Automatica 97 (2018) 200–213

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Linear quadratic mean field Stackelberg differential games[☆]

Jun Moon[a,*], Tamer Başar[b]

[a] School of Electrical and Computer Engineering, Ulsan National Institute of Science and Technology (UNIST), Ulsan, 44919, South Korea
[b] Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, 1308 West Main Street, Urbana, IL, 61801, USA

ARTICLE INFO

Article history:
Received 1 August 2017
Received in revised form 29 April 2018
Accepted 27 July 2018
Available online 22 August 2018

Keywords:
Mean field theory
Stackelberg and Nash games
Stochastic optimal control
Forward–backward stochastic differential equations

ABSTRACT

We consider linear–quadratic mean field Stackelberg differential games with the adapted open-loop information structure of the leader. There are one leader and N followers, where N is arbitrarily large. The leader holds a dominating position in the game in the sense that the leader first chooses and then announces the optimal strategy to which the N followers respond by playing a Nash game, i.e., choosing their optimal strategies noncooperatively and simultaneously based on the leader's observed strategy. In our setting, the followers are coupled with each other through the mean field term included in their cost functions, and are strongly influenced by the leader's open-loop strategy included in their cost functions and dynamics. From the leader's perspective, he is coupled with the N followers through the mean field term included in his cost function. To circumvent the complexity brought about by the coupling nature among the leader and the followers with large N, which makes the use of the direct approach almost impossible, our approach in this paper is to characterize an approximated stochastic mean field process by solving a local optimal control problem of the followers with leader's control taken as an exogenous stochastic process. We show that for each fixed strategy of the leader, the followers' local optimal decentralized strategies lead to an ϵ-Nash equilibrium. The paper then solves the leader's local optimal control problem, as a nonstandard constrained optimization problem, with constraints being induced by the approximated mean field process determined by Nash followers (which also depend on the leader's control). We show that the local optimal decentralized controllers for the leader and the followers constitute an (ϵ_1, ϵ_2)-Stackelberg–Nash equilibrium for the original game, where ϵ_1 and ϵ_2 both converge to zero as $N \to \infty$. Numerical examples are provided to illustrate the theoretical results.

1. Introduction

Recently, games with a large number of agents have been studied extensively within the mean field game framework. In this setting, the individual agents interact with each other through a mean field term included in the individual cost functions and/or controlled stochastic differential (or dynamic) equations, which captures the average behavior of the agents. Since there is a large number of agents, complexity issues arise from the dimension of the state space and the heterogeneity of the agents. In such cases, computing Nash equilibria for the corresponding game using direct methods as discussed in standard texts for dynamic games, such as Başar and Olsder (1999), may be cumbersome and complicated. To resolve this difficulty, the mean field analysis has been introduced to obtain the best estimate of the actual mean field behavior, which leads to optimal decentralized strategies that are functions of local information and constitute an ϵ-Nash equilibrium (Huang, Caines, & Malhame, 2007; Li & Zhang, 2008). Lasry and Lions (2007) developed independently a different approach to obtain the mean field equilibrium, which entails solving coupled forward–backward partial differential equations where the former is related to optimal control with the Hamilton–Jacobi–Bellman equation, and the latter is related to the mean field distribution with the Fokker–Planck–Kolmogorov equation. Both of these approaches are built on a platform that utilizes the fact that the impact of the individual agents on the mean field behavior becomes negligible when the number of agents goes to infinity.

There are various applications of mean field games. In Yin, Mehta, Meyn, and Shanbhag (2012), the problem of a large number of coupled oscillators has been formulated within the mean field game framework, where the decentralized optimal strategies were

[☆] This research was supported in part by the National Research Foundation of Korea (NRF) Grant funded by the Ministry of Science and ICT, South Korea (NRF-2017R1A5A1015311), in part by Institute for Information & communications Technology Promotion (IITP) grant funded by the Korea government (MSIT), South Korea (No. 2018-0-00958), and in part by the Office of Naval Research (ONR), USA, MURI Grant N00014-16-1-2710. The material in this paper was partially presented at the 54th IEEE Conference on Decision and Control, Dec. 15–18, 2015. This paper was recommended for publication in revised form by Associate Editor Dario Bauso under the direction of Editor Ian R. Petersen.

* Corresponding author.
E-mail addresses: junmoony@unist.ac.kr (J. Moon), basar1@illinois.edu (T. Başar).

https://doi.org/10.1016/j.automatica.2018.08.008
0005-1098/© 2018 Elsevier Ltd. All rights reserved.

characterized to obtain an ϵ-Nash equilibrium. The problem of charging of a large number of Plug-in Electric Vehicles was studied independently in Couillet, Perlaza, Tembine, and Debbah (2012) and Ma, Callaway, and Hiskens (2013). In addition to this, there are several application domains of mean field games; for example, engineering, finance, economics with a large number of firms, systems biology, etc; see Bauso, Zhang, and Papachristodoulou (2017), Cardaliaguet and Lehalle (2017), Couillet et al. (2012), Festa and Gottlich (2017), Firoozi and Caines (2016), Huang, Caines, and Malhame (2003), Kizilkale and Malhame (2014), Lasry and Lions (2007), Weintraub, Benkard, and Van Roy (2008) and Zhu, Tembine, and Başar (2011).

The mean field game framework has been extended to various different settings. Linear–quadratic mean field games were studied in Ahuja (2016), Bardi and Priuli (2014), Bensoussan, Sung, Yam, and Yung (2014), Huang and Li (2018), Huang et al. (2007) and Li and Zhang (2008). Risk-sensitive and/or robust mean field games were introduced and discussed in Moon and Başar (2014, 2017), Tembine (2015) and Tembine, Zhu, and Başar (2014). Mean field consensus games were studied in Nourian, Caines, Malhame, and Huang (2013), and mean field games with Markov jump parameters were considered in Moon and Başar (2016) and Wang and Zhang (2012). Mean field games were studied within the probabilistic approach in Bardi and Fischer (2018) and Carmona and Delarue (2013).

In contrast to the types of mean field games referenced above, Huang (2010) and Nguyen and Huang (2012) considered the situation when there are one major agent and a large number of minor agents, where each minor agent is affected by the major agent's Brownian motion through its state included in the minor agent's dynamics and cost function. This can be viewed as strong influence of the major agent on minor agents, in view of which unlike Carmona & Delarue (2013), Huang et al. (2007), Li & Zhang (2008) and Moon & Başar (2017), stochastic mean field approximation was introduced. Specifically, due to the strong influence of the leader, the approximated mean field coupling term is no longer deterministic, but a stochastic process driven by the Brownian motion of the leader. In Huang (2010), the state augmentation method was developed via the strong law of large numbers to characterize the best stochastic mean field process when the followers are heterogeneous with K distinct models. In Nguyen and Huang (2012), fixed point analysis was applied to obtain similar results as in Huang (2010) when the dynamics and costs of the followers are parametrized within a continuum set. These two different approaches lead to (different) decentralized optimal strategies for the individual agents that constitute an (different) ϵ-Nash equilibrium. The nonlinear counterpart of mean field games with major and minor agents was studied in Bensoussan, Chau, and Yam (2015a) and Nourian and Caines (2013). Also, probabilistic mean field games with major and minor players were studied in Carmona and Zhu (2014) and Huang, Wang, and Wu (2016), and finite-state mean field games with major and minor players were considered in Carmona and Wang (2016).

We should mention that mean field games with major and minor players, as discussed above are *Nash* games. That is, each agent determines his optimal strategy noncooperatively and all simultaneously, which lead to ϵ-Nash equilibria, and there is no hierarchy of decision making between the agents. On the other hand, if one wants to model a certain hierarchical structure in mean field games, the corresponding problem can be formulated by employing the *Stackelberg* setting. Classical Stackelberg games are hierarchical decision making problems, where there is a leader with a dominant position over the follower (Von Stackelberg, 1952). The leader first announces his optimum strategy by taking into account the rational reactions of the followers. The follower then chooses his optimal strategy based on the leader's strategy, and finally the leader comes back and implements his announced strategy, thus generating his action. When there is such a solution, the resulting optimum strategies for the leader and the follower form a Stackelberg equilibrium (Başar & Olsder, 1999).

Stackelberg differential and dynamic games have been studied extensively in the literature since 1970, and detailed expositions can be found in Başar, Bensoussan, and Sethi (2010), Başar and Olsder (1999), Başar and Selbuz (1979), Bensoussan, Chen, and Sethi (2015b), Freiling, Jank, and Lee (2001), Papavassilopoulos and Cruz (1979), Simaan and Cruz (1973), Yong (2002), and the references therein. Stackelberg games have a wide range of applications. In a communication network, there is a single service provider and a (large) number of users, where the service provider sets the usage price(s) for the Nash followers (Başar & Srikant, 2002). Moreover, in the smart grid, the optimal demand response management can be studied within the framework of Stackelberg games, where the utility companies are leader, and the users are followers (Maharjan, Zhu, Zhang, Gjessing, & Başar, 2013).

1.1. Problem statement and main contributions

In this paper, we consider mean field Stackelberg differential games when there are one leader and a large number, say N, of followers. The leader *globally* dominates over the followers for the entire duration in the sense that before the start of the game he[1] chooses and then announces his strategy to the N followers who play a Nash game. The N number of Nash followers choose their optimal strategies noncooperatively and simultaneously based on the leader's observed strategy. In this paper, the information structures for both the leader and the followers are taken to be *adapted open-loop*, that is, the information structure that defines admissible controls for the leader and the followers is the filtration generated by each agent's initial condition and Brownian motion (the leader, of course, knows everything that the followers know).[2] Moreover, in this setting, the followers are coupled with each other through the mean field term included in each follower's cost function, and is strongly influenced by the leader's strategy included in each agent's cost function and dynamics. From the leader's perspective, he is coupled with the N followers through the mean field term included in his cost function. We also consider the heterogeneous case of the followers with K distinct models, that is, each follower belongs to a finite model set $\mathcal{K} = \{1, 2, \ldots, K\}$. We note that the classical Stackelberg (differential and dynamics) games with the leader(s) and Nash followers without the mean field framework were studied in Nie, Chen, and Fukushima (2006) and Simaan and Cruz (1973), where (necessary or sufficient) conditions were provided for characterization of the Stackelberg equilibrium.

Since there is a large number of followers, complexity issues arise from the mean field coupling term and heterogeneity of Nash followers. In addition to this, solving the leader's optimal control problem becomes complicated, since it depends on a large number of Nash followers. Therefore, computing an exact Stackelberg–Nash solution is cumbersome and complicated. To circumvent this difficulty, our approach in this paper is to apply the stochastic mean field approximation to characterize the best estimate of the actual mean field behavior.

We first consider the mean field Nash game for the N followers given an arbitrary strategy of the leader. We solve a local optimal control problem of the followers with leader's control taken as an exogenous stochastic process. We characterize the best estimate

[1] In this paper, we use a "he" when referring to a specific player (the leader and the followers). It could equally have been a "she".

[2] The precise notion of (adapted) open-loop information structure for (stochastic) Stackelberg games can be found in Başar and Olsder (1999), Bensoussan et al. (2015b) and Yong (2002).