EXTENDING THE REALM OF OPTIMIZATION FOR COMPLEX SYSTEMS: UNCERTAINTY, COMPETITION, AND DYNAMICS DOE DE:SC0003879 ARRA ($ 1,031,997)

UDAY V. SHANBHAG (PI)
T. BASAR, S. P. MEYN, AND P. G. MEHTA

This research grant was motivated by a range of problems drawn from optimization, control, and game theory in regimes complicated by dynamics, uncertainty, and nonsmoothness. Apart from the providing summer and travel support for the PIs, this grant has contributed towards the support of several graduate students, both through tuition assistantships and conference support. Briefly, this grant has allowed for presenting ongoing and completed research at a range of meetings including the Informs annual meetings, Optimization meetings (SIAM Conference on Optimization and the Mathematical Programming Symposium), Control theory meetings (American Control Conference and the IEEE Conference on Decision and Control), amongst others. In the remainder of this report, we will outline the various students and briefly review their research activity.

Huibing Yin. Huibing Yin graduated from the department of Mechanical Science and Engineering in 2011 and was advised by PIs Mehta and Shanbhag with PIs Basar and Meyn serving on the doctoral committee. Huibing's early research focused on the development of analytical and algorithmic tools for distributed computation of Nash equilibria while much of Huibing's doctoral research examined the question of synchronization in mean-field oscillator games, with an emphasis on learning and efficiency analysis. Here is a brief summary of his research papers.

(J1) Huibing Yin, Uday V. Shanbhag, Prashant G. Mehta: Nash Equilibrium Problems With Scaled Congestion Costs and Shared Constraints. IEEE Trans. Automat. Contr. 56(7): 1702-1708 (2011)

Abstract: We consider a class of convex Nash games where strategy sets are coupled across agents through a common constraint and payoff functions are linked via a scaled congestion cost metric. A solution to a related variational inequality problem provides a set of Nash equilibria characterized by common Lagrange multipliers for shared constraints. While this variational problem may be characterized by a non-monotone map, it is shown to admit solutions, even in the absence of restrictive compactness assumptions on strategy sets. Additionally, we show that the equilibrium is locally unique both in the primal space as well as in the larger primal-dual space. The existence statements can be generalized to accommodate a piecewise-smooth congestion metric while affine restrictions, surprisingly, lead to both existence and global uniqueness guarantees. In the second part of the technical note, we discuss distributed computation of such equilibria in monotone regimes via a distributed iterative Tikhonov regularization (ITR) scheme. Application to a class

Date: October 8, 2013.

of networked rate allocation games suggests that the ITR schemes perform better than their two-timescale counterparts.

(J2) Huibing Yin, Prashant G. Mehta, Sean P. Meyn, Uday V. Shanbhag: Synchronization of Coupled Oscillators is a Game. IEEE Trans. Automat. Contr. 57(4): 920-935 (2012)

Abstract: The purpose of this paper is to understand phase transition in noncooperative dynamic games with a large number of agents. Applications are found in neuroscience, biology, and economics, as well as traditional engineering applications. The focus of analysis is a variation of the large population linear quadratic Gaussian (LQG) model of Huang et al. 2007, comprised here of a controlled N-dimensional stochastic differential equation model, coupled only through a cost function. The states are interpreted as phase angles for a collection of heterogeneous oscillators, and in this way the model may be regarded as an extension of the classical coupled oscillator model of Kuramoto. A deterministic PDE model is proposed, which is shown to approximate the stochastic system as the population size approaches infinity. Key to the analysis of the PDE model is the existence of a particular Nash equilibrium in which the agents "opt out' of the game, setting their controls to zero, resulting in the "incoherence' equilibrium. Methods from dynamical systems theory are used in a bifurcation analysis, based on a linearization of the partial differential equation (PDE) model about the incoherence equilibrium. A critical value of the control cost parameter is identified: above this value, the oscillators are incoherent; and below this value (when control is sufficiently cheap) the oscillators synchronize. These conclusions are illustrated with results from numerical experiments.

(J3) Huibing Yin, Prashant G. Mehta, Sean P. Meyn, Uday V. Shanbhag: Learning in mean-field oscillator games. (To appear) IEEE Trans. Automat. Control.

Abstract: The purpose of this paper is to show how insight obtained from a mean-field model can be used to create an architecture for approximate dynamic programming (ADP) for a certain class of games comprising of a large number of agents. The general technique is illustrated with the aid of a mean-field oscillator game model introduced in our prior work [2]. The states of the model are interpreted as the phase angles for a collection of non-homogeneous oscillators, and in this way the model may be regarded as an extension of the classical coupled oscillator model of Kuramoto. The paper introduces ADP techniques for design and adaptation (learning) of approximately optimal control laws for this model. For this purpose, a parameterization is proposed, based on an analysis of the mean-field PDE model for the game. In an offline setting, a Galerkin procedure is introduced to choose the optimal parameters while in an online setting, a steepest descent algorithm is proposed. The paper provides detailed analysis of the optimal parameter values as well as the Bellman error with both the Galerkin approximation and the online algorithm. Finally, a phase transition result is described for the large population limit when each oscillator uses the approximately optimal control law. A critical value of the control cost parameter is identified: Above this value, the oscillators are incoherent; and below this value (when control is sufficiently cheap) the oscillators synchronize. These conclusions are illustrated with results from numerical experiments.

(J4) Huibing Yin, Prashant G. Mehta, Sean P. Meyn, Uday V. Shanbhag: On the Efficiency of Mean-Field OScillator Games. (To appear) Dynamic Games and Applications

Abstract: A key question in the design of engineered competitive systems has been that of the endciency loss of the associated equilibria. Yet, there is little known in this regard in the context of stochastic dynamic games, particularly in a large population regime. In this paper, we revisit a class of noncooperative games, arising from the synchronization of a large collection of heterogeneous oscillators. In [30], we derived a PDE model for analyzing the associated equilibria in large population regimes through a mean field approximation. Here, we examine the endciency of the associated mean-field equilibria with respect to a related welfare optimization problem. We construct constrained variational problems both for the noncooperative game and its centralized counterpart and derive the associated nonlinear eigenvalue problems. A relationship between the solutions of these eigenvalue problems is observed and allows for deriving an expression for endciency loss. By leveraging bifurcation analysis, a local bound on endciency loss is derived under an assumption that oscillators share the same frequency. Through numerical case studies, the analytical results are validated in the homogeneous frequency regime; analogous numerical results are provided for the heterogeneous frequency regime.

(C1) Huibing Yin, Uday V. Shanbhag, Prashant G. Mehta: Nash equilibrium problems with congestion costs and shared constraints. CDC 2009: 4649-4654

(C2) Huibing Yin, Prashant G. Mehta, Sean P. Meyn, Uday V. Shanbhag: Learning in mean-field oscillator games. CDC 2010: 3125-3132

(C3) Adam K. Tilton, Tao Yang, Huibing Yin, Prashant G. Mehta: Feedback particle filter-based multiple target tracking using bearing-only measurements. FUSION 2012: 2058-2064

Hao Jiang. Hao Jiang is in his final year of his doctoral program in the department of Industrial and Enterprise Systems Engineering. He was advised by PI Shanbhag and his committee will potentially include PI Meyn (not been formed yet). His research has examined a range of questions that combine learning and computation. Some of this research activity has been captured by the following papers.

(J1) Distributed computation of equilibria in misspecified convex stochastic Nash games Hao Jiang, Uday V. Shanbhag, Sean P. Meyn (submitted)

Abstract: The distributed computation of Nash equilibria is assuming growing relevance in the engineering community where such problems emerge in the context of distributed control. Accordingly, we consider static stochastic convex games complicated by a parametric misspecification, a natural concern when competing in any large-scale networked engineered system. In this paper, we present two sets of schemes in which firms learn the equilibrium strategy and correct the misspecification by leveraging noisecorrupted observations, which are subsequently supported by numerics. (1) Monotone stochastic Nash games: We present a scheme that combines a distributed stochastic approximation scheme with a learning update that leverages a least-squares estimator derived from a collection of observations. We proceed to show that the resulting sequence of estimators converges to the true equilibrium strategy and the true parameter in mean-square and in probability, respectively. (2) Monotone

Nash-Cournot games with unobservable aggregate output: We refine this scheme to a Cournot setting where the tuple of strategies is not observable and assume that payoff functions and strategy sets are public knowledge (a common knowledge assumption). When noise-corrupted prices are observable, distributed best response schemes are developed, allowing for simultaneously learning the equilibrium strategy and the parameter in an almost-sure sense. Furthermore, these statements may be extended to accommodate nonlinear price functions.

(C1) Hao Jiang, Uday V. Shanbhag: On the convergence of joint schemes for online computation and supervised learning. CDC 2012: 4462-4467

(C2) Hao Jiang, Uday V. Shanbhag, Sean P. Meyn: Learning equilibria in constrained Nash-Cournot games with misspecified demand functions. CDC-ECE 2011: 1018-1023

(C3) Hao Jiang and Uday V. Shanbhag: On the Solution of Stochastic Optimization Problems in Imperfect Information Regimes. To appear in the Winter Simulation Conference, 2013.

Quanyan Zhu. Quanyan Zhu has graduated from the department of Electrical and Computer Engineering in 2011 and was advised by PIs Basar. His research has captured a breadth of questions including stochastic and mean-field games. Here is a brief summary of his research papers.

(C1) Quanyan Zhu, Hamidou Tembine, Tamer Basar: Heterogeneous learning in zero-sum stochastic games with incomplete information. CDC 2010: 219-224

Abstract:Learning algorithms are essential for the applications of game theory in a networking environment. In dynamic and decentralized settings where the traffic, topology and channel states may vary over time and the communication between agents is impractical, it is important to formulate and study games of incomplete information and fully distributed learning algorithms which for each agent requires a minimal amount of information regarding the remaining agents. In this paper, we address this major challenge and introduce heterogeneous learning schemes in which each agent adopts a distinct learning pattern in the context of games with incomplete information. We use stochastic approximation techniques to show that the heterogeneous learning schemes can be studied in terms of their deterministic ordinary differential equation (ODE) counterparts. Depending on the learning rates of the players, these ODEs could be different from the standard replicator dynamics, (myopic) best response (BR) dynamics, logit dynamics, and fictitious play dynamics. We apply the results to a class of security games in which the attacker and the defender adopt different learning schemes due to differences in their rationality levels and the information they acquire.

(C2) Qi Zhang, Quanyan Zhu, Raouf Boutaba: Dynamic Resource Allocation for Spot Markets in Cloud Computing Environments. UCC 2011: 178-185

Abstract: The advent of cloud computing promises to provide computational resources to customers like public utilities such as water and electricity. To deal with dynamically fluctuating resource demands, market-driven resource allocation has been proposed and recently implemented by public Infrastructure-as-a-Service (IaaS) providers like Amazon EC2. In this environment, cloud resources are offered

in distinct types of virtual machines (VMs) and the cloud provider runs an auction-based market for each VM type with the goal of achieving maximum revenue over time. However, as demand for each type of VMs can fluctuate over time, it is necessary to adjust the capacity allocated to each VM type to match the demand in order to maximize total revenue while minimizing the energy cost. In this paper, we consider the case of a single cloud provider and address the question how to best match customer demand in terms of both supply and price in order to maximize the providers revenue and customer satisfactions while minimizing energy cost. In particular, we model this problem as a constrained discrete-time optimal control problem and use Model Predictive Control (MPC) to find its solution. Simulation studies using real cloud workloads indicate that under dynamic workload conditions, our proposed solution achieves higher net income than static allocation strategies and minimizes the average request waiting time.

(C3) Quanyan Zhu, Tamer Basar: A multi-resolution large population game framework for smart grid demand response management. NetGCoop 2011: 1-8

Abstract: Dynamic demand response (DR) management is becoming an integral part of power system and market operational practice. Motivated by the smart grid DR management problem, we propose a multi-resolution stochastic differential game-theoretic framework to model the players' intra-group and inter-group interactions in a large population regime. We study the game in both risk-neutral and risk-sensitive settings, and provide closed-form solutions for symmetric mean-field responses in the case of homogenous group population, and characterize the symmetric mean-field Nash equilibrium using the Hamilton-Jacobi-Bellman (HJB) equation together with the Fokker-Planck-Kolmogorov (FPK) equation. Finally, we apply the framework to the smart grid DR management problem and illustrate with a numerical example.

(C4) Quanyan Zhu, Tamer Basar: Robust and resilient control design for cyber-physical systems with an application to power systems. CDC-ECE 2011: 4066-4071

Abstract: The tradeoff between robustness and resilience is a pivotal design issue for modern industrial control systems. The trend of integrating information technologies into control system infrastructure has made resilience an important dimension of the critical infrastructure protection mission. It is desirable that systems support state awareness of threats and anomalies, and maintain acceptable levels of operation or service in the face of unanticipated or unprecedented incidents. In this paper, we propose a hybrid theoretical framework for robust and resilient control design in which the stochastic switching between structure states models unanticipated events and deterministic uncertainties in each structure represent the known range of disturbances. We propose a set of coupled optimality criteria for a holistic robust and resilient design for cyber-physical systems. We apply this method to a voltage regulator design problem for a synchronous machine with infinite bus and illustrate the solution methodology with numerical examples.

(C5) Quanyan Zhu, Zhu Han, Tamer Basar: A differential game approach to distributed demand side management in smart grid. ICC 2012: 3345-3350

Abstract: Smart grid is a visionary user-centric system that will elevate the conventional power grid system to one which functions more cooperatively, responsively, and economically. Dynamic demand side management is one of the key issues that enable the implementation of smart grid. In this paper, we use the framework of

dynamic games to model the distribution demand side management. The market price is characterized as the dynamic state using a sticky price model. A two-layer optimization framework is established. At the lower level, for each player (such as one household), different appliances are scheduled for energy consumption. At the upper level, the dynamic game is used to capture the interaction among different players in their demand responses through the market price. We analyze the N-person nonzero-sum stochastic differential game and characterize its feedback Nash equilibrium. A special case of homogeneous users is investigated in detail and we provide a closed-form solution for the optimal demand response. From the simulation results, we demonstrate the use of demand response strategy from the game-theoretic framework and study the behavior of market price and demand responses to different parameters.

(C6) Quanyan Zhu, Hamidou Tembine, Tamer Basar: Hybrid risk-sensitive mean-field stochastic differential games with application to molecular biology. CDC-ECE 2011: 4491-4497

Abstract: We consider a class of mean-field nonlinear stochastic differential games (resulting from stochastic differential games in a large population regime) with risk-sensitive cost functions and two types of uncertainties: continuous-time disturbances (of Brownian motion type) and event-driven random switching. Under some regularity conditions, we first study the best response of the players to the mean field, and then characterize the (strongly time-consistent Nash) equilibrium solution in terms of backward-forward macroscopic McKean-Vlasov (MV) equations, Fokker-Planck-Kolmogorov (FPK) equations, and Hamilton-Jacobi-Bellman (HJB) equations. We then specialize the solution to linear-quadratic mean-field stochastic differential games, and study in this framework the optimal transport of the GlpF transmembrane channel of Escherichia coli, where glycerol molecules (as players in the game) choose forces to achieve optimal transport through the membrane. Simulation studies show that GlpF improves the glycerol conduction more in a higher periplasmic glycerol concentration, which is consistent with observations made in the biophysics literature.

(C7) Tamer Basar, Quanyan Zhu: Prices of Anarchy, Information, and Cooperation in Differential Games. CoRR abs/1103.2579 (2011)

Anupama Kowli. Anupama Kowli has graduated from the department of Electrical and Computer Engineering in 2012. Her research placed an accent on the modeling and control in the context of power markets. She was advised by PI Meyn and a brief summary of her research follows.

(C1) Gui Wang, Matias Negrete-Pincetic, Anupama Kowli, Ehsan Shafieepoorfard, Sean P. Meyn, Uday V. Shanbhag: Real-time prices in an entropic grid. ISGT 2012: 1-8

(C2) Sean P. Meyn, Matias Negrete-Pincetic, Gui Wang, Anupama Kowli, Ehsan Shafieepoorfard: The value of volatile resources in electricity markets. CDC 2010: 1029-1036

Abstract:While renewable resources most certainly provide environmental benefits, and also help to meet aggressive renewable energy targets, their deployment has pronounced impacts on system operations. There is an acute need to understand these impacts in order to fully harness the benefits of renewable resource integration. In this paper we focus on the integration of wind energy resources in a multi-settlement electricity market structure. We study the dynamic competitive equilibrium for a

stochastic market model, and obtain closed form expressions for the supplier and consumer surpluses. Numerical results based on these formulae show that the value of wind generation to consumers, under the current operational practices, falls dramatically with volatility. In fact, we can establish thresholds for the coefficient of variation beyond which the value of wind is questionable. These findings emphasize the need to investigate operational schemes that address volatility.

(C3) Anupama Kowli, Sean P. Meyn, Supporting wind generation deployment with demand response, Power and Energy Society General Meeting, 2011 IEEE, 1-8

(C4) Gui Wang, Anupama Kowli, Matias Negrete-Pincetic, Ehsan Shafieepoorfard, Sean P. Meyn, A control theorist's perspective on dynamic competitive equilibria in electricity markets, Proc. 18th World Congress of the International Federation of Automatic Control

Abstract: If you talk to a control theorist today about the road to achieve an efficient, reliable, and affordable electricity supply you will most likely be told that real-time prices will be a part of its realization. Perhaps this is true. However, we argue that appropriate design using economic models that can capture the emerging physical realities is a key requirement for achieving a reliable, and smart electrical grid. To capture the potential pitfalls of real-time prices, we present an extension of our earlier work on dynamic markets to general network settings, allowing for more general constraints on generation and transmission. We conclude in wide generality that in the economic ideal of the competitive equilibrium, the standard results follow - the equilibrium is efficient, and average prices coincide with average marginal cost. However, these conclusions hold only on average. More importantly, we find that in the competitive equilibrium, (a) prices can be negative, (b) prices can go well above the choke up price - which is the maximum price consumers are willing to pay, and (c) the variance of the price decreases with increasing demand response. We illustrate these finding through numerical experiments.

Ehsan Shafieepoorfard. Ehsan Shafieepoorfard is in the department of Electrical and Computer Engineering and is advised by PI Meyn. Similar to A. Kowli, his interests also lie in control/optimization in the context of power systems and markets and his research activity is captured next.

(C1) Gui Wang, Matias Negrete-Pincetic, Anupama Kowli, Ehsan Shafieepoorfard, Sean P. Meyn, Uday V. Shanbhag: Real-time prices in an entropic grid. ISGT 2012: 1-8

(C2) Sean P. Meyn, Matias Negrete-Pincetic, Gui Wang, Anupama Kowli, Ehsan Shafieepoorfard: The value of volatile resources in electricity markets. CDC 2010: 1029-1036

Abstract: While renewable resources most certainly provide environmental benefits, and also help to meet aggressive renewable energy targets, their deployment has pronounced impacts on system operations. There is an acute need to understand these impacts in order to fully harness the benefits of renewable resource integration. In this paper we focus on the integration of wind energy resources in a multi-settlement electricity market structure. We study the dynamic competitive equilibrium for a stochastic market model, and obtain closed form expressions for the supplier and consumer surpluses. Numerical results based on these formulae show that the value of wind generation to consumers, under the current operational practices, falls dramatically with volatility. In fact, we can establish thresholds for the

coefficient of variation beyond which the value of wind is questionable. These findings emphasize the need to investigate operational schemes that address volatility.

(C3) Anupama Kowli, Sean P. Meyn, Supporting wind generation deployment with demand response, Power and Energy Society General Meeting, 2011 IEEE, 1-8

(C4) Gui Wang, Anupama Kowli, Matias Negrete-Pincetic, Ehsan Shafieepoorfard, Sean P. Meyn, A control theorist's perspective on dynamic competitive equilibria in electricity markets, Proc. 18th World Congress of the International Federation of Automatic Control

Abstract:If you talk to a control theorist today about the road to achieve an efficient, reliable, and affordable electricity supply you will most likely be told that real-time prices will be a part of its realization. Perhaps this is true. However, we argue that appropriate design using economic models that can capture the emerging physical realities is a key requirement for achieving a reliable, and smart electrical grid. To capture the potential pitfalls of real-time prices, we present an extension of our earlier work on dynamic markets to general network settings, allowing for more general constraints on generation and transmission. We conclude in wide generality that in the economic ideal of the competitive equilibrium, the standard results follow - the equilibrium is efficient, and average prices coincide with average marginal cost. However, these conclusions hold only on average. More importantly, we find that in the competitive equilibrium, (a) prices can be negative, (b) prices can go well above the choke up price - which is the maximum price consumers are willing to pay, and (c) the variance of the price decreases with increasing demand response. We illustrate these finding through numerical experiments.