



Stackelberg security games: Computing the shortest-path equilibrium



Julio B. Clempner^{a,*}, Alexander S. Poznyak^b

^a Center for Economics, Management and Social Research, National Polytechnic Institute, Lauro Aguirre 120, col. Agricultura, Del. Miguel Hidalgo 11360, Mexico City, Mexico

^b Center for Research and Advanced Studies, National Polytechnic Institute, Av. IPN 2508, Col. San Pedro Zacatenco 07360, Mexico City, Mexico

ARTICLE INFO

Article history:

Available online 17 January 2015

Keywords:

Security games
Strong Stackelberg equilibrium
Shortest-path games
Lyapunov equilibrium
Lyapunov games
Complexity analysis
Repeated Markov games

ABSTRACT

In this paper we consider a game theory approach for representing a real-world attacker–defender Stackelberg security game. In this novel approach the behavior of an ergodic system (repeated stochastic Markov chain game) is represented by a Lyapunov-like function non-decreasing in time. Then, the representation of the Stackelberg security game is transformed in a potential game in terms of Lyapunov. We present a method for constructing a Lyapunov-like function: the function replaces the recursive mechanism with the elements of the ergodic system seeking to drive the underlying finite-state Stackelberg game to an equilibrium point along a least expected cost path. The proposed method analyzes both pure and mixed stationary strategies to find the strong Stackelberg equilibrium. Mixed strategies are found when the resources available for both the defender and the attacker are limited. Lyapunov games model how players are likely to behave in one-shot games and allow finishing during the game whether the applied best-reply strategy (pure or mixed) provides the convergence to a shortest-path equilibrium point (or not). We prove that Lyapunov games truly fit into the framework for deterministic and stochastic shortest-path security games. The convergence rate of the proposed method to a Stackelberg/Nash equilibrium is analyzed. Validity of the proposed method is successfully demonstrated both theoretically and by a simulated experiment.

© 2014 Elsevier Ltd. All rights reserved.

1. Introduction

1.1. Brief review

In the last years, there has been a substantial effort in the application of Stackelberg game-theoretic approaches in the security arena (Pita et al., 2009; Pita, Jain, Ordoñez, Tambe, & Kraus, 2010; Jain, Kardes, Kiekintveld, Ordoñez, & Tambe, 2010), in which security agencies implement patrols and checkpoints to protect targets from criminal attacks. The motivation for solving the security problem is related to the fact that it is not possible to cover all targets at the time (there exists limited defense resources). The game-theoretic approach employed successful to solve security games is that of a Stackelberg game between a defender (leader) and an attacker (follower).

Our goal is to analyze a two-player Stackelberg security game, a defender and an attacker, set of possible targets in a class of ergodic controllable finite Markov chains (Poznyak, Najim, & Gomez-Ramirez, 2000). The defender commits first to a given strategy in

the Markov chains game. He/she decides upon a randomized policy of defending the targets, possibly with limited defense resources. Then, the attacker realizes a Nash solution (chooses a target so as to maximize its expected utility). The optimal actions of the defender is then chosen to minimize its payoff. The main concern about Stackelberg games is as follows: the highest leader payoff is obtained when the followers always reply in the best possible way for the leader (this payoff is at least as high as any Nash payoff).

The solutions to security games are the Strong Stackelberg Equilibria (SSE) (Korzhyk, Yin, Kiekintveld, Conitzer, & Tambe, 2011c), where both the defender and attacker choose a best-reply strategy and, in addition, the follower breaks ties optimally for the leader. The defender's task is to pick an optimal (pure or mixed) strategy given that the attacker is going to play a best-reply to it (Basilico, Gatti, and Villa (2010), Dickerson et al. (2010), Letchford, MacDermed, Conitzer, Parr, and Isbell (2012), Letchford and Conitzer (2013), Letchford and Vorobeychik (2013), Korzhyk, Conitzer, and Parr (2011a, 2011b). In this sense, we present a method for constructing a Lyapunov-like function which replaces the recursive mechanism for one-shot strategies and converges to a Lyapunov equilibrium point. It is the case that a Lyapunov equilibrium is a SSE (Clempner & Poznyak, 2011, 2013).

* Corresponding author.

E-mail addresses: julio@clempner.name (J.B. Clempner), apoznyak@ctrl.cinvestav.mx (A.S. Poznyak).

1.2. Related work

There were several real-world security situations that happened to provide the impetus for the work to be described in this paper (Conitzer & Korzhyk, 2011; Conitzer & Sandholm, 2006; Hernández, Barrientos, & del Cerro, 2014; Jain et al., 2010; Jain, Kiekintveld, & Tambe, 2011; Kiekintveld et al., 2009; Korzhyk, Conitzer, & Parr, 2010; Korzhyk et al., 2011c; Letchford & Conitzer, 2010; Letchford, Conitzer, & Munagala, 2009; Letchford et al., 2012; Paruchuri et al., 2008; Pita et al., 2010; Pita et al., 2009; von Stengel & Zamir, 2010). In these efforts, it has become increasingly clear that the lack of perfect information and deception play critical roles in the development of useful security strategies. Also, there are important computing game-theoretic real-world applications to security (Tambe, 2011): (1) ARMOR security system, including the randomized allocation of checkpoints and canine units at the Los Angeles International Airport (Pita et al., 2009; Jain et al., 2010), (2) IRIS, deployed by the Federal Air Marshals Service (Tsai, Rathi, Kiekintveld, Ordoñez, & Tambe, 2009; Jain et al., 2011), (3) GUARDS: Game-theoretic Unpredictable and Randomly Deployed Security (An et al., 2011a; Pita, Tambe, Kiekintveld, Cullen, & Steigerwald, 2011) under development for the Transportation Security Administration (TSA), (4) PROTECT: Port Resilience Operational/ Tactical Enforcement to Combat Terrorism under development for the United States Coast Guard (An et al., 2011a; Shieh et al., 2012; Tambe, 2011; Fang, Jiang, & Tambe, 2013) and, (5) RaPtoR (Varakantham, Lau, & Yuan, 2013) used to generate schedules for patrolling public transit systems (LA Metro and Singapore Metro system). These games are security games between a defender (allocates defensive resources), and an attacker (decide on targets to attack): (1) the defender considers what the target (best-reply) of the attacker is; (2) then, holding the attacked target fixed, the defender picks a quantity that minimizes its payoff; (3) the attacker actually observes this and in equilibrium picks the expected quantity that maximizes its payoff as a response.

Among the different of approaches that are used to tackle security problems uncertainty has become one of the principal challenges (domains where defenders may face unanticipated disruptions of their schedules): bounded rationality in computing optimal strategies (Yang, Kiekintveld, Ordonez, Tambe, & John, 2011; Nguyen, Yang, Azaria, Kraus, & Tambe, 2013), the possibility that attackers have unknown capability constraints that restrict the set of targets they can feasibly attack (An, Tambe, Ordonez, Shieh, & Kiekintveld, 2011b; An et al., 2012), noise in the defender's execution of the suggested mixed strategy and/or the observations made by an attacker can be noise (Yin, Jain, Tambe, & Ordonez, 2011; Yin & Tambe, 2012), uncertainty handled using Markov decision process using a Bayesian Stackelberg game model (Yin et al., 2011; Yin & Tambe, 2012; Jain et al., 2011; Jiang, Yin, Zhang, Tambe, & Kraus, 2013). A particular case of Stackelberg security games considers the problem of multi-robot patrolling against intrusions around a given area with the existence of an attacker attempting to penetrate into the area (Agmon, Kaminka, & Kraus, 2011; Basilio et al., 2010; Hernández et al., 2014).

A remarkable solution to Stackelberg security games was proposed by Trejo, Clempner, and Poznyak (2015b, in press), consisting of a leader and multiple followers, using the extraproximal method for computing the Stackelberg/Nash equilibria in a class of ergodic controlled finite Markov chains games. The extraproximal method is a two-step iterated procedure for solving the game. For solving the problem the original game is formulated in terms of coupled nonlinear programming problems implementing the Lagrange principle. In addition, the Tikhonov's regularization method is employed to ensure the convergence of the cost-functions to a Stackelberg/Nash equilibrium point.

Since part of this work deals with deploying a game-theoretic approach related to Shortest-Path (SP) games, it is important to discuss existing literature that has addressed a similar challenge. Security games are a class of games related to SP games where a “minimizing” player seeks to drive a finite-state dynamic system to reach a terminal state along a least expected cost path, and a “maximizer” player that looks for maximizing the expected total cost interfering with the minimizer's progress. Patek and Bertsekas (Patek, 1997; Patek & Bertsekas, 1999) analyzed the case of two players, where one player seeks to drive the system to termination along a least cost path and the other seeks to prevent termination altogether. They did not assume non-negativity of the costs, the analysis is much more complicated than the corresponding analysis of Kushner and Chamberlain (1969) and generalize (to the case of two players) those for stochastic SP problems (Bertsekas & Tsitsiklis, 1991). Patek and Bertsekas (1999) proposed alternative assumptions which guarantee that, at least under optimal policies, the terminal state is reached with probability one. They considered undiscounted additive cost games without averaging, admitting that there are policies for the minimizer which allow the maximizer to prolong the game indefinitely at infinite cost to the minimizer. Under assumptions which generalize deterministic SP problems, they established (i) the existence of a real-valued equilibrium cost vector achievable with stationary policies for the opposing players and (ii) the convergence of value iteration and policy iteration to the unique solution of Bellman's equation (Bellman, 1962). The results of Patek and Bertsekas did imply the results of Shapley (Shapley, 1953), as well as those of Kushner and Chamberlain (1969). Because of their assumptions relating to termination, they were able to derive stronger conclusions than those made by Kumar and Shiao (1981) for the case of a finite state space. In a subsequent work, Patek (2001) reexamined the stochastic SP formulation in the context of Markov decision processes with an exponential utility function. In a context related with Petri nets, Clempner (2006) change the traditional cost function by a trajectory-tracking function which is a Lyapunov-like function, introducing the concept of Lyapunov games and making a significant difference in the conceptualization of the problem domain, allowing the replacement of the Nash equilibrium by the Lyapunov equilibrium point in game theory (Clempner & Poznyak, 2011, 2013).

1.3. Main contribution

In this paper are investigated the Stackelberg security games in the context of SP problems represented by a Lyapunov game theoretical approach (Clempner, 2006; Clempner & Poznyak, 2011, 2013). We represent the game using an ergodic class of controlled Markov chain. In Lyapunov games, the Lyapunov-like functions are used to compute players' utilities applying local-optimal policies that produces a monotonic progress towards a Lyapunov equilibrium point. In this work, we present a method for constructing a Lyapunov-like function that explains the behavior of players in a repeated Markov chain game. The Lyapunov-like function replaces the recursive mechanism with the elements of the ergodic system that model how players are likely to behave in one-shot games. Lyapunov games allow finishing during the game whether the applied discrete or mixed best-reply strategy provides the convergence to an equilibrium point. The paper confronts fundamental questions of how a player should compute his/her pure and mixed stationary local-optimal strategies in order to converge to a strong Stackelberg equilibrium. The model presented in the formulation of the mixed strategies is not linear. However, we introduce a joint c -variable that enables our approach to be reformulated as a linear programming problem with linear restrictions. Then, the problem of finding mixed stationary strategies for computing the SP of the

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات