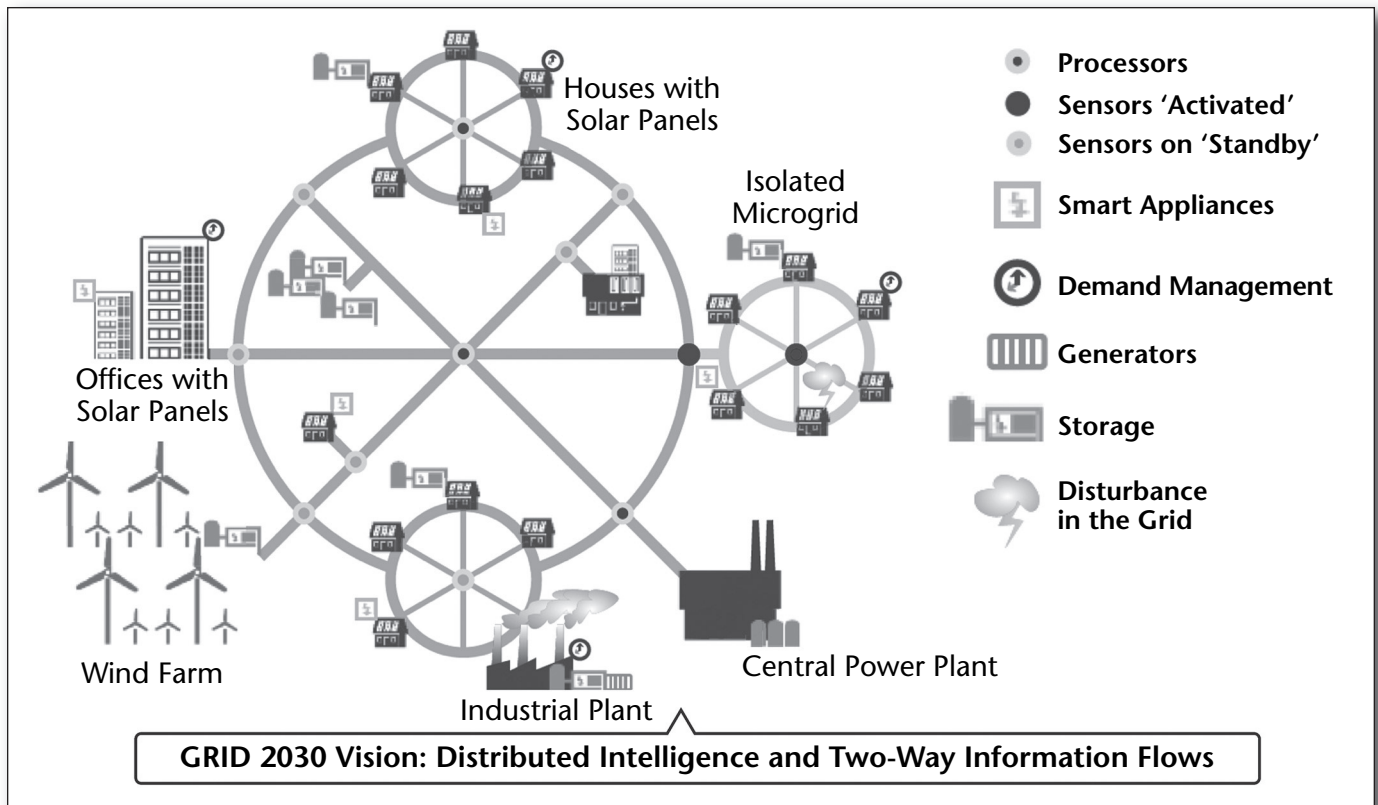


Figure 4. Smart Grid Management Would Require Coordination Between These Multiple Entities.

ciency only under very strong assumptions. Alternatively, some heuristic approaches are used for supply and demand matching in electricity networks. PowerMatcher (Kok, Warmer, and Kamphuis 2005; Kok 2010) is a general-purpose coordination mechanism for balancing demand and supply in clusters of distributed energy resources. The heart of the system is an electronic market on which local control agents negotiate using strategies based on short-term microeconomics. Different scheduling strategies have also been proposed for reducing imbalance costs in a smart grid due to unpredictable changes in production and consumption (Vandael et al. 2011).

Minimize Building Energy Consumption

In addition to saving energy in its distribution stage through techniques such as smart grid management, minimizing the energy consumption is also important to achieve the goal of sustainable energy. In the United States, about 40 percent of energy consumption is from buildings, of which 25 percent is associated with heating and cooling at an annual cost of \$40 billion. Furthermore, on an annual basis, buildings in the United States consume 73 percent of its electricity. Multiagent technology, together with existing IT solutions/infra-

structure, has been identified as a promising approach to achieve greater energy efficiency in buildings (Rogers et al. 2011). Kwak et al (2011) present a novel multiagent system based on distributed coordination reasoning under uncertainty for sustainability called SAVES. SAVES is capable of generating plans to minimize the energy consumption while satisfying the comfort level of occupants in the buildings. SAVES is currently being tested in simulations, but the goal is to deploy in buildings for a proof of concept demonstration.

In related research, Rogers and colleagues (Rogers et al. 2011) address the challenge of adaptively controlling a home heating system in order to minimize cost and carbon emissions within a smart grid. The designed energy management agent learns the thermal properties of the home and uses Gaussian processes to predict the environmental parameters over the next 24 hours, allowing it to adjust the timing of heater use in order to satisfy preferences for comfort while minimizing cost and carbon emissions.

Multiagent Traffic Management

The increasing demand for mobility in our society has led to the more serious problem of traffic con-

gestion. Traffic causes air pollution and decrease in speed, which is directly linked to energy (for example, fuel) consumption. A more efficient use of the available transportation infrastructure is necessary and this relates closely to multiagent systems as many problems in traffic management and control are inherently distributed (Bazzan 2009). AI and multiagent techniques have been proposed for traffic management (see Klugl and Bazzan [2012]; Bazzan [2009] for a survey). A reservation-based intersection control approach with a communication protocol is proposed in Dresner and Stone (2008). In the reservation-based approach, autonomous guided vehicles report information (for example, the velocity, direction, maximum/minimum acceleration) to intersection managers, which later decide on rejection or acceptance of requests based on their knowledge of other vehicles.

Pulter, Schepperle, and Bohm (2011) quantify the fuel consumption with existing agent-based approaches for intersection control and propose an agent-based mechanism for intersection control, with minimization of fuel consumption as an explicit design objective. Simulations show that the proposed mechanism could reduce fuel consumption by up to 26 percent and waiting time by up to 98 percent, compared to traffic lights.

Agent technology has also been used to offer support for commercial aviation transportation. An air traffic control system based on adjustable autonomy has been created to support the optimal allocation of tasks (functions) between the system and the human operators (Schurr, Picciano, and Marecki 2010). The system includes (1) a simulation environment, (2) a DFAS algorithm for providing adjustable autonomy strategies, and (3) the agents for executing the strategies and measuring system efficiency. An initial pilot study shows some promising results.

Disaster Management

Efficient and effective disaster management is becoming increasingly important for the world given the major disasters in recent years, ranging from natural disasters such as the Tohoku earthquake, Haiti earthquake, Asian tsunami, and hurricane Katrina, to the man-made disasters such as the 9/11 attack and the London terrorist attacks. Disaster management is a significantly challenging research topic. Agents face a highly dynamic and uncertain environment, which makes it difficult for agents to make the optimal decisions in the long term. For disasters, new tasks may continually appear or disappear, and thus timely response is crucial. In addition, there are often a large number of complex rescue tasks, each requiring multiple agents (or other entities) to act together since agents often have limited capabilities.

Efficient task (resource) allocation is a critical factor in any successful disaster management. Since agents' capabilities are often limited, coordination is often necessary through forming teams for coalitions. Effective coordination ensures that tasks are allocated so that efforts are not duplicated and all resources (including time) are used in an efficient way. Coordination can be done in either a centralized way or a distributed way. Distributed mechanisms have many useful properties (for example, robustness, flexibility, lower overheads) and are more appropriate for complex dynamic environments. A variety of distributed coordination mechanisms have been proposed, for example, DCOP-based approaches (Scerri et al. 2005; Dos Santos and Bazzan 2011) and max-sum algorithms (Ramchurn et al. 2010; Farinelli et al. 2008).

Another important line of work is simulating pedestrian behavior in disaster scenarios. Agent-based simulation allows for each pedestrian to be modeled as an autonomous entity. Under this model, pedestrians are represented as agents capable of perceiving and interacting with their environment as well as other agents. Recent research on agent modeling includes the ESCAPES system (Tsai et al. 2011), which is concerned with the interactions between agents and the resulting group dynamics. Additionally, ESCAPES focuses on domains including airports, malls, and museums. To accurately represent these types of environments, ESCAPES considered it particularly important to model the influence of families, emotional contagion, social comparison, and spread of knowledge, which past work had not cohesively addressed.

Health Applications

The use of modern information and communication technologies can aid to decrease both the cost of prenatal health-care services and also the load of medical practitioners. One key example at AAMAS is the SUAP project, which provides a multiagent system for supporting and monitoring prenatal care (Nunes et al. 2010). SUAP uses agent technology to manage health-care records, to act as a clinical decision support system, and to handle the logistics of high-risk pregnancy cases. The first version of the SUAP system was deployed in July 2009 and it was composed of the core functionalities that provided the prenatal care systematization.

Conclusions and Future Work

Key applications of multiagent systems highlighted at the AAMAS conference illustrate that researchers are making huge strides in the areas of security, sustainability, and safety. Many applications are already in use, with more in the pipeline.

These applications have also led to some fundamental research challenges in many different areas of multiagent systems. In terms of future directions, some of the domains where such agent-based approaches could and should have a significant impact can be found in the grand challenges listed by the National Academy of Engineering.⁷ For example, one of the challenges is to secure cyberspace, where current agent-based solution methods can be extended and applied. Researchers in cyber security have already started investigating game-theoretic approaches similar to the ones mentioned earlier for routing packets and scheduling packet inspections (Alpcan 2010; Kodialam and Lakshman 2003). Similarly, agent-based techniques can be very useful in addressing another grand challenge of improving urban infrastructure. Indeed, as mentioned earlier, disaster response and energy management in urban settings are active areas of research focused on improving urban societies. A third important example research area where agent technology can have a very significant impact is advanced personalized learning. This requires the development of an agent or multiagent system that can identify individual preferences and aptitudes of each student, such that instruction can be tailored to a student's individual needs. Indeed, while significant research challenges remain to be addressed, the trajectory of use-inspired research at the AAMAS conference is extremely promising; and thus as it has already begun to do, research in agents and multiagent systems could have significant societal impact in the near future.

Acknowledgement

We would like to thank Dr. Sarvapali Ramchurn for the example of the smart grid.

Notes

1. See www.legion.com/legion-software.
2. See www.tsa.gov/lawenforcement/programs/fams.shtm.
3. See www.tsa.gov and www.tsa.gov/what_we_do/layers/index.shtm.
4. See Air Traffic Control: By the Numbers; www.zfwnat-ca.com.
5. The leader can always induce the follower strictly to break ties in favor of the leader by perturbing the leader's strategy by an infinitesimal amount (von Stengel and Zamir 2004)
6. A map illustrating the terrorist attacks of 2008 in Mumbai can be found at kids.britannica.com/comptons/art-153427/Armed-terrorists-attacked-several-sites-in-Mumbai-India-on-Nov.
7. See www.engineeringchallenges.org.

References

Ali, S. Ahmed. 2009. Rs 18l Seized in Naka Bandi at Vile

Parle. *Times of India* (Aug. 4). (articles.timesofindia.indiatimes.com/2009-08-04/mumbai/28170719_1_nakabandi-rs-18l-mumbai-police).

Alpcan, T. 2010. *Network Security: A Decision and Game-Theoretic Approach*. Cambridge, UK: Cambridge University Press.

An, B.; Pita, J.; Shieh, E.; Tambe, M.; Kiekintveld, C.; and Marecki, J. 2011a. GUARDS and PROTECT: Next Generation Applications of Security Games. *ACM SIGecom Exchanges* 10(1): 31–34.

An, B.; Tambe, M.; Ordóñez, F.; Shieh, E.; and Kiekintveld, C. 2011b. Refinement of Strong Stackelberg Equilibria in Security Games. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*. Menlo Park, CA: AAAI Press.

Avenhaus, R.; von Stengel, B.; and Zamir, S. 2002. Inspection Games. In *Handbook of Game Theory*, ed. R. J. Aumann and S. Hart, volume 3, 1947–1987. Amsterdam: North-Holland.

Basilico, N.; Gatti, N.; and Amigoni, F. 2009. Leader-Follower Strategies for Robotic Patrolling in Environments with Arbitrary Topologies. In *Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems*, 500–503. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Bazzan, A. L. 2009. Opportunities for Multiagent Systems and Multiagent Reinforcement Learning in Traffic Control. In *Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems*, 342–375. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Brown, G.; Carlyle, M.; Salmeron, J.; and Wood, K. 2006. Defending Critical Infrastructure. *Interfaces* 36(6): 530–544.

Camerer, C. F.; Ho, T.; and Chong, J. 2004. A Cognitive Hierarchy Model of Games. *The Quarterly Journal of Economics* 119(3): 861–898.

Chalkiadakis, G.; Robu, V.; Kota, R.; Rogers, A.; and Jennings, N. 2011. Cooperatives of Distributed Energy Resources for Efficient Virtual Power Plants. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems*, 787–794. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Conitzer, V., and Sandholm, T. 2006. Computing the Optimal Strategy to Commit To. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, 82–90. New York: Association for Computing Machinery.

Dos Santos, F., and Bazzan, A. L. 2011. Towards Efficient Multiagent Task Allocation in the RoboCup Rescue: A Biologically-inspired Approach. *Autonomous Agents and Multi-Agent Systems* 22(3): 465–486.

Dresner, K., and Stone, P. 2008. A Multiagent Approach to Autonomous Intersection Management. *Journal of Artificial Intelligence Research* 31: 591–656.

Farinelli, A.; Rogers, A.; Petcu, A.; and Jennings, N. R. 2008. Decentralised Coordination of Low-Power Embedded Devices Using the Max-Sum Algorithm. In *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems*, 639–646. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Fox, C. R., and Rottenstreich, Y. 2003. Partition Priming

- in Judgment Under Uncertainty. *Psychological Science* 14(3): 195–200.
- Gerding, E.; Robu, V.; Stein, S.; Parkes, D.; Rogers, A.; and Jennings, N. 2011. Online Mechanism Design for Electric Vehicle Charging. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems*, 811–818. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Griffin, D.; Johnston, K.; and Schwarzschild, T. 2008. Sources: Air Marshals Missing from Almost all Flights. *CNN* (March 25).
- Harsanyi, J., and Selten, R. 1972. A Generalized Nash Solution for Two-Person Bargaining Games with Incomplete Information. *Management Science* 18(5, Part 2): 80–106.
- Hastie, R., and Dawes, R. M. 2001. *Rational Choice in an Uncertain World: The Psychology of Judgment and Decision Making*. Thousand Oaks, CA: Sage Publications.
- Jain, M.; Kardes, E.; Kiekintveld, C.; Ordóñez, F.; and Tambe, M. 2010a. Security Games with Arbitrary Schedules: A Branch and Price Approach. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*. Menlo Park, CA: AAAI Press.
- Jain, M.; Tsai, J.; Pita, J.; Kiekintveld, C.; Rathi, S.; Tambe, M.; and Ordóñez, F. 2010b. Software Assistants for Randomized Patrol Planning for the LAX Airport Police and the Federal Air Marshals Service. *Interfaces* 40(4): 267–290.
- Jain, M.; Kiekintveld, C.; and Tambe, M. 2011. Quality-Bounded Solutions for Finite Bayesian Stackelberg Games: Scaling Up. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Jain, M.; Korzhyk, D.; Vanek, O.; Conitzer, V.; Pechoucek, M.; and Tambe, M. 2011. A Double Oracle Algorithm for Zero-Sum Security Games on Graphs. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Johnson, W. L. (ed). 1997. *Proceedings of the First International Conference on Autonomous Agents*. New York: Association for Computing Machinery.
- Kahneman, D., and Tversky, A. 1979. Prospect Theory: An Analysis of Decision Under Risk. *Econometrica* 47(2): 263–292.
- Kamboj, S.; Kempton, W.; and Decker, K. S. 2011. Deploying Power Grid-Integrated Electric Vehicles as a Multi-agent System. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems*, 13–20. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Kiekintveld, C.; Jain, M.; Tsai, J.; Pita, J.; Tambe, M.; and Ordóñez, F. 2009. Computing Optimal Randomized Resource Allocations for Massive Security Games. In *Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems*, 689–696. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Kiekintveld, C.; Marecki, J.; and Tambe, M. 2011. Approximation Methods for Infinite Bayesian Stackelberg Games: Modeling Distributional Payoff Uncertainty. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Klügl, F., and Bazzan, A. L. C. 2012. Agent-Based Modeling and Simulation. *AI Magazine* 33(3).
- Kodialam, M., and Lakshman, T. 2003. Detecting Network Intrusions via Sampling: A Game Theoretic Approach. In *Proceedings of the 22nd Annual Joint Conference of the IEEE Computer and Communications Societies*. Piscataway, NJ: Institute of Electrical and Electronics Engineers.
- Kok, K. 2010. Multi-agent Coordination in the Electricity Grid, from Concept Towards Market Introduction. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, 1681–1688. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Kok, J. K.; Warmer, C. J.; and Kamphuis, I. G. 2005. PowerMatcher: Multiagent Control in the Electricity Infrastructure. In *Proceedings of the 4th International Conference on Autonomous Agents and Multiagent Systems*, 75–82. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Korzhyk, D.; Conitzer, V.; and Parr, R. 2010. Complexity of Computing Optimal Stackelberg Strategies in Security Resource Allocation Games. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence*, 805–810. Menlo Park, CA: AAAI Press.
- Kwak, J.; Varakantham, P.; Tambe, M.; Klein, L.; Jazizadeh, F.; Kavulya, G.; Gerber, B. B.; and Gerber, D. J. 2011. Towards Optimal Planning for Distributed Coordination Under Uncertainty in Energy Domains. Paper presented at the 2nd International Workshop on Agent Technologies for Energy Systems, Taipei, Taiwan, May 2.
- Lamparter, S.; Becher, S.; and Fischer, J.-G. 2010. An Agent-Based Market Platform for Smart Grids. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 1689–1696. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- McKelvey, R. D., and Palfrey, T. R. 1995. Quantal Response Equilibria for Normal Form Games. *Games and Economic Behavior* 10(1): 6–38.
- Murr, A. 2007. The Element of Surprise. *Newsweek National News*. 28 September.
- Nunes, I.; Choren, R.; Nunes, C.; Fábri, B.; Silva, F.; Carvalho, G.; and de Lucena, C. J. P. 2010. Supporting Prenatal Care in the Public Health Care System in a Newly Industrialized Country. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, 1723–1730. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Osbourne, M. J., and Rubinstein, A. 1994. *A Course in Game Theory*. Cambridge, MA: The MIT Press.
- Paruchuri, P.; Tambe, M.; Ordóñez, F.; and Kraus, S. 2005. Safety in Multiagent Systems by Policy Randomization. Paper presented at the International Workshop on Safety and Security in Multi-Agent Systems. Utrecht, The Netherlands, 26 July.
- Paruchuri, P.; Tambe, M.; Ordóñez, F.; and Kraus, S. 2006. Security in Multiagent Systems by Policy Randomization. In *Proceedings of the 5th International Conference on Autonomous Agents and Multiagent Systems*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Paruchuri, P.; Pearce, J. P.; Tambe, M.; Ordóñez, F.; and Kraus, S. 2007. An Efficient Heuristic Approach for Security Against Multiple Adversaries. In *Proceedings of the 6th International Conference on Autonomous Agents and Multiagent Systems*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- 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 *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems*, 895–902. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Pita, J.; Jain, M.; Western, C.; Portway, C.; Tambe, M.; Ordóñez, F.; Kraus, S.; and Paruchuri, P. 2008. Deployed ARMOR Protection: The Application of a Game-Theoretic Model for Security at the Los Angeles International Airport. In *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems*, 125–132. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

- Pita, J.; Jain, M.; Ordóñez, F.; Tambe, M.; and Kraus, S. 2010. Robust Solutions to Stackelberg Games: Addressing Bounded Rationality and Limited Observations in Human Cognition. *Artificial Intelligence Journal* 174(15): 1142–1171.
- Pita, J.; Kiekintveld, C.; Tambe, M.; Steigerwald, E.; and Cullen, S. 2011. GUARDS Game Theoretic Security Allocation on a National Scale. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Pulter, N.; Schepperle, H.; and Bohm, K. 2011. How Agents Can Help Curbing Fuel Combustion — A Performance Study of Intersection Control. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems*, 795–802. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Ramchurn, S.; Farinelli, A.; Macarthur, K.; Polukarov, M.; and Jennings, N. 2010. Decentralised Coordination in RoboCup Rescue. *The Computer Journal* 53(9): 1–15.
- Ramchurn, S. D.; Vytelingum, P.; Rogers, A.; and Jennings, N. R. 2011. Agent-Based Control for Decentralised Demand Side Management in the Smart Grid. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems*, 5–12. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Rogers, A.; Maleki, S.; Ghosh, S.; and Nicholas R, J. 2011. Adaptive Home Heating Control Through Gaussian Process Prediction and Mathematical Programming. Paper presented at the Second International Workshop on Agent Technologies for Energy Systems. Taipei, Taiwan, 2 May.
- Scerri, P.; Farinelli, A.; Okamoto, S.; and Tambe, M. 2005. Allocating Tasks in Extreme Teams. In *Proceedings of the 4th International Conference on Autonomous Agents and Multiagent Systems*, 727–734. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Schurr, N.; Picciano, P.; and Marecki, J. 2010. Function Allocation for NextGen Airspace via Agents. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, 1731–1738. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Shoham, Y., and Leyton-Brown, K. 2008. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge, UK: Cambridge University Press.
- Stevens, D.; Hamilton, T.; Schaffer, M.; Dunham-Scott, D.; Medby, J. J.; Chan, E. W.; Gibson, J.; Eisman, M.; Mesic, R.; Kelley, C. T. Jr.; Kim, J.; LaTourrette, T.; Riley, K. J. 2006. Implementing Security Improvement Options at Los Angeles International Airport. Technical Report DB-499-1-LAWA. Santa Monica, CA: Rand Corporation.
- Tambe, M. 2011. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge, UK: Cambridge University Press.
- Tsai, J.; Rathi, S.; Kiekintveld, C.; Ordóñez, F.; and Tambe, M. 2009. IRIS: a Tool for Strategic Security Allocation in Transportation Networks. In *Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems*, 37–44. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Tsai, J.; Fridman, N.; Bowring, E.; Brown, M.; Epstein, S.; Kaminka, G.; Marsella, S.; Ogden, A.; Rika, I.; Sheel, A.; Taylor, M.; Wang, X.; Zilka, A.; and Tambe, M. 2011. ESCAPES — Evacuation Simulation with Children, Authorities, Parents, Emotions, and Social comparison. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- U.S. Government Accountability Office. 2009. *Federal Air Marshal Service Has Taken Actions to Fulfill Its Core Mission and Address Workforce Issues, but Additional Actions Are Needed to Improve Workforce Survey*. Washington, DC: Government Printing Office.
- Vandael, S.; Craemer, K. D.; Boucke, N.; Holvoet, T.; and Deconinck, G. 2011. Decentralized Coordination of Plug-in Hybrid Vehicles for Imbalance Reduction in a Smart Grid. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems*, 803–810. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- von Stengel, B., and Zamir, S. 2004. Leadership with Commitment to Mixed Strategies. Technical Report LSE-CDAM-2004-01, CDAM Research Report. London: London School of Economics and Political Science.
- Vytelingum, P.; Ramchurn, S. D.; Voice, T. D.; Rogers, A.; and Jennings, N. R. 2010a. Trading Agents for the Smart Electricity Grid. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, 897–904. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Vytelingum, P.; Voice, T. D.; Ramchurn, S. D.; Rogers, A.; and Jennings, N. R. 2010b. Agent-Based Micro-Storage Management for the Smart Grid. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, 39–46. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Wagenaar, W. A. 1972. Generation of Random Sequences by Human Subjects: A Critical Survey of Literature. *Psychological Bulletin* 77(1): 65–72.
- Wooldridge, M. 2009. *An Introduction to MultiAgent Systems*, 2nd ed. New York: Wiley.
- Yang, R.; Kiekintveld, C.; Ordóñez, F.; Tambe, M.; and John, R. 2011. Improving Resource Allocation Strategy Against Human Adversaries in Security Games. In *Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Yin, Z.; Jain, M.; Tambe, M.; and Ordóñez, F. 2011. Risk-Averse Strategies for Security Games with Execution and Observational Uncertainty. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence*. Menlo Park, CA: AAAI Press.

Manish Jain is currently a Ph.D. candidate in the Department of Computer Science at the University of Southern California (USC). He is a part of the TEAMCORE Research group. His work is on the applications of game-theoretic and large-scale optimization techniques for Homeland Security, including deployed software tools for the Federal Air Marshal Service and Los Angeles World Airports police. He has coauthored papers on the subject of security games that have been presented at AAMAS and AAAI conferences, in both the main track and the industry track. He has also organized tutorials on the topic of security games at both AAMAS and AAAI conferences. His work was recently award the RIST Prize by the Military Operations Research Society. His work published at Interfaces was nominated as a finalist for the EURO excellence in practice award. He has also won the best research assistant award at the University of Southern California.

Bo An is a postdoctoral researcher in the Department of Computer Science at the University of Southern California. He received his Ph.D. degree in computer science from the University of Massachusetts, Amherst. Prior to that, he received B.Sc. and M.Sc. degrees in computer science from Chongqing University, China. His research interests include artificial intelligence, multiagent systems, game theory, automated negotiation, resource allocation, distributed systems, and optimization. He has published more than 20 papers at conferences including AAMAS, IJCAI, and AAAI and

**Conference Chair**

Phil McCarthy
pmccarthy@decooda.com
Decooda

Program Chairs

Chutima Boonthum-Denecke
chutima.boonthum@gmail.com
Hampton University

G. Michael Youngblood
Michael.Youngblood@parc.com
PARC, a Xerox company

Special Tracks Coordinator

Bill Eberle
weberle@tntech.edu
Tennessee Technological University

Important Dates

Paper Submission: Nov. 19, 2012
Author Notification: Jan. 21, 2013
Abstract Only Submission: Feb. 11, 2013
Abstract only Notification: Feb. 17, 2013
Camera-Ready Copy: Feb. 25, 2013

The 26th International FLAIRS Conference

TradeWinds Island Resorts, St. Pete Beach, Florida, USA
May 22-24, 2013

The 26th International FLAIRS Conference (FLAIRS-26) will be held May 22-24, 2013 at the TradeWinds Island Resorts in St. Pete Beach, Florida, USA. FLAIRS-26 continues a tradition of presenting and discussing artificial intelligence research in a convivial atmosphere within a beautiful setting. Events will include invited speakers, special tracks, discussion panels, presentations of papers and posters, and impromptu presentations. As always, there will be a Best Paper award and a Best Poster award. In addition, because FLAIRS has a rich tradition of encouraging student authors, there will be a Best Student Paper award for the best paper written primarily by a student. Submissions are now invited for full papers (6 pages), short papers to be presented as a poster (4 pages), and poster abstracts (250 words). The proceedings will be published by the AAAI. The conference is hosted by the Florida Artificial Intelligence Research Society in cooperation with AAAI.

Topics of interest are in all areas of AI, including but not limited to

- **Foundations:** Knowledge Representation, Cognitive Modeling, Perception, Reasoning & Programming, Constraints, Satisfiability, Search, Learning
- **Architectures:** Agents and Distributed AI, Intelligent User Interfaces, Natural Language Systems, Information Retrieval, Robotics
- **Applications:** Aviation and Aerospace, Education, Entertainment, Medicine, Management and Manufacturing, World Wide Web
- **Implications:** Philosophical Foundations, Social Impact and Ethics, Evaluation of AI Systems, Teaching AI
- **Special Tracks:** Numerous special tracks offer opportunities for focused interaction. All special track papers are published in the proceedings.

<http://www.flairs-26.info/>

In cooperation with the Association for the Advancement of Artificial Intelligence

journals including *JAAMAS* and *IEEE Transactions*. He is cochair for the AAAI 2012 Spring Symposium on Game Theory for Security, Sustainability, and Health and the recipient of the 2010 IFAAMAS Victor Lesser Distinguished Dissertation Award.

Milind Tambe is a professor of computer science and industrial and systems engineering at the University of Southern California. He leads the TEAMCORE Research Group at USC, with research focused on agent-based and multiagent systems. He is a fellow of AAAI (Association for Advancement of Artificial Intelligence) and recipient of the ACM (Association for Computing Machinery) Autonomous Agents Research Award. He is also the recipient of the Homeland Security Award from the Christopher Columbus Fellowship Foundation, a special commendation given by the Los Angeles World Airports police from the city of Los Angeles, USC Viterbi School of Engineering use-inspired research award, Okawa foundation faculty research award, the RoboCup scientific challenge award, USC Steven B. Sample Teaching and Mentoring award, and the ACM recognition of service award. Tambe received his Ph.D. from the School of Computer Science at Carnegie Mellon University.

CALL FOR PAPERS

The Twenty Sixth International Conference on Industrial, Engineering & Other Applications of Applied Intelligent Systems (IEA/AIE-2013) – Amsterdam, the Netherlands – June 17 – 21, 2013

IEA/AIE-2013 Conference website:
<http://iea-aie2013.few.vu.nl/>

Sponsored by: International Society of Applied Intelligence (ISAI) and *In Cooperation with:* Association for the Advancement of Artificial Intelligence (AAAI); Association for Computing Machinery (ACM/SIGART, SIGKDD); Austrian Association for Artificial Intelligence (OeGAI); Catalan Association for Artificial Intelligence (ACIA); Delft University of Technology; European Neural Network Society (ENNS); International Neural Network Society (INNS); Italian Artificial Intelligence Association (AI*IA); Japanese Society for Artificial Intelligence (JSAI); Lithuanian Computer Society—Artificial Intelligence Section (LIKS-AIS); Slovenian Artificial Intelligence Society (SLAIS); Spanish Society for Artificial Intelligence (AEPIA); Taiwanese Association for Artificial Intelligence (TAAI); Taiwanese Association for Consumer Electronics (TACE); Texas State University-San Marcos, VU University Amsterdam.

IEA/AIE-2013 continues the tradition of emphasizing Industrial, Engineering and Other applications of Applied Intelligent Systems' technology in solving real-life problems. Paper submission is required by **November 9, 2012**. Submission instructions and additional details may be obtained from the website <http://iea-aie2013.few.vu.nl> or from the Organizing Chair, Dr. Tibor Bosse (tbosse@few.vu.nl).