Passive Dynamics in Mean Field Control

Ana Bušić and Sean Meyn

***Abstract*— Mean-field models are a popular tool in a variety of fields. They provide an understanding of the impact of interactions among a large number of particles or people or other "self-interested agents", and are an increasingly popular tool in distributed control.**

This paper considers a particular randomized distributed control architecture introduced in our own recent work. In numerical results it was found that the associated mean-field model had attractive properties for purposes of control. In particular, when viewed as an input-output system, its linearization was found to be minimum phase.

In this paper we take a closer look at the control model. The results are summarized as follows:

(i) The Markov Decision Process framework of Todorov is extended to continuous time models, in which the "control cost" is based on relative entropy. This is the basis of the construction of a family of Markovian generators, parameterized by a scalar $\zeta \in \mathbb{R}$.

(ii) A decentralized control architecture is proposed in which each agent evolves as a controlled Markov process. A central authority broadcasts a common control signal $\{\zeta_t\}$ to each agent. The central authority chooses $\{\zeta_t\}$ based on an aggregate scalar output of the Markovian agents.
This is the basis of the mean field model.

(iii) Provided the control-free system (with $\zeta \equiv 0$) is a reversible Markov process, the following identity holds for the transfer function G obtained from the linearization,

$$\mathbf{Re}\,(G(j\omega)) = \mathbf{PSD}_Y(\omega) \geq 0 \qquad \omega \in \mathbb{R}\,,$$

where the right hand side denotes the power spectral density for the output of any one of the individual Markov processes (with $\zeta \equiv 0$).

I. Introduction

Mean field models are a standard tool in physics when analyzing a large number of particles, where an individual particle has negligible impact upon the ensemble. Similar models are the foundation of competitive equilibrium theory in economics, and mean field models are increasingly popular in control theory [1], [4]–[6].

The present work considers application to distributed control, inspired by numerical results in our prior work [12] on automated demand response for a large collection of loads. The goal was to obtain *ancillary service* to help regulate the power grid, as in many prior works [2], [11].

This research is supported by the French National Research Agency grant ANR-12-MONU-0019, NSF grants CPS-0931416 and CPS-1259040, and US-Israel BSF Grant 2011506.

A.B. is with Inria and the Computer Science Dept. of École Normale Supérieure, Paris, France.

S.M. is with the Department of Electrical and Computer Engg. at the University of Florida, Gainesville.

The paper [12] focused on a large population of "on-off" loads, with special attention to residential pool pumps. The normal operation of a pool pump was modeled as a Markov decision process, which included as an exogenous input a regulation signal from a balancing authority. This resulted in an input-output system with input equal to the regulation signal, and output equal to the number of pools in operation. In the numerical example considered, the control system had some very attractive properties: Its linearization was stable, and simulations of 100,000 pools resulted in behavior very closely matched to the deterministic linear model obtained from linearization of the Markovian dynamics. Most important was the finding that the linearization was *minimum phase*. This is a valuable property in any control system.

In this paper we set out to see why these conclusions might be expected in greater generality.

To explain the goals of the paper we take a high-level look at the prior work [12]. Shown in Fig. 1 is a state transition diagram for the discrete-time Markovian model considered in [12]. The variables $p^\oplus$ and $p^\ominus$ indicate the probability of turning a pool pump on (respectively, off), which depends upon how long the pool has been off (respectively, on).

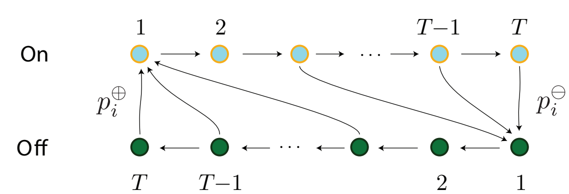

Fig. 1. State transition diagram for the pool-pump model.

A continuous time counterpart is described by a model on a continuous state space,

$$\mathsf{X} = \{(m,\tau) : m \in \{\oplus,\ominus\},\ \tau \geq 0\}.$$

If $X_t = (\oplus,\tau)$, this means that the pool pump has been operating for exactly τ seconds. If $t > \tau$, this implies that the pump was turned on at time $t-\tau$. The differential generator for this Markovian model is defined for functions $f\colon \mathsf{X} \to \mathbb{R}$ that are differentiable in their second variable. There are functions $q^\oplus(\,\cdot\,)$ and $q^\ominus(\,\cdot\,)$ such that for any such f,

$$\mathcal{D}f\,(x) = \begin{cases} q^\oplus(\tau)[f(\ominus,0) - f(\oplus,\tau)] + \frac{\partial}{\partial\tau} f\,(\oplus,\tau), \\ \hfill x = (\oplus,\tau) \\ q^\ominus(\tau)[f(\oplus,0) - f(\ominus,\tau)] + \frac{\partial}{\partial\tau} f\,(\ominus,\tau), \\ \hfill x = (\ominus,\tau) \end{cases}$$

Hence $q^\oplus(\tau)$ is the *jump rate* to the on-state, for a pool that has been off for τ seconds.

In this prior work the Markovian dynamics were controlled through a signal $\{\zeta_t\}$ that is broadcast to all pools. For this continuous time model, the jump rates would be modified by this signal. With N pools, on letting X^i_t denote the state of the ith pool, the following limit is shown to hold under mild assumptions:

$$\lim_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} \mathbb{I}\{X^i_t \in A\} = \mu_t(A)\,, \qquad A \subset \mathsf{X}. \tag{1}$$

If $\mathcal{D}_{\zeta_t}$ denotes the transformed generator at time t, then the limit is the solution to the differential equation,

$$\frac{d}{dt}\mu_t = \mu_t \mathcal{D}_{\zeta_t} \tag{2}$$

This means that for functions f satisfying the conditions above,

$$\frac{d}{dt}\int f(x)\mu_t(dx) = \int (\mathcal{D}_{\zeta_t} f\,(x))\mu_t(dx)\,.$$

The output is defined by a linear function of μ_t: For a function $\mathcal{U}\colon \mathsf{X}\to\mathbb{R}$,

$$y_t = \int \mathcal{U}(x)\mu_t(dx) \tag{3}$$

The coupled equations (2,3) describe a nonlinear input-output model of the form considered in this paper. The input ζ_t and output y_t are assumed to be real-valued.

Because (2) is linear in the "state" μ_t, and (3) is also linear in μ_t, it is easy to obtain a linearized model given some structure on the controlled generator. The question addressed in this paper is, *why should the linearized system have good properties for the purposes of control*?

We address this question for models in continuous time, since the analysis is most elegant in this setting. While many of the results in this paper can be extended to a general state space setting, for the remainder of the paper we restrict to a finite state space, $\mathsf{X} = \{x^1,\dots,x^d\}$. The family of generators $\{\mathcal{D}_\zeta : \zeta\in\mathbb{R}\}$ is a collection of $d\times d$ matrices, which are assumed to be a smooth function of the scalar parameter ζ.

The linear model is intended to approximate the nonlinear model near an equilibrium. To define the equilibrium we let π denote an invariant probability measure for the Markov process with generator $\mathcal{D}=\mathcal{D}_0$. This satisfies the invariance equation,

$$\pi\mathcal{D}\,(x^j) = \sum_{i=1}^{d} \pi(x^i)\mathcal{D}(x^i,x^j) = 0, \qquad x^j\in\mathsf{X}.$$

For the nonlinear model (2), if $\zeta_t\equiv 0$ and if $\mu_0=\pi$, then $\mu_t=\pi$ for all t.

The linear model evolves according to the d-dimensional linear state space equations,

$$\frac{d}{dt}\Phi_t = A\Phi_t + B\zeta_t, \qquad \gamma_t = C\Phi_t \tag{4}$$

The ith component of Φ_t is intended to approximate $\mu_t(x^i) - \pi(x^i)$, and $\gamma_t = C\Phi_t$ is intended to approximate $y_t - y^0$, with $y^0 = \sum_i \pi(x^i)\mathcal{U}(x^i)$. The $d\times d$ matrix A is the transpose of $\mathcal{D}_0=\mathcal{D}$, and $C_i = \mathcal{U}(x^i)-\pi(x^i)$ for each i. The matrix B is obtained from the derivative of $\mathcal{D}_\zeta$, at $\zeta=0$:

$$B_j = \sum_{i=1}^{d} \pi(x^i)\mathcal{D}'_0(x^i,x^j) \tag{5}$$

where $\mathcal{D}'_0(x^i,x^j)$ denotes the derivative of $\mathcal{D}_\zeta(x^i,x^j)$ with respect to ζ, evaluated at $\zeta=0$.

The transfer function for this model is $G(s) = C[Is-A]^{-1}B$, $s\in\mathbb{C}$. The minimum phase condition means that all zeros of G lie in the strict left half plane. In this paper we establish a stronger condition on the transfer function, under the assumption that the nominal Markov model is *reversible*. Through the procedure introduced in this paper, the linear dynamics satisfy the *positive real* condition,

$$\mathrm{Re}\,(G(j\omega)) \ge 0, \qquad \omega\in\mathbb{R}. \tag{6}$$

We obtain positivity by establishing the following identity,

$$\mathrm{Re}\,(G(j\omega)) = \mathrm{PSD}_Y(\omega) \qquad \omega\in\mathbb{R}\,, \tag{7}$$

where the right hand side denotes the power spectral density for $\{Y_t = \mathcal{U}(X_t)\}$ with $\boldsymbol{X}$ the stationary Markov process with generator $\mathcal{D}$.

The positive real condition is established only when the family of generators is constructed using the optimal control approach described in Section III. This recalls a similar result from linear optimal control theory, where it is known that the positive real condition holds for a certain transfer function, provided the system is controlled using state feedback based on linear-quadratic optimal control [7], [15]. We do not know if there is any connection between the main results of this paper, and these famous results from linear control theory.

The remainder of the paper is organized as follows: Section II contains an extension of Todorov's optimal control framework to Markovian models in continuous time. This is the basis of the mean-field model in Section III, and the main result that establishes the identity (7). An example is given to show that reversibility of the nominal model is necessary in general. Conclusions and discussion are contained in Section IV.

II. Construction of the controlled generator

Here we describe a stochastic optimal control problem in which the input is completely unconstrained. The optimization criterion will include a scalar weighting term ζ. The optimal solution will define the generator $\mathcal{D}_\zeta$ that is used in the mean field analysis that follows.

We consider a model in continuous time, with finite state space $\mathsf{X} = \{x^1,\dots,x^d\}$. The optimization is based on a *nominal* Markov process on this state space. Its generator (i.e., rate matrix) is defined for functions $f\colon\mathsf{X}\to\mathbb{R}$ via,

$$\begin{aligned}\mathcal{D}f\,(x) &= \sum_{x'} \mathcal{D}(x,x')f(x') \\ &= \lim_{t\downarrow 0}\frac{1}{t}\mathsf{E}[f(X_t)-f(X_0)\mid X_0=x]\,, \quad x\in\mathsf{X}.\end{aligned} \tag{8}$$

The transition semigroup is the exponential, $P^t = e^{t\mathcal{D}}$, $t \ge 0$. The Markov process is assumed to be irreducible, so that there is a unique invariant probability measure π: Interpreted as a row vector, it satisfies $\pi\mathcal{D} = 0$ and $\pi P^t = \pi$ for $t \ge 0$.

For fixed T and fixed initial condition $X(0) = x$, let p^0 denote the probability distribution for the stochastic process $\{X_t : 0 \le t \le T\}$ for the nominal model.

This is an unusual stochastic control problem because there is no explicit "input". Any modification p of p^0 is permitted. A particular optimization objective will ensure that an optimal solution is Markovian.

It is assumed that a utility function $\mathcal{U}\colon \mathsf{X} \to \mathbb{R}$ is given, that represents some benefit as a function of state. The cost of deviation $p \ne p^0$ is defined by Kullback-Leibler divergence, denoted $D(p\|p^0)$. The T-stage welfare is defined as the difference,

$$\mathcal{W}_T(p) = \zeta \mathsf{E}_p\Big[\int_0^T \mathcal{U}(X_t)\,dt\Big] - D(p\|p^0)$$

where in the expectation $\{X_t : 0 \le t \le T\}$ is distributed according to p. The maximizer exists, and is denoted p^* (or p^*_T when the time horizon is emphasized).

The parameter ζ is a real scalar. For notational simplicity, until Section III we take $\zeta = 1$.

Before proceeding with the formula for p^*, it is helpful to recall the definition of divergence in this sample-path setting. We let $\mathcal{X}_T$ denote the sigma algebra generated by the stochastic process $\{X_t : t \le T\}$. A log-likelihood ratio is interpreted as an $\mathcal{X}_T$-measurable random variable: If p admits a log likelihood ratio $L = \log(p/p^0)$, this means that for any $\mathcal{X}_T$-measurable random variable F we can write,

$$\mathsf{E}_p[F] = \mathsf{E}[e^L F] \tag{9}$$

where the expectation on the left is under p, and the expectation on the right is under p^0 (the subscript is not used for the nominal model). The K-L divergence is then defined to be,

$$D(p\|p^0) = \mathsf{E}_p[L]$$

If L does not exist, then $D(p\|p^0) = \infty$.

Proposition 2.1: Suppose that the nominal is irreducible, and that $X(0) = x$ is specified. Then p^*_T is unique, and is given by the twisted (or 'tilted') distribution that is uniquely defined by the log likelihood ratio,

$$L^* = -\Lambda^*_T + \int_0^T \mathcal{U}(X_t)\,dt \tag{10}$$

The optimal welfare $\mathcal{W}_T(p^*_T)$ coincides with the constant Λ^*_T appearing in (10), which is equal to the cumulative log-moment generating function,

$$\Lambda^*_T = \log\Big(\mathsf{E}\Big[\exp\Big(\int_0^T \mathcal{U}(X_t)\,dt\Big)\Big]\Big) \tag{11}$$

where the expectation is w.r.t. the nominal model.

Proof of Proposition 2.1: Optimality of p^* (with log likelihood ratio (10)) is a consequence of Kullback's inequality (see eqn (4.5) of [10]). See also Theorem 3.1.2 of [3] for a version of this result on a finite probability space. The papers [13], [16] contain more background and other applications of this result.

An explicit value for the optimal welfare follows: We have,

$$D(p^*\|p) = \mathsf{E}_{p^*}[L^*] = -\Lambda^*_T + \mathsf{E}_{p^*}\Big[\int_0^T \mathcal{U}(X_t)\,dt\Big]$$

and consequently,

$$\max_p \mathcal{W}_T(p) = \mathcal{W}_T(p^*) = \Lambda^*_T$$

The formula (11) follows from the fact that $p^* = e^{L^*}p^0$ defines a probability distribution:

$$1 = \mathsf{E}[e^{L^*}] = e^{-\Lambda^*_T}\mathsf{E}\Big[\exp\Big(\int_0^T \mathcal{U}(X_t)\,dt\Big)\Big] \tag{12}$$

□

While the optimal probability measure p^* is Markovian, it is not time-homogeneous.

We now consider an infinite horizon optimization problem: Find a Markov process for which the associated family of distributions $\{\check{p}_T\}$, with initial condition $X(0) = x$, attain the limit,

$$\lim_{T\to\infty} \frac{1}{T}\mathcal{W}_T(\check{p}_T) = \mathcal{W}^*_\infty := \lim_{T\to\infty} \frac{1}{T}\mathcal{W}_T(p^*_T) \tag{13}$$

This has a solution defined by a time-homogeneous Markov process, whose generator is denoted $\check{\mathcal{D}}$. Its construction is based on the solution to an eigenvector problem: Let $I_{\mathcal{U}}$ denote the diagonal matrix $I_{\mathcal{U}} = \mathrm{diag}\,(\mathcal{U}(x^1), \dots, \mathcal{U}(x^d))$, and let $v\colon \mathsf{X} \to (0,\infty)$ denote a non-trivial solution to the eigenvector problem,

$$[\mathcal{D} + I_{\mathcal{U}}]v = \Lambda v \tag{14}$$

where Λ is the eigenvalue of $\mathcal{D} + I_{\mathcal{U}}$ with maximal real-part (the Perron-Froebenius eigenvalue [9]).

Proposition 2.2: The following hold under the assumptions of Proposition 2.1:

(i) $\mathcal{W}^*_\infty = \Lambda$; the eigenvalue appearing in (14).

(ii) The generator for the Markov process that attains the optimal average welfare $\mathcal{W}^*_\infty$ is obtained by normalizing $\mathcal{D} + I_{\mathcal{U}}$, and applying a similarity transformation using I_v:

$$\check{\mathcal{D}} = \mathbb{I}_v^{-1}[\mathcal{D} + I_{\mathcal{U}} - \Lambda I]\mathbb{I}_v \tag{15}$$

(iii) For each T, the welfare for the distribution $\check{p}_T$ is given by,

$$\mathcal{W}_T(\check{p}_T) = T\mathcal{W}^*_\infty - \mathsf{E}\Big[\log\Big(\frac{v(X_T)}{v(X_0)}\Big)\Big]$$

The eigenvector equation (14) implies that $\sum_{x'} \check{\mathcal{D}}(x,x') = 0$ for all $x \in \mathsf{X}$, as required for a Markovian generator.

Proof of Proposition 2.2: Part (i) is essentially known: From (12) it follows that $\mathcal{W}_\infty^*$ is the multiplicative-ergodic limit,

$$\mathcal{W}_\infty^* = \lim_{T\to\infty} \frac{1}{T} \log\Bigl(\mathsf{E}\Bigl[\exp\Bigl(\int_0^T \mathcal{U}(X_t)\, dt\Bigr)\Bigr]\Bigr) \tag{16}$$

The right hand side is denoted $\Lambda(\mathcal{U})$ in [9], where it is shown in far greater generality that the multiplicative-ergodic limit $\Lambda(\mathcal{U})$ coincides with the eigenvector Λ.

However, the proof that $\Lambda = \Lambda(\mathcal{U})$ and the proof of the remaining claims of the theorem will follow from a representation of the Markov process with generator $\check{\mathcal{D}}$.

Let $\check{p}$ denote the probability measure on sample paths, with given initial condition $X(0) = x$. For finite T, if F is an $\mathcal{X}_T$-measurable functional then, for the Markovian model with generator $\check{\mathcal{D}}$ we have (exactly as in (9)),

$$\mathsf{E}_{\check{p}}[F] = \mathsf{E}[e^{\check{L}_T} F]$$

where the expectation is with respect to p^0, and

$$\check{L}_T = \log\Bigl(\frac{v(X_T)}{v(X_0)}\Bigr) + \int_0^T [\mathcal{U}(X_t) - \Lambda]\, dt$$

Using the fact that $\check{p}_T$ is a probability distribution gives,

$$1 = \mathsf{E}[e^{\check{L}_T}] = e^{-\Lambda T}\mathsf{E}\Bigl[\frac{v(X_T)}{v(X_0)} \exp\Bigl(\int_0^T \mathcal{U}(X_t)\, dt\Bigr)\Bigr]$$

The identity $\Lambda = \Lambda(\mathcal{U})$ follows: Since v is strictly positive and finite-valued, it follows that $\check{p}$ is infinite-horizon optimal:

$$\lim_{T\to\infty} \frac{1}{T}\mathcal{W}_T(\check{p}) = \Lambda(\mathcal{U}) = \lim_{T\to\infty} \frac{1}{T}\Lambda_T^* = \mathcal{W}_\infty^*$$

This establishes (ii).

Given the formula for $\check{L}_T$, the total welfare at time T using $\check{p}$ is thus,

$$\begin{aligned}\mathcal{W}_T(\check{p}) &= \mathsf{E}_{\check{p}}\Bigl[\int_0^T \mathcal{U}(X_t)\, dt\Bigr] - D(\check{p}_T \| p_T^0) \\ &= -\mathsf{E}\Bigl[\log\Bigl(\frac{v(X_T)}{v(X_0)}\Bigr)\Bigr] + \Lambda T\end{aligned}$$

which establishes (iii). □

III. Mean field model and its linear approximation

Up to now we have only two generators: The nominal generator $\mathcal{D}$, and its transformation (15) obtained as the solution to an optimal control problem. We next construct a parameterized family of generators denoted $\{\mathcal{D}_\zeta : \zeta \in \mathbb{R}\}$. For each ζ, this is obtained as the infinite-horizon optimal control solution of the previous section, with the finite-horizon welfare modified as follows,

$$\mathcal{W}_T(p) = \zeta\mathsf{E}_p\Bigl[\int_0^T \mathcal{U}(X_t)\, dt\Bigr] - D(p\|p^0)$$

If $\zeta = 0$ then $p^* = p^0$. For arbitrary ζ, the generator $\mathcal{D}_\zeta$ that solves the infinite-horizon optimal control problem is of the form (15), in which $v = v_\zeta$ is a solution to the eigenvector problem,

$$[\mathcal{D} + \zeta I_{\mathcal{U}}]v_\zeta = \Lambda_\zeta v_\zeta \tag{17}$$

where

$$\Lambda_\eta = \Lambda(\zeta\mathcal{U}) := \lim_{T\to\infty} \frac{1}{T} \log \mathsf{E}\Bigl[\exp\Bigl(\zeta \int_0^T \mathcal{U}(X_t)\, dt\Bigr)\Bigr]$$

The solutions to (17) are used to define the continuous family of generators, $\mathcal{D}_\zeta = \mathbb{I}_{v_\zeta}^{-1}[\mathcal{D} + \zeta I_{\mathcal{U}} - \Lambda_\zeta I]\mathbb{I}_{v_\zeta}$, or componenetwise,

$$\mathcal{D}_\zeta(x^i, x^j) = \frac{v_\zeta(x^j)}{v_\zeta(x^i)}[\mathcal{D}(x^i, x^j) + (\zeta\mathcal{U}(x^i) - \Lambda_\zeta)\mathbb{I}\{x^j = x^i\}] \tag{18}$$

The mean field model is the nonlinear state space model defined by (2), which in this finite state space setting becomes,

$$\frac{d}{dt}\mu_t(x) = \sum_{x^i \in \mathsf{X}} \mu_t(x^i)\mathcal{D}_{\zeta_t}(x^i, x) \tag{19}$$

Recall that π denotes the unique invariant probability measure for the nominal model. The nominal model is called reversible if the detailed-balance equations hold:

$$\pi(x^i)\mathcal{D}_0(x^i, x^j) = \pi(x^j)\mathcal{D}_0(x^j, x^i), \qquad x^i, x^j \in \mathsf{X}$$

Theorem 3.1: Suppose that the nominal model is reversible. Then its linearization (4) satisfies,

$$\mathrm{Re}\, G(j\omega) = \mathrm{PSD}_Y(\omega), \qquad \omega \in \mathbb{R}, \tag{20}$$

where

$$G(s) = C[Is - A]^{-1}B \text{ for } s \in \mathbb{C}. \tag{21}$$

The proof of the proposition involves a sequence of steps. The first steps are contained in Section III-A: The power spectral density for $\boldsymbol{Y}$ can be expressed in terms of the family of *resolvent matrices* for the Markov process, and these can be interpreted as a component of the transfer function G. The formula (20) is based on these results, and a closer look at the linearization contained in Section III-B, which closes with a proof of Thm. 3.1.

Before proceeding with these technical arguments, we give an example to show that the positive real condition may not hold if the nominal model is not reversible. Moreover, this example shows that without reversibility, the linearization may not be minimum phase.

Consider the Markov chain with eight states, whose transition diagram and generator are shown in Fig. 2. This Markov process cannot be reversible because some transitions are uni-directional. For example, an immediate transition from 3 to 1 is possible, but not from 1 to 3.

In the notation of this paper we have $d = 8$, and we take $x^i = i$ for $1 \le i \le 8$. The utility function $\mathcal{U}\colon \mathsf{X} \to \mathbb{R}$ is taken to be $\mathcal{U}(x) = x$.

Shown in Fig. 3 is a Nyquist plot and pole-zero plot for the transfer function $G(s) = C[Is - A]^{-1}B$ with $a = c = 10$ and $b = 1$. The Nyquist plot shows that the system is not positive real, and the pole-zero plot shows that the system is

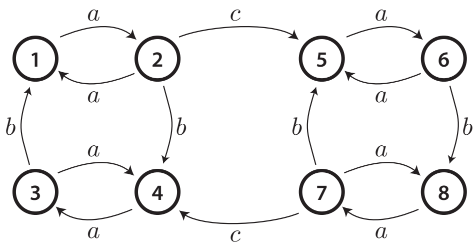

$$\mathcal{D} = \begin{bmatrix} -a & a & 0 & 0 & 0 & 0 & 0 & 0 \\ a & -(a+b+c) & 0 & b & c & 0 & 0 & 0 \\ b & 0 & -(a+b) & a & 0 & 0 & 0 & 0 \\ 0 & 0 & a & -a & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & -a & a & 0 & 0 \\ 0 & 0 & 0 & 0 & a & -(a+b) & 0 & b \\ 0 & 0 & 0 & c & b & 0 & -(a+b+c) & a \\ 0 & 0 & 0 & 0 & 0 & 0 & a & -a \end{bmatrix}$$

Fig. 2. State transition diagram and generator for a Markov process that is not reversible. The transfer function for the linearized mean field model is neither positive real nor minimum phase.

not minimum phase: there is a zero in the right half plane, at approximately $s = +9$. A common pole-zero pair at $s = 0$ is not shown in the plot.

This example demonstrates that the positive real condition obtained in Thm. 3.1 requires assumptions on the nominal model. Reversibility is used to obtain the identity (20) that implies the positive real condition. We know of no alternate assumptions that imply the positive-real condition.

A. Resolvents and transfer functions

The family of resolvent matrices are defined for $s \in \mathbb{C}$ by the integral,

$$R_s = \int_0^\infty e^{-st} P^t \, dt$$

This is well defined whenever $\mathrm{Re}\,(s) < 0$. It can be shown using (8) (or the representation $P^t = e^{t\mathcal{D}}$) that the resolvent equation holds $\mathcal{D}R_s = sR_s - I$; equivalently,

$$R_s = [sI - \mathcal{D}]^{-1} \tag{22}$$

We already see that this forms a component of $G(s)$ in (21) (recall that $A = \mathcal{D}^{\mathsf{T}}$). Consequently, for each s satisfying $\mathrm{Re}\,(s) < 0$,

$$G(s) = C[Is - A]^{-1}B = B^{\mathsf{T}}[Is - A^{\mathsf{T}}]^{-1}C^{\mathsf{T}} = B^{\mathsf{T}} R_s C^{\mathsf{T}} \tag{23}$$

Based on this identity, the following result shows that the frequency response is similar to a cross-power spectral density:

Proposition 3.2: The frequency response for the transfer function (21) with $A = \mathcal{D}^{\mathsf{T}}$ can be expressed, for $\omega \in \mathbb{R}$, by

$$G(j\omega) = \int_0^\infty e^{-j\omega t} \mathsf{E}_\pi[f(X_0)g(X_t)]\, dt$$

where $f(x^i) = B_i/\pi(x^i)$ and $g(x^i) = C_i = \widetilde{\mathcal{U}}(x^i)$, $x^i \in \mathsf{X}$.

Proof: The proof begins with the representation (23), which holds by definition whenever $\mathrm{Re}\,(s) < 0$. From the definition of the resolvent matrix, (23) gives,

$$\begin{aligned} G(s) &= \int_0^\infty e^{-st} \Big[\sum_{i,j} P^t(x^i, x^j) B_i C_j \Big]\, dt \\ &= \int_0^\infty e^{st} \mathsf{E}_\pi[f(X_0)g(X_t)]\, dt \end{aligned}$$

where the final equality follows from the definition of f and g.

To complete the proof we must extend (23) to $s = -j\omega$, for which $\mathrm{Re}\,(s) = 0$. For this we note that $g(x^i) = C_i = \widetilde{\mathcal{U}}(x^i)$, so that

$$\lim_{t\to\infty} \mathsf{E}_\pi[f(X_0)g(X_t)] = 0,$$

where the convergence rate is exponential. □

B. Linearization

To apply (3.2) we require a representation of the matrix B defined in (5). For this we normalize the eigenvector so that $v_\zeta(x^1) = 1$ for all ζ; this is without loss of generality since the components $\mathcal{D}_\zeta(x^i, x^j)$ of the generator (18) are defined in terms of the ratio $v_\zeta(x^j)/v_\zeta(x^i)$.

Let h_0 denote the solution to Poisson's equation,

$$\mathcal{D}_0 h_0 = -\widetilde{\mathcal{U}} \tag{24}$$

with boundary condition $h_0(x^1) = 0$. For a finite-state space Markov process, one solution to Poisson's equation is given by

$$h(x) = R_0\widetilde{\mathcal{U}}\,(x) = \int_0^\infty \mathsf{E}[\widetilde{\mathcal{U}}(X_t) \mid X_0 = x]\, dt, \quad x \in \mathsf{X} \tag{25}$$

and then we take $h_0(x) = h(x) - h(x^1)$, $x \in \mathsf{X}$.

Let $\mathcal{D}^\dagger$ denote the generator for the time-reversed process,

$$\mathcal{D}^\dagger(x^i, x^j) = \pi(x^j)\mathcal{D}(x^j, x^i)\frac{1}{\pi(x^i)}, \qquad x^i, x^j \in \mathsf{X}$$

Proposition 3.3: The entries of B are given by,

$$B_i = -\pi(x^i)[\mathcal{D}^\dagger h_0\,(x^i) + \mathcal{D}h_0\,(x^i)]$$

If the process is reversible, then $B_i = 2\pi(x^i)C_i = 2\pi(x^i)\widetilde{\mathcal{U}}(x^i)$.

To prove the proposition we first need the following formulae for the derivatives of Λ_ζ and v_ζ. We omit the proof, which is similar to the discrete-time case [8], [9].

Lemma 3.4: The log-moment generating function has derivative at the origin given by,

$$\frac{d}{d\zeta}\Lambda_\zeta\Big|_{\zeta=0} = y^0 = \sum_i \pi(x^i)\mathcal{U}(x^i)$$

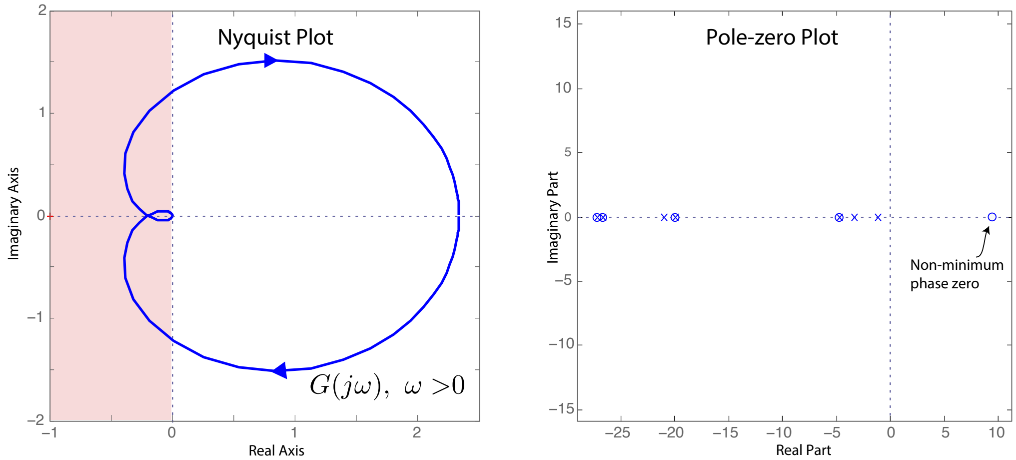

Fig. 3. Nyquist plot and pole-zero plot for linearization

The derivative of the eigenvector is the solution to Poisson's equation,

$$\frac{d}{d\zeta} v_\zeta(x^i)\Big|_{\zeta=0} = h_0(x^i), \qquad x^i \in \mathsf{X}.$$

□

Proof of Prop. 3.3: Applying the lemma to (18) gives, for all $x^i, x^j \in \mathsf{X}$,

$$\begin{aligned}\mathcal{D}_\zeta(x^i,x^j) &= [1-\zeta h_0(x^i)]\mathcal{D}(x^i,x^j)[1+\zeta h_0(x^j)] \\ &\quad + [\zeta\mathcal{U}(x^i) - \Lambda_\zeta]\mathbb{I}\{x^i = x^j\} + o(\zeta) \\ &= [1-\zeta h_0(x^i)]\mathcal{D}(x^i,x^j)[1+\zeta h_0(x^j)] \\ &\quad + \zeta[\mathcal{U}(x^i) - y^0]\mathbb{I}\{x^i = x^j\} + o(\zeta)\end{aligned}$$

From the definition $\widetilde{\mathcal{U}} = \mathcal{U} - y^0$, we conclude that the derivative is given by,

$$\begin{aligned}\frac{d}{d\zeta}\mathcal{D}_\zeta(x^i,x^j)\Big|_{\zeta=0} &= -h_0(x^i)\mathcal{D}(x^i,x^j) + \mathcal{D}(x^i,x^j)h_0(x^j) \\ &\quad + \widetilde{\mathcal{U}}(x^i)\mathbb{I}\{x^i = x^j\}\end{aligned}$$

The entries of the matrix B are thus given by,

$$\begin{aligned}B_j &= \sum_{x^i} \pi(x^i)\Big(-h_0(x^i)\mathcal{D}(x^i,x^j) + \mathcal{D}(x^i,x^j)h_0(x^j) \\ &\qquad + \widetilde{\mathcal{U}}(x^i)\mathbb{I}\{x^i = x^j\}\Big) \\ &= -\sum_{x^i} h_0(x^i)\pi(x^i)\mathcal{D}(x^i,x^j) + \pi(x^j)\widetilde{\mathcal{U}}(x^j)\end{aligned}$$

where in the second identity we used the invariance equation, $\sum_{x^i}\pi(x^i)\mathcal{D}(x^i,x^j) = 0$. The second identity is equivalent to the desired representation. □

Proof of Thm. 3.1: Prop. 3.3 tells us that under reversibility we have $B_i = 2\pi(x^i)\widetilde{\mathcal{U}}(x^i)$, and hence in the notation of Prop. 3.2,

$$f(x^i) = B_i/\pi(x^i) = 2\widetilde{\mathcal{U}}(x^i), \qquad g(x^i) = C_i = \widetilde{\mathcal{U}}(x^i)$$

Prop. 3.2 and Prop. 3.3 then give,

$$\begin{aligned}G(j\omega) &= \int_0^\infty e^{-j\omega t}\mathsf{E}_\pi[f(X_0)g(X_t)]\,dt \\ &= 2\int_0^\infty e^{-j\omega t}\mathsf{E}_\pi[\widetilde{\mathcal{U}}(X_0)\widetilde{\mathcal{U}}(X_t)]\,dt\end{aligned}$$

Thm. 3.1 thus follows:

$$\begin{aligned}\mathrm{Re}\,G(j\omega) &= 2\mathrm{Re}\int_0^\infty e^{-j\omega t}\mathsf{E}_\pi[\widetilde{\mathcal{U}}(X_0)\widetilde{\mathcal{U}}(X_t)]\,dt \\ &= \int_{-\infty}^\infty e^{-j\omega t}\mathsf{E}_\pi[\widetilde{\mathcal{U}}(X_0)\widetilde{\mathcal{U}}(X_t)]\,dt\end{aligned}$$

□

IV. Conclusions

This paper gives a general condition under which the linearization of a mean field model is positive-real.

The linearization around $\zeta = 0$ is a natural choice, but the main result of the paper can be extended to any constant value: If $\mathcal{D}$ is reversible, then so is $\mathcal{D}_\zeta$ for each fixed $\zeta \in \mathbb{R}$. Based on this observation, it is possible to show that the linearization about any fixed value of ζ is positive real under the assumptions of Thm. 3.1. This suggests an open question: Is the nonlinear model with state equation (2) passive? Passivity would be a valuable property for the purposes of control.

There are many open questions in the context of design. Can we obtain more general sufficient conditions for the positive real condition, the weaker minimum phase condition, or the stronger passivity condition for the nonlinear model?

To relax the assumptions of Thm. 3.1, it is likely that we will require application of the Kalman-Yakubovich-Popov Lemma, which provides an algebraic characterization of the passive real condition [14].

We are currently considering these theoretical questions, and applications to problems in decentralized control, especially in power systems settings.

REFERENCES

[1] V.S. Borkar and R. Sundaresan. Asympotics of the invariant measure in mean field models with jumps. *Stochastic Systems*, 2(2):322–380, 2012.

[2] R. Couillet, S.M. Perlaza, H. Tembine, and M. Debbah. Electrical vehicles in the smart grid: A mean field game analysis. *Selected Areas in Communications, IEEE Journal on*, 30(6):1086–1096, 2012.

[3] A. Dembo and O. Zeitouni. *Large Deviations Techniques And Applications*. Springer-Verlag, New York, second edition, 1998.

[4] N. Gast, B. Gaujal, and J.-Y. Le Boudec. Mean field for Markov decision processes: From discrete to continuous optimization. *IEEE Trans. Automat. Control*, 57(9):2266–2280, 2012.

[5] Peng Guan, M. Raginsky, and R. Willett. Online Markov decision processes with Kullback-Leibler control cost. In *American Control Conference (ACC)*, pages 1388–1393, 2012.

[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] Rudolf Emil Kalman. When is a linear control system optimal? *Journal of Basic Engineering*, 86:51, 1964.

[8] I. Kontoyiannis and S. P. Meyn. Spectral theory and limit theorems for geometrically ergodic Markov processes. *Ann. Appl. Probab.*, 13:304–362, 2003. Presented at the INFORMS Applied Probability Conference, NYC, July, 2001.

[9] I. Kontoyiannis and S. P. Meyn. Large deviations asymptotics and the spectral theory of multiplicatively regular Markov processes. *Electron. J. Probab.*, 10(3):61–123 (electronic), 2005.

[10] S. Kullback. Certain inequalities in information theory and the Cramer-Rao inequality. *Ann. Math. Statist.*, 25(4):745–751, 1954.

[11] Zhongjing Ma, D. Callaway, and I. Hiskens. Decentralized charging control for large populations of plug-in electric vehicles: Application of the Nash certainty equivalence principle. In *Control Applications (CCA), 2010 IEEE International Conference on*, pages 191 –195, sept. 2010.

[12] Sean Meyn, Prabir Barooah, Ana Bušić, and Jordan Ehren. Ancillary service to the grid from deferrable loads: the case for intelligent pool pumps in Florida (Invited). In *Proceedings of the 52nd IEEE Conf. on Decision and Control*, 2013.

[13] Paolo Pra, Lorenzo Meneghini, and WolfgangJ. Runggaldier. Connections between stochastic control and dynamic games. *Mathematics of Control, Signals and Systems*, 9(4):303–326, 1996.

[14] Anders Rantzer. On the Kalman-Yakubovich-Popov lemma. *Systems & Control Letters*, 28(1):7 – 10, 1996.

[15] M Safonov and Michael Athans. Gain and phase margin for multiloop lqg regulators. *IEEE Trans. Automat. Control*, 22(2):173–179, 1977.

[16] J. Unnikrishnan, D. Huang, S. P. Meyn, A. Surana, and V. V. Veeravalli. Universal and composite hypothesis testing via mismatched divergence. *IEEE Trans. Inform. Theory*, 57(3):1587 –1603, 2011.