Bayesian Stackelberg Games and their Application for Security at Los Angeles International Airport

MANISH JAIN, JAMES PITA, MILIND TAMBE, FERNANDO ORDONEZ,
*PRAVEEN PARUCHURI, **SARIT KRAUS
University of Southern California, CA, USA ({manish.jain, jpita, tambe, fordon}@usc.edu)
* Intelligent Automation Inc., Rockville, MD, USA (pparuchuri@i-a-i.com)
** Bar-Ilan University, Israel (sarit@macs.biu.ac.il)

Many multiagent settings are appropriately modeled as Stackelberg games [Fudenberg and Tirole 1991; Paruchuri et al. 2007], where a leader commits to a strategy first, and then a follower selfishly optimizes its own reward, *considering the action chosen by the leader*. Stackelberg games are commonly used to model attacker-defender scenarios in security domains [Brown et al. 2006] as well as in patrolling [Paruchuri et al. 2007; Paruchuri et al. 2008]. For example, security personnel patrolling an infrastructure commit to a patrolling strategy first, before their adversaries act taking this committed strategy into account. Indeed, Stackelberg games are being used at the Los Angeles International Airport to schedule security checkpoints and canine patrols [Murr 2007; Paruchuri et al. 2008; Pita et al. 2008a]. They could potentially be used in network routing, pricing in transportation systems and many other situations [Korilis et al. 1997; Cardinal et al. 2005].

Although the follower in a Stackelberg game is allowed to observe the leader's strategy before choosing its own strategy, there is often an advantage for the leader over the case where both players must choose their moves simultaneously. To see the advantage of being the leader in a Stackelberg game, consider the game with the payoff as shown in Table I. The leader is the row player and the follower is the column player. The only pure-strategy Nash equilibrium for this game is when the leader plays *a* and the follower plays *c* which gives the leader a payoff of 2. However, if the leader commits to a mixed strategy of playing *a* and *b* with equal (0.5) probability, then the follower will play *d*, leading to an expected payoff for the leader of 3.5.

	c	d
a	2,1	4,0
b	1,0	3,2

Table I. Payoff table for example normal form game.

In real scenarios, the leader may face uncertainty over multiple types of followers, which

---

can be modeled using a Bayesian Stackelberg game. The motivation behind [Paruchuri et al. 2008] is to determine the optimal mixed strategy for a leader to commit to in a Bayesian Stackelberg game, i.e. a Stackelberg game where the leader may face one of multiple follower types. Such a Bayesian game arises in security domains because for example, when patrolling a region, a security robot may only have uncertain knowledge about different robber types it may face. Unfortunately, the problem of finding optimal leader strategies in Bayesian Stackelberg games is NP-hard, and transforming the game into a normal form game using Harsanyi transformation [Harsanyi and Selten 1972] loses the compactness of the game.

In [Paruchuri et al. 2008], an efficient exact method for finding the optimal leader strategy in such games, known as DOBSS (Decomposed Optimal Bayesian Stackelberg Solver) has been introduced. This method has three key advantages. First, the method allows for a Bayesian game to be expressed compactly without requiring conversion to a normal-form game via the Harsanyi transformation. Second, the method requires only one mixed-integer linear program (MILP) to be solved, rather than a set of linear programs as in [Conitzer and Sandholm 2006] thus leading to a further performance improvement. Third, it directly searches for an optimal leader strategy, rather than a Bayes-Nash equilibrium, allowing it to find high-reward non-equilibrium strategies (thus exploiting the advantage of being the leader). DOBSS is orders of magnitude faster than the previously published exact methods. Compared to previous heuristical methods [Paruchuri et al. 2007], it is not only faster but also avoids the significant problems of infeasibility faced by those methods.

Stackelberg games and DOBSS are at the heart of the ARMOR system *deployed* for the past 8 months at the Los Angeles International Airport to schedule security personnel [Murr 2007; Paruchuri et al. 2008; Pita et al. 2008a; Paruchuri et al. 2008; Pita et al. 2008b] and randomize canine patrols. For our application at the Los Angeles Airport, the leader has 784 actions and there may be up to 4 adversary types each with 8 actions. While DOBSS can solve such problems in a reasonable time (about 80 seconds), as we go on to future work involving larger scale applications, scaling up DOBSS continues to be an important challenge.

Acknowledgements. The research was supported by Department of Homeland Security through Center for Risk and Economic Analysis of Terrorism Events (CREATE) center at USC.

REFERENCES

BROWN, G., CARLYLE, M., SALMERN, J., AND WOOD, K. 2006. *Defending Critical Infrastructure*. Interfaces.

CARDINAL, J., LABBE, M., LANGERMAN, S., AND PALOP, B. 2005. Pricing of geometric transportation networks. In *17th Canadian Conference on Computational Geometry*.

CONITZER, V. AND SANDHOLM, T. 2006. Computing the optimal strategy to commit to. In *EC*.

FUDENBERG, D. AND TIROLE, J. 1991. *Game Theory*. MIT Press.

HARSANYI, J. C. AND SELTEN, R. 1972. A generalized Nash solution for two-person bargaining games with incomplete information. *Management Science 18,* 5, 80–106.

KORILIS, Y. A., LAZAR, A. A., AND ORDA, A. 1997. Achieving network optima using stackelberg routing strategies. In *IEEE/ACM Transactions on Networking*.

MURR, A. 2007. Random checks. *Newsweek National News*, http://www.newsweek.com/id/43401.

PARUCHURI, P., MARECKI, J., PEARCE, J., TAMBE, M., AND KRAUS, S. 2008. An efficient exact algorithm for solving bayesian stackelberg games. In *AAMAS 2008*.

PARUCHURI, P., PEARCE, J. P., MARECKI, J., TAMBE, M., ORDONEZ, F., AND KRAUS, S. 2008. Efficient algorithms to solve bayesian stackelberg games for security applications. In *AAAI*.

PARUCHURI, P., PEARCE, J. P., TAMBE, M., ORDÓÑEZ, F., AND KRAUS, S. 2007. An efficient heuristic approach for security against multiple adversaries. In *AAMAS*.

PITA, J., JAIN, M., ORDONEZ, F., PORTWAY, C., TAMBE, M., WESTERN, C., PARUCHURI, P., AND KRAUS, S. 2008a. Armor security for los angeles international airport. In *AAMAS Industry Track*.

PITA, J., JAIN, M., ORDONEZ, F., PORTWAY, C., TAMBE, M., WESTERN, C., PARUCHURI, P., AND KRAUS, S. 2008b. Armor security for los angeles international airport. In *AAAI Demo*.