Synchronization of Coupled Oscillators is a Game

Huibing Yin, Prashant G. Mehta, Sean P. Meyn and Uday V. Shanbhag

***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, economics, as well as traditional engineering applications. The focus of analysis is a variation of the large population LQG model of Huang et. al. 2007 [6], comprised here of a controlled nonlinear N-dimensional stochastic differential equation model, coupled only through a nonlinear cost function. The states are interpreted as the phase angle 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.**

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 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.

I. Introduction

The dynamics of a large population of coupled heterogeneous nonlinear systems is of interest in a number of applications (e.g., neuroscience, communication networks, power systems, markets). One important and possibly the simplest class of models for this problem is the Winfree model of N coupled oscillators,

$$\mathrm{d}\theta_i(t) = \omega_i\,\mathrm{d}t + \frac{\kappa}{N}\sum_{j=1}^{N}\psi^{\bullet}(\theta_j(t),\theta_i(t))\,\mathrm{d}t + \sigma\,\mathrm{d}\xi_i(t),$$

where $\theta_i(t)$ is the phase of the i^{th}-oscillator at time t, ω_i is its natural frequency, $\xi_i(t)$ is the standard Wiener process, $\psi^{\bullet}(\theta_j,\theta_i)$ models the influence of the j^{th}-oscillator from the population of N oscillators, and κ is the coupling parameter.

An example is the *Kuramoto model* [7] in which $\psi^{\bullet}(\theta_j,\theta_i) = \sin(\theta_j-\theta_i)$ for each i,j. It is assumed here as in the Kuramoto model that the frequency ω_i is drawn from a distribution $g(\omega)$ with support on $\Omega := [1-\gamma, 1+\gamma]$. The parameters γ and κ are used to model the heterogeneity and the strength of network coupling, respectively.

The dynamics can be visualized using a bifurcation diagram in the (κ,γ) plane, which in particular illustrates the emergence of a phase transition. The stability boundary $\kappa = \kappa_c(\gamma)$ shown on the left hand side of Fig. 1 provides an illustration of the phase transition: The oscillators behave incoherently for $\kappa < \kappa_c(\gamma)$, and synchronize for $\kappa > \kappa_c(\gamma)$. That is, the oscillators synchronize if the coupling is sufficiently large. In the former incoherent setting, the oscillators rotate close to their own natural frequency and hence the trajectory $\theta_i(t)$ is approximately independent of the population. In the synchronized setting each oscillator rotates with a common frequency.

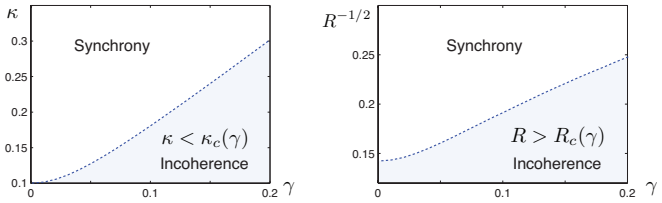

Fig. 1. Bifurcation diagrams. The Kuramoto model with $\sigma^2/2 = 0.05$ (left), and the coupled model considered in this paper with $\sigma^2/2 = 0.05$ (right).

The phase transition is important in a number of applications. For example, in thalamocortical circuits in the brain, transition to the synchronized state is associated with diseased brain states such as epilepsy [13], [15].

The objective of this paper is to model and interpret the phase transition from the perspective of noncooperative game theory. We define the game formally:

Consider a set of N oscillators. The model for the i^{th} oscillator is given by

$$\mathrm{d}\theta_i(t) = (\omega_i + u_i(t))\,\mathrm{d}t + \sigma\,\mathrm{d}\xi_i(t),$$

where $u_i(t)$ is the control input. Suppose the i^{th} oscillator minimizes its own performance objective:

$$\eta_i^{(\text{POP})}(u_i;u_{-i}) = \lim_{T\to\infty}\frac{1}{T}\int_0^T \mathsf{E}[c(\theta_i;\theta_{-i}) + \tfrac{1}{2}Ru_i^2]\,\mathrm{d}s, \qquad (1)$$

where $\theta_{-i} = (\theta_j)_{j\neq i}$, $c(\cdot)$ is a cost function, $u_{-i} = (u_j)_{j\neq i}$ and R models the control penalty. The form of the function c and the value of R are assumed to be common to the entire population. A *Nash equilibrium* in control policies is given by $\{u_i^*\}_{i=1}^N$ such that u_i^* minimizes $\eta_i^{(\text{POP})}(u_i;u_{-i}^*)$ for $i = 1,\dots,N$.

H. Yin and P. G. Mehta are with the Coordinated Science Laboratory and the Department of Mechanical Science and Engineering at the University of Illinois at Urbana-Champaign (UIUC) yin3@illinois.edu; mehtapg@illinois.edu

S. P. Meyn is with the Department of Electrical and Computer Engineering and the Coordinated Science Laboratory at UIUC meyn@illinois.edu

U. V. Shanbhag is with the Department of Industrial and Enterprise Systems Engineering at UIUC udaybag@illinois.edu

Financial support from the CSE fellowship at the University of Illinois, AFOSR grant FA9550-09-1-0190 and NSF grant CCF-0728863 are gratefully acknowledged. The authors thank Roland Malhame for many insightful discussions.

In general, establishing the existence and uniqueness of Nash equilibrium for large N is a challenging problem. In this paper, following an approach first employed in [6], we investigate a distributed control law wherein the i^{th} oscillator optimizes by using local information consisting of (i) its own state (θ_i) and (ii) the *mass-influence* of the population. The idea is that in the limit of large population size (as $N \to \infty$), the population affects the i^{th} oscillator in a nearly deterministic fashion. The distributed control law is obtained via optimization with respect to this deterministic (but not a priori known) mass influence.

Three types of analyses are presented in this paper. We first examine the infinite-oscillator limit, and subsequently investigate the implications for the finite-oscillator model:

1. The infinite oscillator limit. A limiting model is constructing consisting of two partial differential equations (PDEs):

(i) An Hamilton-Jacobi-Bellman (HJB) PDE (12) that describes the solution of minimizing (1) under the assumption of a known deterministic mass influence.

(ii) A Fokker-Planck-Kolmogorov (FPK) equation (15) that describes the evolution of the population density with optimal control input obtained from the solution of (i).

The two PDEs are coupled via the mass influence term (16). It arises as an averaged cost function, where the average is based on the solution of (ii). The averaged cost function is used in the HJB equation in (i). The solution of the HJB equation describes the distributed control law.

2. ε-Nash equilibrium for finite N. Following the methodology outlined in [6], we establish that the distributed control law is an ε-Nash equilibrium for the stochastic dynamic game with a finite number of oscillators ($N < \infty$). This implies that any unilateral deviation by an individual oscillator can at best improve the performance by a small ($\varepsilon = \mathcal{O}(\frac{1}{\sqrt{N}})$) amount when the population size N is large.

The final analysis is grounded in the large population limit:

3. Transition from incoherence to synchrony. A bifurcation diagram is obtained in the (R,γ) plane for the infinite limit model. The plot shown on the right in Fig. 1 depicts a phase transition: For $R > R_c$, the oscillators are incoherent, and for $R < R_c$ the oscillators synchronize. That is, the oscillators synchronize when the control is sufficiently cheap.

In addition to the work of [6], methods to construct approximate solutions to distributed control problems or dynamic games using related methods can be found in [1], [8], [16], [3], [11], [5], [9].

The remainder of this paper is organized as follows. A description of the SDE and PDE models is contained in Sec. II, and Sec. III contains analysis of the game for a finite number of oscillators. Bifurcation analysis appears in Sec. IV, which is illustrated with results from numerical experiments in Sec. V. Conclusions are contained in Sec. VI.

II. Oblivious Equilibria

We begin with a description of the coupled oscillator model, associated optimal control problems, and the proposed infinite-oscillator approximation.

A. Finite oscillator model

We consider a population of N oscillators competing in a noncooperative game as defined in the Introduction (see (1)). The dynamics of the i^{th} oscillator are described by the stochastic differential equation,

$$\mathrm{d}\theta_i = (\omega_i + u_i(t))\,\mathrm{d}t + \sigma\,\mathrm{d}\xi_i, \quad i = 1,\dots,N, \quad t \ge 0 \qquad (2)$$

where $\theta_i(t)$ is the phase of the i^{th} oscillator at time t, $u_i(t)$ is the control input, and $\{\xi_i\}$ are mutually independent standard Wiener processes. The standard deviation σ is independent of i.

For each i, the constant frequency ω_i is independent of time — It is assumed that at time $t = 0$, the N scalars $\{\omega_i\}$ are chosen independently according to a fixed distribution with density g, which is supported on an interval of the form $\Omega = [1-\gamma, 1+\gamma]$ for some $\gamma < 1$. In the numerical calculations that follow the density was taken to be uniform, namely $g(\omega) = (2\gamma)^{-1}$ for $|\omega - 1| \le \gamma$.

We seek control that is decentralized and of the following form: For each i and t the control $u_i(t)$ depends only on $\theta_i(t)$, and perhaps some aggregate information, such as the mean value of $\{\theta_j(t)\}_{j=1}^N$. This amounts to a dynamic game, whose exact solution is infeasible for large N.

Instead we seek an approximation of the form described in [16] or [6]. This approximation is based on an infinite-population limit similar to those introduced in this prior work and others (e.g., [14]). The approximation is based on a sequence of steps:

(i) We construct a density function p that is intended to approximate the probability density function for the individual oscillators. For any i and any $t > 0$, the density $p(\cdot, t, \omega_i)$ is intended to approximate the probability density of the random variable $\theta_i(t)$, evolving according to the stochastic differential equation (2).

(ii) We obtain an approximation for the cost function c. It is assumed that the cost function c appearing in (1) is separable, as shown below:

$$c(\theta_i; \theta_{-i}) = \frac{1}{N}\sum_{j\ne i} c^\bullet(\theta_i, \theta_j), \qquad (3)$$

with $c^\bullet$ a non-negative function on $\mathbb{R}^2$. If N is large, the sum in (3) is expected to be nearly deterministic when the frequencies $\{\omega_i\}$ are independently sampled according to the density g. The law of large numbers suggests the approximation of $c(\vartheta; \theta_{-i}(t))$ by $\bar{c}(\vartheta, t)$, where

$$\bar{c}(\vartheta, t) := \int_{\omega\in\Omega}\int_{\theta=0}^{2\pi} c^\bullet(\vartheta,\theta)\,p(\theta; t, \omega)\,g(\omega)\,\mathrm{d}\theta\,\mathrm{d}\omega. \qquad (4)$$

(iii) For the scalar model (2) with cost $\bar{c}(\vartheta, t)$ depending only on $\vartheta = \theta_i$, the game reduces to independent optimal control problems, and the Nash equilibrium reduces to an oblivious equilibrium [16].

Remark 1: The notion of an *oblivious equilibrium* was introduced by Weintraub et al. [16] as a means of approximating a Markov perfect equilibrium (MPE) of a dynamic

game with a large number of agents. The individual agents are oblivious to the state of the entire system and make their control decisions based only on local state variables, together with a consistently defined average. In the limit of large population size, this is justified using the LLN.

In the following subsection we develop the "oblivious" solution described in (iii). We then turn to the PDE approximation in (i) that defines the approximate cost in (ii).

B. Optimal control of a single oscillator

The cost function is time dependent, of the form,

$$\bar{c}(\theta_i,t)+\tfrac{1}{2}Ru_i^2,\quad \theta,u\in\mathbb{R}. \tag{5}$$

It is assumed that $\bar{c}$ is a periodic function of t, with period denoted τ. The dynamics remain of the affine form,

$$\mathrm{d}\theta_i=(\omega_i+u_i(t))\,\mathrm{d}t+\sigma\,\mathrm{d}\xi,\quad t\ge 0. \tag{6}$$

The average cost is defined as the limit supremum,

$$\eta_i(u_i;\bar{c})=\limsup_{T\to\infty}\frac{1}{T}\int_0^T \mathsf{E}[\bar{c}(\theta_i(s),s)+\tfrac{1}{2}Ru_i^2(s)]\,\mathrm{d}s \tag{7}$$

The goal is to minimize η_i over all adapted controls. We let η_i^* denote the minimal cost.

The pair $X(t)=(\theta_i(t),t)$ may be viewed as a controlled Markov process on the product space $[0,2\pi]\times[0,\tau]$, so that cost is only a function of this state and the control u. Moreover, for continuous state feedback $u(t)=\varphi(X(t))$ this Markov process is *hypoelliptic*, for which there is a rich ergodic theory. In particular, because of the compact state space, η_i^* exists and is independent of the initial state [10].

The associated average-cost optimality equations (or HJB equations) are given by,

$$\min_{u_i}\{\bar{c}(\theta,t)+\tfrac{1}{2}Ru_i^2+\mathscr{D}_{u_i}h_i(\theta,t)\}=\eta_i^* \tag{8}$$

where $h_i(\theta,t)$ is the relative value function and $\mathscr{D}_u$ denotes the controlled generator, defined for C^2 functions g via,

$$\mathscr{D}_u g=\partial_t g+(\omega_i+u)\partial_\theta g+\frac{\sigma^2}{2}\partial^2_{\theta\theta}g.$$

where ∂_t and ∂_θ denote the partial derivative with respect to t and θ, respectively, and $\partial^2_{\theta\theta}$ denotes the second derivative with respect to θ.

The relative value function $h_i(\theta,t)$ can be expressed as the integral,

$$h_i(\theta,t)=\int_t^\infty \mathsf{E}[\bar{c}(\theta_i(s),s)+\tfrac{1}{2}Ru_i^{*2}(s)-\eta_i^*\mid\theta_i(t)=\theta]\,\mathrm{d}s,$$

with $u_i^*(s)$ the optimal control. Because the cost is quadratic in u_i, and the dynamics linear in u_i, the optimal control in state feedback form is expressed $u_i^*(t)=\varphi_i(\theta,t)$, where

$$\varphi_i(\theta,t):=-\frac{1}{R}\partial_\theta h_i(\theta_i,t). \tag{9}$$

Substituting $u_i^*(t)$ into (8) gives the nonlinear PDE,

$$\partial_t h_i+\omega\partial_\theta h_i=\frac{1}{2R}(\partial_\theta h_i)^2-\bar{c}(\theta,t)+\eta_i^*-\frac{\sigma^2}{2}\partial^2_{\theta\theta}h_i. \tag{10}$$

C. PDE model

We now provide a complete description of the PDE model that is intended to approximate the stochastic model for large N. We begin by noting that for a single oscillator model, the evolution of the density p with state-feedback control $u=\varphi(\theta,t)$ is defined by the controlled PDE,

$$\partial_t p+\partial_\theta\left((\omega+\varphi(\theta,t))p\right)=\frac{\sigma^2}{2}\partial^2_{\theta\theta}p. \tag{11}$$

The notation in the large population limit is a minor variant of the $N=1$ solution: The relative value function is denoted by $h(\theta,t,\omega)$, which is a solution to the HJB equation,

$$\partial_t h+\omega\partial_\theta h=\frac{1}{2R}(\partial_\theta h)^2-\bar{c}(\theta,t)+\eta^*-\frac{\sigma^2}{2}\partial^2_{\theta\theta}h. \tag{12}$$

The average optimal cost is a function of ω:

$$\eta^*(\omega)=\lim_{T\to\infty}\frac{1}{T}\int_0^T\int_0^{2\pi}[\bar{c}(\theta,t)+\frac{1}{2R}(\partial_\theta h)^2]p(\theta;t,\omega)\,\mathrm{d}\theta\,\mathrm{d}t. \tag{13}$$

Using the associated optimal feedback control law,

$$\varphi(\theta,t,\omega):=-\frac{1}{R}\partial_\theta h(\theta,t,\omega), \tag{14}$$

the FPK equation that defines the evolution of density (denoted as $p(\theta;t,\omega)$) is given by,

$$\partial_t p+\omega\partial_\theta p=\frac{1}{R}\partial_\theta\left[p(\partial_\theta h)\right]+\frac{\sigma^2}{2}\partial^2_{\theta\theta}p. \tag{15}$$

The only difference thus far is notational: $h_i(\theta,t)$ is the value function for $N=1$ with a single frequency ω_i, and $h(\theta,t,\omega)$ is the value function for a continuum of oscillators, distinguished by their natural frequency ω. Such is the case because we have *assumed* $\bar{c}(\theta,t)$ on the right hand-side of (12) is a known deterministic periodic function that is furthermore consistent across the population.

The consistency is enforced here using the integral (4). The two PDEs are coupled through this integral that defines the relationship between the cost $\bar{c}$ and the density p:

$$\bar{c}(\vartheta,t)=\int_\Omega\int_0^{2\pi}c^\bullet(\vartheta,\theta)p(\theta;t,\omega)g(\omega)\,\mathrm{d}\theta\,\mathrm{d}\omega. \tag{16}$$

D. Incoherence

The system of equations (12) - (16) may have multiple solutions. Suppose that the cost $c^\bullet$ introduced in (3) is of the form $c^\bullet(\vartheta,\theta)=c^\bullet(\vartheta-\theta)$. In this case we single out the *incoherence* solution defined by

$$h(\theta,t,\omega)=h_0(\theta):=0\qquad p(\theta;t,\omega)=p_0(\theta):=\frac{1}{2\pi} \tag{17}$$

The control law (14) sets $u(t)\equiv 0$.

Consider the special case,

$$c^\bullet(\vartheta,\theta)=\frac{1}{2}\sin^2\left(\frac{\vartheta-\theta}{2}\right). \tag{18}$$

The cost $\bar{c}$ defined in (4) is constant in this solution,

$$\bar{c}(\vartheta,t)=\frac{1}{2}\frac{1}{2\pi}\int_\Omega\int_0^{2\pi}\sin^2\left(\frac{\vartheta-\theta}{2}\right)g(\omega)\,\mathrm{d}\theta\,\mathrm{d}\omega=\frac{1}{4},$$

which coincides with the average cost $\eta^*(\omega) = \eta_0 := \bar{c}$ for all $\omega \in \Omega$. This value is approximately consistent with the finite-N model. When each control is set to zero we obtain $\mathrm{d}\theta_i(t) = \omega_i \,\mathrm{d}t + \sigma \,\mathrm{d}\xi_i(t)$ for each i, which results in average cost independent of i,

$$\lim_{T\to\infty} \frac{1}{T}\int_0^T \mathsf{E}[c(\theta_i(t);\theta_{-i}(t))]\,\mathrm{d}t = \frac{N-1}{N}\eta_0$$

We return to this example in the bifurcation analysis of Sec. IV. There is a trade-off between reducing the cost associated with $\theta_i \neq \theta_j$, and reducing the cost of control. These competing costs suggest that a qualitative change in optimal control may arise when the parameter R varies from ∞ to 0. In Appendix VII-B we propose a numerical algorithm to compute a candidate optimal solution for the PDE model. We find that the algorithm converges to the incoherence solution whenever R is above a critical threshold (see Fig. 3).

III. ε-Nash equilibrium

In this section we assume that we have solved the optimal control problem for the PDE model described in Sec. II-C. We assume moreover that the resulting average cost $\bar{c}$ given in (16) is periodic in t. We show that in the stochastic model with $N < \infty$, the resulting control solution φ given in (14) defines an ε-Nash equilibrium, with $\varepsilon \to 0$ as $N \to \infty$.

A. Infinite population limit

To obtain the limiting model we impose additional assumptions on the stochastic model. Recall that the $\{\omega_i\}$ are chosen to be i.i.d. (and independent of $\{\xi_i\}$) from a distribution with density g. We assume moreover that the initial conditions are chosen randomly, with $\{(\theta_i(0),\omega_i)\}$ i.i.d., independent of $\{\xi_i\}$, with common marginal distribution $(\theta_i(0),\omega_i) \sim p(\theta;0,\omega)g(\omega)$.

Suppose that $N < \infty$, and that each oscillator is controlled using the control solution in (14):

$$u_i^o(t) = -\frac{1}{R}\partial_\theta h(\theta(t),t,\omega)\Big|_{\omega=\omega_i}. \tag{19}$$

By construction,

$$\eta_i(u_i^o;\bar{c}) \le \eta(u_i;\bar{c}) \tag{20}$$

for all adapted controls u_i. On account of the decentralized nature of the control law, the processes $\{\theta_i(t)\}$ are themselves independent. Hence analysis is brought down to the single oscillator of the form described in Sec. II-B.

For a given i with a fixed $\omega = \omega_i$, the evolution of the conditional distribution $p(\cdot;t,\omega_i)$ of the phase is described by the FPK equation (11). The evolution of the unconditional distribution of a particular θ_i is thus given by averaging this over ω with respect to the density g.

These arguments justify the averaging described in the previous section. Define the sequence of empirical distributions via,

$$\Gamma_N(A\times B) := \frac{1}{N}\sum_{i=1}^N \mathbb{1}\{\omega_i(t)\in A,\ \theta_i(t)\in B\}, \qquad \begin{matrix} A\subset\Omega, \\ B\subset[0,2\pi]. \end{matrix}$$

With this decentralized control law we obtain from the LLN,

$$\lim_{N\to\infty}\Gamma_N(A\times B) = \int_{\omega\in A}\int_{\vartheta\in B} g(\omega)p(\vartheta,t,\omega)\,\mathrm{d}\vartheta\,\mathrm{d}\omega.$$

B. Comparison of Nash and oblivious control laws

We denote by $\theta_i^o(t)$ the solution to the SDE (2), obtained using the oblivious control $u_i = u_i^o(t)$ defined in (19), and $\theta_{-i}^o = (\theta_1^o,\dots,\theta_{i-1}^o,\theta_{i+1}^o,\dots,\theta_N^o)$. The conditional mean of (3) is used as an approximate cost function for a single oscillator,

$$\bar{c}_i^{(N)}(\vartheta,t) := \mathsf{E}[c(\theta_i(t);\theta_{-i}^o(t)) \mid \theta_i(t)=\vartheta)]. \tag{21}$$

In Lemma 3.1 this approximate cost function is used to make precise the nature of approximation in going from cost as the summation (3) in the finite oscillator case, to the integral (16) in the PDE limit. The proof follows from the CLT applied to the samples $\{(\theta_i(0)\}$.

Lemma 3.1: Consider $\bar{c}(\vartheta,t)$ in (16) and $\bar{c}_i^{(N)}(\vartheta,t)$ in (21). In the limit of large N, for each $i = 1,\dots,N$,

$$\max_{\vartheta\in[0,2\pi]}\left[\limsup_{T\to\infty}\frac{1}{T}\int_0^T \big|\bar{c}(\vartheta,s)-\bar{c}_i^{(N)}(\vartheta,s)\big|\,\mathrm{d}s\right] = O(\frac{1}{\sqrt{N}})$$

■

The proof of the next result is straightforward (see also [6]).

Lemma 3.2: Consider a single oscillator control problem (7) with cost function $\bar{c}(\theta_i,t)+\frac{1}{2}Ru_i^2$ and the associated optimal control u_i^*. Let $\bar{c}^\varepsilon(\theta,t)$ denote some perturbation of $\bar{c}(\theta,t)$ for which there is an $\varepsilon > 0$ such that,

$$\max_{\vartheta\in[0,2\pi]}\left[\limsup_{T\to\infty}\frac{1}{T}\int_0^T\Big[\big|\bar{c}(\vartheta,s)-\bar{c}^\varepsilon(\vartheta,s)\big|\Big]\,\mathrm{d}s\right] \le \varepsilon.$$

Then $\eta_i(u_i^*;\bar{c}^\varepsilon) \le \eta(u_i;\bar{c}^\varepsilon)+2\varepsilon$, for any adapted control u_i. ■

We can now establish the main result of this section:

Theorem 3.3: The oblivious control $\{u_i^o\}$ is an ε-Nash equilibrium for (1): For any adapted control u_i,

$$\eta_i^{(\mathrm{POP})}(u_i^o;u_{-i}^o) \le \eta_i^{(\mathrm{POP})}(u_i;u_{-i}^o) + O(\frac{1}{\sqrt{N}}).$$

Proof: From (20), $\eta_i(u_i^o;\bar{c}) \le \eta(u_i;\bar{c})$. The ε-Nash property follows from Lemma 3.2 because $\bar{c}_i^{(N)}(\vartheta,s)$ approximates $\bar{c}(\vartheta,s)$ in the limit of large N (see Lemma 3.1). ■

IV. Bifurcation analysis of PDEs

In the remainder of the paper we present a finer analysis of the coupled equations (12) – (16) for a particular choice of $c^\bullet$. Our main goal is to establish a transition from incoherence to synchrony as the control penalty parameter R is decreased beyond a critical value. The analytical conclusions are illustrated with results from numerical experiments.

Throughout the remainder of the paper we restrict to the cost function $c^\bullet$ defined in (18).

Solutions to the equations (12) - (16) are investigated here using the method of bifurcation theory; the parameter R is used as the bifurcation parameter. The following assumptions are imposed on the model, and on any solution (p,h) considered in our analysis:

(A1) The density g is uniform on $\Omega := [1-\gamma, 1+\gamma]$.
(A2) The functions p, h are periodic in θ:

$$h(\theta,t,\omega) = h(\theta+2\pi,t,\omega),$$
$$p(\theta;t,\omega) = p(\theta+2\pi,t,\omega), \qquad \theta\in[0,2\pi], t\ge 0, \omega\in\Omega.$$

We single out one solution obtained in Sec. II-D: The incoherence solution. We denote this solution by $z_0(\theta) := (h_0(\theta), p_0(\theta))^{\mathrm{T}}$. The existence of a bifurcating solution branch is investigated via analysis of a linearization about z_0. The spectral analysis of the linearization is used to obtain the bifurcation point as a critical value of $R = R_c(\gamma)$.

To verify the conclusions of bifurcation analysis, the solution of the PDE is obtained numerically by using an algorithm presented in Appendix VII-B. Numerical results described in Sec. V show that the incoherent solution is a limiting fixed-point of the algorithm when $R > R_c$. Below the critical value of R, the incoherent solution is no longer 'stable.' The numerical algorithm yields a periodic traveling wave solution that is interpreted as the synchrony solution.

For both types of solutions, the cost function $\bar{c}$ is periodic in time and θ: For some value of $\tau > 0$,

$$\bar{c}(\theta+2\pi, t+\tau) = \bar{c}(\theta,t), \qquad \theta\in\mathbb{R}, t\ge 0\,.$$

We find that $\tau = 2\pi$ in the numerical experiments considered below, which under (A1) coincides with the mean value of ω over Ω.

A. Linear analysis

The linearization of the equations (12) - (16) is taken about the equilibrium incoherence solution $z_0 = (h_0, p_0)$. A perturbation of this solution is denoted $z_0 + \tilde{z} = (h_0, p_0) + (\tilde{h}, \tilde{p})$. Since $p = p_0 + \tilde{p}$ is a probability density, the perturbation satisfies the normalization condition $\int_0^{2\pi} \tilde{p}(\theta;t,\omega)\,\mathrm{d}\theta = 0$ for any t, ω. Since the relative value function is only defined to a constant, we also impose a similar normalization condition for h: $\int_0^{2\pi} \tilde{h}(\theta,t,\omega)\,\mathrm{d}\theta = 0$ for any t, ω.

When $\tilde{z}$ is small, its evolution is approximated by the linear equation,

$$\frac{\partial}{\partial t}\tilde{z}(\theta,t,\omega) = \mathscr{L}_R\tilde{z}(\theta,t,\omega),$$

where

$$\mathscr{L}_R\tilde{z}(\theta,t,\omega) := \begin{pmatrix} -\omega\partial_\theta\tilde{h} - \bar{c} - \frac{\sigma^2}{2}\partial^2_{\theta\theta}\tilde{h} \\ -\omega\partial_\theta\tilde{p} + \frac{1}{2\pi R}\partial^2_{\theta\theta}\tilde{h} + \frac{\sigma^2}{2}\partial^2_{\theta\theta}\tilde{p} \end{pmatrix} \tag{22}$$

and $\bar{c}(\theta,t) =$

$$\frac{1}{2}\int_\Omega\int_0^{2\pi} \sin^2\left(\frac{\theta-\vartheta}{2}\right)\tilde{p}(\vartheta;t,\omega)g(\omega)\,\mathrm{d}\vartheta\,\mathrm{d}\omega \tag{23}$$

On taking the Laplace transform of (22) we can obtain the representation $Z = [I\lambda - \mathscr{L}_R]^{-1}\tilde{z}(0)$, whenever the inverse exists. We say that $\lambda\in\mathbb{C}$ is in the spectrum of $\mathscr{L}_R$ if the inverse $[I\lambda - \mathscr{L}_R]^{-1}$ does not exist as a bounded linear operator on $L^2([0,2\pi]\times\Omega)$.

The associated eigenvector problem is given by,

$$\lambda Z = \mathscr{L}_R Z,$$

where $Z(\theta,\lambda,\omega) = (H(\theta,\lambda,\omega), P(\theta,\lambda,\omega))$. Key to analysis is the Fourier series expansion with respect to θ,

$$H = \sum_{k=-\infty}^{+\infty} H_k(\omega)e^{ik\theta}, \quad P = \sum_{k=-\infty}^{+\infty} P_k(\omega)e^{ik\theta}, \tag{24}$$

where dependence on λ is suppressed for notational ease. The two normalization conditions give $P_0(\omega) = H_0(\omega) = 0$.

Let $\bar{C}$ denote the Laplace transform of $\bar{c}$ (see (23)). Using the Fourier series expansion

$$\bar{C}(\theta) = -\frac{\pi}{2}\sum_{k=\{1,-1\}} e^{ik\theta}\int_\Omega P_k(\omega)g(\omega)\,\mathrm{d}\omega. \tag{25}$$

Using (24) - (25) yields a diagonal decomposition of the linear operator

$$\mathscr{L}_R = \bigoplus_k \mathscr{L}_R^{(k)},$$

where each $\mathscr{L}_R^{(k)}$ acts on the pair $(H_k, P_k)^{\mathrm{T}}$. The individual operators have the explicit form

$$\mathscr{L}_R^{(1)} := \begin{pmatrix} \frac{\sigma^2}{2} - \omega i & \frac{\pi}{4}\int_\Omega g(\omega)\,\mathrm{d}\omega \\ -\frac{1}{2\pi R} & -\frac{\sigma^2}{2} - \omega i \end{pmatrix}$$

$$\mathscr{L}_R^{(k)}Z_k := \begin{pmatrix} \frac{\sigma^2}{2}k^2 - k\omega i & 0 \\ -\frac{k^2}{2\pi R} & -\frac{\sigma^2}{2}k^2 - k\omega i \end{pmatrix}, \quad k\ge 2,$$

and $\mathscr{L}_R^{(-k)} = \overline{\mathscr{L}^{(k)}}_R$.

The spectrum of $\mathscr{L}_R$ is given by the union of spectrum of $\mathscr{L}_R^{(k)}$, $k = \pm1, \pm2, \ldots$. In general, the spectrum include both continuous and discrete parts. The continuous spectrum of $\mathscr{L}_R^{(k)}$ correspond to points λ in the spectrum such that,

(i) The operator $\lambda I - \mathscr{L}_R^{(k)}$ is injective,
(ii) The range of operator $\lambda I - \mathscr{L}_R^{(k)}$ is dense in $L^2(\Omega)$.

For background see [12].

In Appendix VII-A we establish the following characterization of the spectrum:

Theorem 4.1: For the linear operator $\mathscr{L}_R : L^2([0,2\pi]\times\Omega) \to L^2([0,2\pi]\times\Omega)$,

(i) The continuous spectrum equals the union of sets $\{S^{(k)}\}_{k=-\infty}^{\infty}$ where

$$S^{(k)} := \left\{\lambda\in\mathbb{C}\,\middle|\,\lambda = \pm\frac{\sigma^2}{2}k^2 - k\omega i \text{ for all } \omega\in\Omega\right\}.$$

(ii) The discrete spectrum coincides with the discrete spectrum of $\mathscr{L}_R^{(k)}$ for $k = \pm1$.

■

The points in $S^{(k)}$ are in one-one correspondence with the frequencies in the distribution $g(\omega)$. That is, for each $\omega_0\in\Omega$, the point $\pm\frac{\sigma^2}{2}k^2 - k\omega_0 i \in S^{(k)}$ lies in the continuous spectrum. On the complex plane, $S^{(k)}$ comprises of two line segments, one in the left half-plane and the other in the right half-plane. The main thing to note is that the continuous spectrum does not change with the value of R and is moreover bounded away from the imaginary axis for $k = \pm1, \pm2, \cdots$. So, the focus of the analysis and the

numerical study that follows is on the discrete spectrum for $k=\pm 1$.

For $k=1$, let $(H_1,P_1)^{\mathrm{T}}$ denote the eigenvector corresponding to an eigenvalue λ. We assume $\lambda \notin S^{(1)}$, the set contained in continuous spectrum. We have:

$$\lambda H_1(\omega) = \left(\frac{\sigma^2}{2} - \omega i\right) H_1(\omega) + \frac{\pi}{4}\int_\Omega P_1(\omega) g(\omega)\,\mathrm{d}\omega, \quad (26)$$

$$\lambda P_1(\omega) = -\frac{1}{2\pi R} H_1(\omega) - \left(\frac{\sigma^2}{2} + \omega i\right) P_1(\omega). \quad (27)$$

We formally obtain from (27),

$$P_1(\omega) = -\frac{1}{2\pi R}\frac{H_1(\omega)}{\lambda + \frac{\sigma^2}{2} + \omega i}$$

and on substituting this into (26),

$$H_1(\omega) = -\frac{1}{8R(\lambda - \frac{\sigma^2}{2} + \omega i)}\int_\Omega \frac{H_1(\omega)}{\lambda + \frac{\sigma^2}{2} + \omega i} g(\omega)\,\mathrm{d}\omega. \quad (28)$$

The solution $H_1, P_1 \in L^2(\Omega)$ because $\lambda \notin S^{(1)}$, i.e., $\lambda \pm \sigma^2/2 + \omega i \neq 0$ for all values of $\omega \in \Omega$. Denote $b := \int_\Omega \frac{H_1(\omega)}{\lambda + \frac{\sigma^2}{2} + \omega i} g(\omega)\,\mathrm{d}\omega$ which is a constant independent of ω. This gives $H_1(\omega) = -b\big(8R(\lambda - \frac{1}{2}\sigma^2 + \omega i)\big)^{-1}$. Substituting this into (28) yields the characteristic equation for λ:

$$\frac{1}{8R}\int_\Omega \frac{g(\omega)}{(\lambda - \frac{\sigma^2}{2} + \omega i)(\lambda + \frac{\sigma^2}{2} + \omega i)}\,\mathrm{d}\omega + 1 = 0. \quad (29)$$

For $k=-1$, the eigenvalue is complex conjugate $\bar{\lambda}$.

V. Numerical results

We present here results from computational experiments based on the coupled equations (12) - (16).

We fix $\sigma^2/2 = 0.05$ throughout, and the cost parameter R is treated as a variable in the bifurcation analyses that follow.

A. Eigenvalue as a function of R

The characteristic equation (29) was solved numerically to obtain a path of eigenvalues as a function of R.

Fig. 2 (a) depicts a locus of eigenvalues obtained for a family of models parameterized by R; γ is fixed at 0.05. For $R \sim \infty$ there are a pair of complex eigenvalues at $\pm\frac{\sigma^2}{2} - i$ (for $k=1$). As the parameter R decreases, these eigenvalues move continuously towards the imaginary axis. The critical value R_c is defined as the value of R at which these two eigenvalue paths collide on the imaginary axis, resulting in an eigenvalue pair of multiplicity 2. The eigenvalues split as R is decreased further, and remain on the imaginary axis for $R < R_c$. The real and the imaginary part of the two eigenvalue paths originating at $\pm\frac{\sigma^2}{2} - i$ are depicted in Fig. 2 (b). This eigenvalues also have their complex conjugate counterparts (for $k=-1$) that are not depicted for the sake of clarity.

In (a) and (b) the value of γ is fixed at 0.05. The critical value R_c is a function of the parameter γ. Fig. 2 (c) depicts a plot of $R_c(\gamma)$ as a function of γ. In the (R,γ) plane, this defines a boundary separating potentially two kinds of behavior – incoherence and synchrony.

The linear analysis suggests appearance of the synchrony solution via the Hamiltonian Hopf bifurcation at $R = R_c$ [4]. A rigorous justification for existence will require nonlinear analysis that is not presented here. Instead, we present numerical results on these two types of solution obtained for values of R greater than and less than the critical value R_c.

In the remainder of this section we restrict to $\gamma = 0.05$, which gives $R_c(\gamma) = 39.1$. The computations that follow are based on the waveform relaxation algorithm described in Appendix VII-B.

In numerical experiments the uniform distribution $g(\omega) = (2\gamma)^{-1}$ on the interval $\Omega = [1-\gamma, 1+\gamma]$ is approximated by a uniform distribution on three discrete frequencies $\{1-\gamma, 1, 1+\gamma\}$. The value of $\gamma = 0.05$ is sufficiently small so that the numerical results are very similar to those obtained using a finer discretization of Ω. The PDEs are discretized along the θ coordinate using the method of Fourier collocation [2] with 64 collocation points in the interval $[0, 2\pi]$.

B. Average cost bifurcation diagram

Fig. 3 depicts a numerically obtained bifurcation diagram for the average cost $\eta(\omega)$ as a function of the bifurcation parameter R.

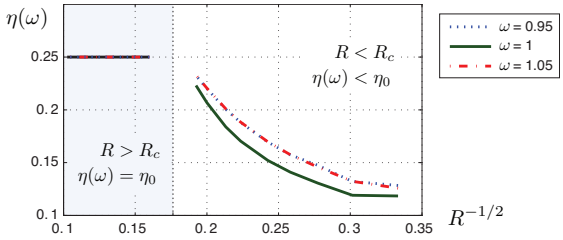

Fig. 3. Bifurcation diagram: the average cost as a function of $1/\sqrt{R}$.

For $R > R_c = 39.1$, the average cost was found to be $\eta(\omega) = \eta_0 = \frac{1}{4}$, which is consistent with the incoherence solution of Sec. II-D. For $R < R_c$ the average cost is reduced, and for such R the value of $\eta(\omega) < \eta_0$ depends upon the frequency ω. Its minimal value is attained uniquely when $\omega = 1$, which is the mean frequency under g.

C. Value functions, control, and density evolution

The relative value function $h(\theta,t,\omega)$ and probability density $p(\theta;t,\omega)$ were computed for a range of values of R.

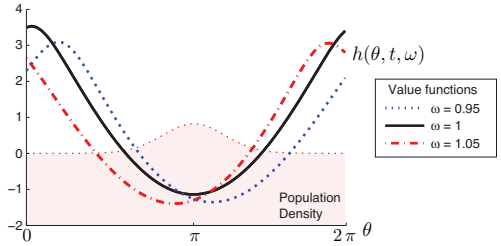

Fig. 4. Relative value function for $R = 10$, and the population density p for a particular value of t.

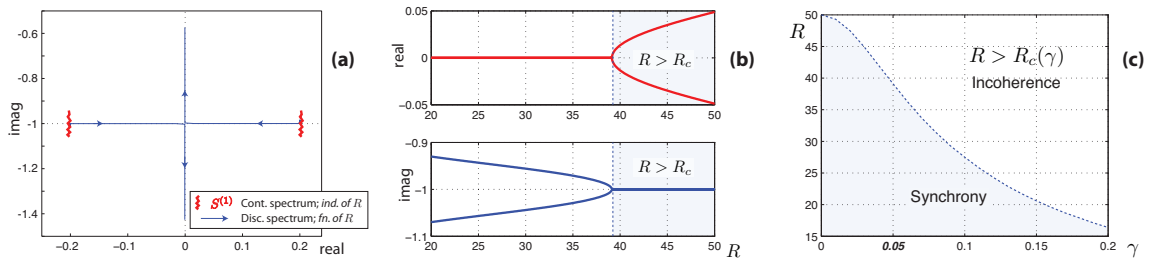

Fig. 2. Spectrum as a function of R. (a) The continuous spectrum for $k = 1$ and 2, along with the two eigenvalue paths as R decreases. (b) The real and imaginary parts of the two eigenvalue paths as R decreases. (c) $R_c(\gamma)$ as a function of γ.

The incoherence solution $h_0 \equiv 0$ was obtained for $R \geq 60$; the algorithm was very slow to converge as R was reduced to values near R_c.

Fig. 4 depicts the relative value function as a function of θ obtained for $R = 10 < R_c$, and for a particular value of t. Experiments revealed that the relative value function and the solution to the FPK equation arise as a traveling wave solution. In particular, the solution p has the form,

$$p(\theta;t,\omega) = p(\theta - at;0,1), \quad h(\theta,t,\omega) = h(\theta - at,0,\omega)$$

Moreover, the wave speed was equal to $a = 1$, independent of ω, which coincides with the mean frequency with respect to the density g.

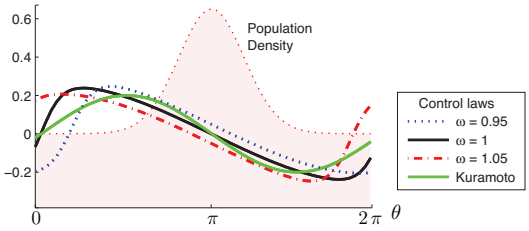

Fig. 5. Comparison of the control obtained from solving (12) - (16) and the Kuramoto model.

Recall the control law (19) is $u^*(t) = -\frac{1}{R}\partial_\theta h^*(\theta,t,\omega)$, which depends upon the frequency ω. The control laws obtained for a fixed t and several values of ω are depicted in Fig. 5 in relation to the population density. Note that the control is zero when $\omega = 1$, and θ lies at its mean value (equal to π in this figure, for the particular value of t chosen).

The control law that gives rise to the Kuramoto oscillator is defined by, $u_i^{(\mathrm{Kur})}(\theta_i,t) = \frac{\kappa}{N}\sum_{j=1}^N \sin(\theta_j(t) - \theta_i)$. Given the previous numerical results using $R = 10$, it is reasonable to conjecture that as N tends to infinity this can be approximated by $u_i^{(\mathrm{Kur})}(\theta_i,t) = \kappa_0 \sin(\vartheta_0 + t - \theta_i)$, for a phase variable ϑ_0 and a gain κ_0 that is proportional to κ. This is because in the synchrony state, the individual oscillators rotate with a common frequency 1. That is, $\theta_j(t) \approx t + \vartheta_{0,j}$ for some $\vartheta_{0,j} \in [0, 2\pi]$. Without noise, such a solution is in fact exact for the Kuramoto oscillators. Fig. 5 shows that the optimal control law is in fact "close" to $u_i^{(\mathrm{Kur})}$ when κ_0, t, and ϑ_0 are chosen appropriately.

VI. Conclusions

This paper aggregates concepts and techniques from nonlinear dynamical systems, game theory, and statistical mechanics to provide new tools for understanding complex interconnected systems, and new bridges with prior research. The key messages are,

(i) Distributed control laws are tractable for a class of large population dynamic games with separable cost structures. This conclusion is based on an approximation of the complex stochastic system using a deterministic PDE model, similar to the mean-field approximation that is central to the study of interacting particle systems.

(ii) The rich theory surrounding the classical Kuramoto model can be extended to the dynamic game setting introduced here to explain phase transitions in these systems. In particular, methods from bifurcation theory can be adopted to analyze multiple equilibria and their stability properties.

The proposed methods are expected to be relevant to applications involving large population of controlled heterogeneous nonlinear systems.

VII. Appendix

A. Proof of theorem 4.1

We provide a proof only for $|k| \geq 2$. The proof for $k = \pm 1$ is conceptually similar but some of the calculations are more complex.

We consider the equation

$$(\lambda I - \mathscr{L}_R^{(k)}) \begin{pmatrix} H(\omega) \\ P(\omega) \end{pmatrix} = \begin{pmatrix} \upsilon(\omega) \\ \zeta(\omega) \end{pmatrix}$$

where $\upsilon(\omega)$, $\zeta(\omega) \in L^2(\Omega)$. Explicitly, this gives

$$(\lambda - \frac{\sigma^2}{2}k^2 + k\omega i)H(\omega) = \upsilon(\omega),$$
$$(\lambda + \frac{\sigma^2}{2}k^2 + k\omega i)P(\omega) = -\frac{1}{2\pi R}H(\omega) + \zeta(\omega).$$

Formally the inverse, if it exists, is given by

$$H(\omega) = \frac{1}{\lambda - \frac{\sigma^2}{2}k^2 + k\omega i}\upsilon(\omega), \tag{30}$$

$$P(\omega) = \frac{1}{\lambda + \frac{\sigma^2}{2}k^2 + k\omega i}\left[-\frac{1}{2\pi R}H(\omega) + \zeta(\omega)\right]. \tag{31}$$

The proof that $\mathscr{L}_R^{(k)}$ is 1-1 for all $\lambda \in \mathbb{C}$ is now straightforward. If $\upsilon(\omega) = 0$ then $H(\omega) = 0$ in $L^2(\Omega)$ using (30) and if additionally $\zeta(\omega) = 0$ then $P(\omega) = 0$ in $L^2(\Omega)$ using (31).

Using the formula for the inverse, the inverse operator is bounded if and only if $\lambda \notin S^{(k)}$. If $\lambda = \lambda_0 = \frac{\sigma^2}{2}k^2 - k\omega_0 i \in S^{(k)}$ for some $\omega_0 \in \Omega$, then $\lambda_0 - \frac{\sigma^2}{2}k^2 + k\omega i = 0$ for $\omega = \omega_0$ and the inverse $\left(\lambda_0 - \frac{\sigma^2}{2}k^2 + k\omega i\right)^{-1}$ in (30) is not bounded. The converse also follows similarly.

Finally, the range of $\lambda_0 I - \mathscr{L}^{(k)}$ is dense in $L^2(\Omega)$ because consider for example, the space of C^1 functions with $\upsilon(\omega_0) = \upsilon'(\omega_0) = 0$. ■

B. Algorithm to solve coupled PDEs

In this section, we present a numerical algorithm to obtain a solution of coupled PDEs (12,15), subject to the static equation (16). For numerical implementation, we consider the finite-horizon optimal control problem on $[0,T]$. The value function is denoted,

$$h^{(T)}(\theta,t;\bar{c}) = \min_{u(\cdot)} \mathsf{E}\int_t^T [\bar{c}(\theta,s) + \tfrac{1}{2}Ru^2(s)]\,\mathrm{d}s. \qquad (32)$$

Comparing (1) or (7) with (32), we are interested in the limit as $T \to \infty$. It follows from the discussion in Sec. II-B that the cost given in (13) is the limit of the normalized value functions,

$$\eta^*(\omega) = \limsup_{T\to\infty} \frac{1}{T} h^{(T)}(\theta,t;\bar{c}).$$

The HJB equation for (32) is given by

$$\partial_t h^{(T)} + \omega\partial_\theta h^{(T)} = \frac{(\partial_\theta h^{(T)})^2}{2R} - \bar{c}(\theta,t) - \frac{\sigma^2}{2}\partial^2_{\theta\theta} h^{(T)}, \quad (33)$$

with boundary condition $h^{(T)}(\theta,T) = 0$. Comparing (33) with (12), we see that the only difference between the two HJB equations is the term η on the RHS of (12).

One issue with numerical solution of (33) is that it must be solved in backward time, starting at $t = T$. To obtain a causal implementation, we introduce a new coordinate $s = T - t$, and denote $V(\theta,s) = h^{(T)}(\theta,t)$ for $t \in [0,T]$.

We thereby obtain the coupled system of equations,

$$\partial_s V = \omega\partial_\theta V - \frac{(\partial_\theta V)^2}{2R} + \bar{c}(\theta,T-s) + \frac{\sigma^2\partial^2_{\theta\theta}V}{2} \quad (34)$$

$$\partial_t p = -\omega\partial_\theta p + \frac{1}{R}\partial_\theta\left[p(\partial_\theta V)\right] + \frac{\sigma^2}{2}\partial^2_{\theta\theta}p, \quad (35)$$

$$\bar{c}(\theta,t) = \int_\Omega\int_0^{2\pi} c^\bullet(\theta,\vartheta)p(\vartheta;t,\omega)g(\omega)\,\mathrm{d}\vartheta\,\mathrm{d}\omega. \quad (36)$$

with initial condition $V(\theta,0) = h^{(T)}(\theta,T) = 0$, and with $p(\theta;0)$ given.

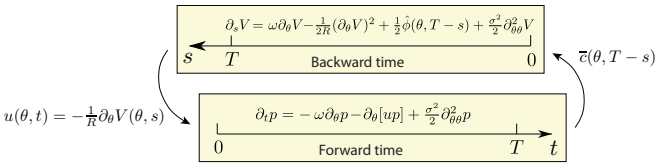

Fig. 6. Information flow of the numerical algorithm.

The system is solved by using a method based on waveform relaxation — the algorithm is illustrated in Fig. 6. It is initialized with some initial guess for the density $p_0(\theta,t)$ over the time horizon $[0,T]$. The following sequence of computations are performed in the k^{th} iteration,

(i) Use $p_k(\theta,t)$ to evaluate the $\bar{c}(\theta,t)$ using (36). Denote it as $\bar{c}_k(\theta,t)$.

(ii) With $\bar{c}_k(\theta,t)$, simulate the causal representation (34) to obtain the solution V_k and hence $h_k^{(T)}$.

(iii) The optimal control law is $u_k(\theta,t) = -\frac{1}{R}\partial_\theta h_k^{(T)}(\theta,t)$.

(iv) Use the control law $u_k(\theta,t)$ to obtain the new density $p_{k+1}(\theta,t)$ from the solution of the FPK equation (35).

This is then repeated with k replaced by $k+1$.

References

[1] S. Adlakha, R. Johari, G. Y. Weintraub, and A. J. Goldsmith. Oblivious equilibrium for large-scale stochastic games with unbounded costs. In *Proc. of the 47th IEEE Conf. On Decision and Control*, pages 5531–5538, Cancun, Mexico, dec 2008.

[2] C. Canuto, M. Y. Hussaini, A. Quarteroni, and T. A. Zang. *Spectral Methods in Fluid Dynamics*. Springer-Verlag, 1988.

[3] R. Cogill, M. Rotkowitz, B. Van Roy, and S. Lall. An approximate dynamic programming approach to decentralized control of stochastic systems. In *Proceedings of the 2004 Allerton Conference on Communication, Control, and Computing*, pages 1040–1049, 2004.

[4] M. Dellnitz, J.E. Marsden, I. Melbourne, and J. Scheurle. Generic bifurcations of pendula. *Int. Series Num. Math.*, 104:111–122, 1992.

[5] V. F. Farias, D. Saure, and G. Y. Weintraub. The linear programming approach to solving large scale dynamic stochastic games (working paper). http://www.stanford.edu/ bvr/publ-all.html, 2008.

[6] M. Huang, P. E. Caines, and R. P. Malhame. Large-population cost-coupled LQG problems with nonuniform agents: Individual-mass behavior and decentralized ε-Nash equilibria. *IEEE Trans. Automat. Control*, 52(9):1560–1571, 2007.

[7] Y. Kuramoto. *International Symposium on Mathematical Problems in Theoretical Physics*, volume 39 of *Lecture Notes in Physics*. Springer-Verlag, 1975.

[8] T. Li and J.-F. Zhang. Asymptotically optimal decentralized control for large population stochastic multiagent systems. *IEEE Trans. Automat. Control*, 53(7):1643–1660, 2008.

[9] P. Mehta and S. Meyn. Q-learning and Pontryagin's Minimum Principle. Accepted for inclusion in the 48th IEEE Conference on Decision and Control, December 16-18 2009.

[10] S. P. Meyn. The policy iteration algorithm for average reward Markov decision processes with general state space. *IEEE Trans. Automat. Control*, 42(12):1663–1680, 1997.

[11] S. P. Meyn and G. Mathew. Shannon meets Bellman: Feature based Markovian models for detection and optimization. In *Proc. 47th IEEE CDC*, pages 5558–5564, 2008.

[12] A. W. Naylor and G. R. Sell. *Linear Operator Theory in Engineering and Science*, volume 40 of *Applied Mathematical Sciences*. Springer-Verlag, 1982.

[13] M. Steriade, D. A. McCormick, and T. J. Sejnowski. Thalamocortical Oscillations in the Sleeping and Aroused Brain. *Science*, 262:679–685, October 1993.

[14] S. H. Strogatz and R. E. Mirollo. Stability of incoherence in a population of coupled oscillators. *Journal of Statistical Physics*, 63:613–635, May 1991.

[15] M. von Krosigk, T. Bal, and D. A. McCormick. Cellular mechanisms of a synchronized oscillation in the thalamus. *Science*, 261:361–364, 1993.

[16] G. Y. Weintraub, L. Benkard, and B. Van Roy. Oblivious equilibrium: A mean field approximation for large-scale dynamic games. In *Advances in Neural Information Processing Systems*, volume 18. MIT Press, 2006.