Optimal steering of Ensembles

Yongxin Chen, Tryphon Georgiou and Michele Pavon

Abstract

We consider ensembles of particles or systems, all obeying the same dynamics or, equivalently, dynamical systems with uncertain states. In either case, we discuss i) the problem of optimally steering ensembles to a desired probability density, and ii) the problem of maintaining an ensemble at a specified admissible stationary distribution. In the context of linear dynamics and Gaussian distributions, our topic relates classical LQR and LQG theory and covariance control, while more generally it pertains to the controllability and optimal control of the Fokker-Planck and Liouville equations.

Introduction

The paradigm that we discuss encompasses various generalization of classical linear quadratic regulator (LQR) and the linear quadratic Gaussian regulator (LQG). More specifically, we focus on optimal control problems, with and without stochastic excitation, that require steering an ensemble of dynamical systems, obeying identical dynamics, between two end-points in time with control input of minimum averaged energy. Typically, besides knowledge of the system dynamics, the data for the problem consist of the starting and terminal distribution of the ensemble. Equivalently, the problem can also be conceived as that of steering a dynamical with state uncertainty between the two end-points in time. Another variation of the problem is to maintain the ensemble at a pre-specified admissible distribution with control input of minimum power. The presentation will be largely based on our recent work in [1], [2], [3], [4], [5], [6], [7], [8] and will highlight the theme and key contributions in these publications.

It is envisioned that the framework in this work will impact problems of steering particle beams via a time-varying potential, swarms (UAV's, large collection of micro-satellites, etc.), as well as the modeling of the flow and collective motion of particles, clouds, insects, birds, etc. between end-point distributions. It is worth underscoring that, as eloquently stated by R. Brockett in [9, page 23], from a controls perspective "important limitations standing in the way of the wider use of optimal control can be circumvented by explicitly acknowledging that in most situations the apparatus implementing the control policy will judged on its ability to cope with a distribution of initial states, rather than a single state." The aim of the research in the publications cited above has been to address relevant questions in this suggested overarching program.

Y. Chen and T.T. Georgiou are with the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, Minnesota MN 55455, USA; email: {chen2468,tryphon}@umn.edu

M. Pavon is with the Dipartimento di Matematica, Università di Padova, via Trieste 63, 35121 Padova, Italy; email: pavon@math.unipd.it

The research was partially supported by the NSF under Grant ECCS-1509387, the AFOSR under Grants FA9550-12-1-0319 and FA9550-15-1-0045, and by the University of Padova Research Project CPDA 140897.

A. Linear-dynamics & Gaussian ensembles

Consider the controlled evolution

$$\begin{aligned} dx^u(t) &= Ax^u(t)dt + Bu(t)dt + B_1 dw(t), \\ x^u(0) &= x_0 \text{ a.s.} \end{aligned} \tag{1}$$

Here A, B, B_1 take values in $\mathbb{R}^{n\times n}$, $\mathbb{R}^{n\times m}$, $\mathbb{R}^{n\times p}$, respectively, x_0 is an n-dimensional, zero-mean Gaussian vector with covariance Σ_0 which is independent of the standard p-dimensional Wiener process $\{w(t) \mid 0 \leq t \leq T\}$, Σ_0 is positive definite, and $T \leq \infty$ represents the end point of a time interval of interest. The standard paradigm of Linear Quadratic theory is to specify a target value for the state vector, e.g., the origin for simplicity, and impose a quadratic penalty on possible deviations. In contrast, the question that we are interested in this context is to determine *the admissible* values of the state covariance $\Sigma(T)$ that can be obtained through feedback control and to specify corresponding control inputs of minimal (quadratic) effort. More generally, we are interested in specifying the terminal distribution.

– Control of state-statistics over a finite time-window:

Assuming the control input in state-feedback form,

$$u(x,t) = -K(t)x, \tag{2}$$

state statistics remain Gaussian and the state covariance $\Sigma(t) := \mathbb{E}\{x(t)x(t)'\}$ of (1) satisfies the Lyapunov differential equation

$$\dot{\Sigma}(t) = (A - BK(t))\Sigma(t) + \Sigma(t)(A - BK(t))' + B_1 B_1' \tag{3}$$

with $\Sigma(0) = \Sigma_0$. Regardless of the choice of $K(\cdot)$, it is easy to see that (3) specifies dynamics that leave invariant the cone of positive semi-definite symmetric matrices. Thus, our interest is in our ability to specify $\Sigma(T) = \Sigma_T$ via a suitable choice of $K(\cdot)$, which therefore necessitates that the solution $\Sigma(t)$ remains in the positive cone for all $t \in [0, T]$.

Letting $U(t) := -\Sigma(t)K(t)'$ and $Q := B_1 B_1'$, we are led to study controllability of the matrix-valued differential Lyapunov system

$$\dot{\Sigma}(t) = A\Sigma(t) + \Sigma(t)A' + BU(t)' + U(t)B' + Q. \tag{4}$$

Interestingly, (A, B) is a controllable pair if and only if we can select $U(\cdot)$ in (4) to steer $\Sigma(\cdot)$ between given positive definite end points Σ_0 and Σ_T while remaining within the non-negative cone [2, Theorem 3], cf. [9], [10]. The theorem in [2] requires time-invariant dynamics, though we believe that the a corresponding statement holds for time-varying systems as well.

We now return to the question of optimality with regard to averaged input energy, namely

$$J(u) := \mathbb{E}\left\{\int_0^T u(t)'u(t)\,dt\right\} < \infty, \tag{5}$$

Resorting to standard LQG arguments it can be shown that optimal solutions are characterized by the following nonlinearly coupled Riccati equations

$$\begin{aligned}
\dot{\Pi} &= -A'\Pi - \Pi A + \Pi BB'\Pi && \text{(6a)}\\
\dot{\mathrm{H}} &= -A'\mathrm{H} - \mathrm{H}A - \mathrm{H}BB'\mathrm{H} && \text{(6b)}\\
&\quad + (\Pi + \mathrm{H})\,(BB' - B_1B_1')\,(\Pi + \mathrm{H})\\
\Sigma_0^{-1} &= \Pi(0) + \mathrm{H}(0) && \text{(6c)}\\
\Sigma_T^{-1} &= \Pi(T) + \mathrm{H}(T) && \text{(6d)}
\end{aligned}$$

where $\mathrm{H}(t) := \Sigma(t)^{-1} - \Pi(t)$. Indeed, assuming that $\{(\Pi(t), \mathrm{H}(t)) \mid 0 \le t \le T\}$ satisfy (6a-6d), the feedback control law

$$u(x,t) = -K(t)x. \tag{7}$$

with $K(t) = B(t)'\Pi(t)$ can be shown to be an optimal solution to the problem and that it is unique.

For the special case where $B(t) \equiv B_1(t)$, that is, the case where the control inputs and the noise enter the system through identical channels, the nonlinear system of equations (6) can be solved in closed form and is detailed in [1]. For the case where control and noise excitation channels differ, [2] provides a numerical approach for constructing suboptimal controls.

– *Admissible stationary state-distributions:*

A stationary counterpart of the above problem is to determine, for A, B, B_1 independent of time and (A, B) controllable, all admissible stationary state covariances Σ corresponding to choices of state-feedback control $u = -Kx$ that ensure $A - BK$ is a Hurwitz matrix.

Interestingly, while any Gaussian distribution can be reached in finite time with a suitable control law, not all non-negative matrices Σ are admissible as covariances of stationary state Gaussian distributions for the given system and, accordingly, a suitable state feedback. For that to be the case, Σ, A, B, B_1 must satisfy the following algebraic condition:

$$\operatorname{rank}\begin{bmatrix} A\Sigma + \Sigma A' + B_1B_1' & B\\ B & 0\end{bmatrix} = \operatorname{rank}\begin{bmatrix} 0 & B\\ B & 0\end{bmatrix}. \tag{8}$$

The precise statement is given in [2, Section III.B] and, once again, a computational approach to obtain suboptimal control laws is provided.

– *An academic example*

In order to illustrate the problem being discussed, we consider the following model for particles experiencing random displacement (i.e., the noise term dw has a direct impact on their position):

$$\begin{aligned}
dx(t) &= v(t)dt + dw(t) && \text{(9a)}\\
dv(t) &= u(t)dt. && \text{(9b)}
\end{aligned}$$

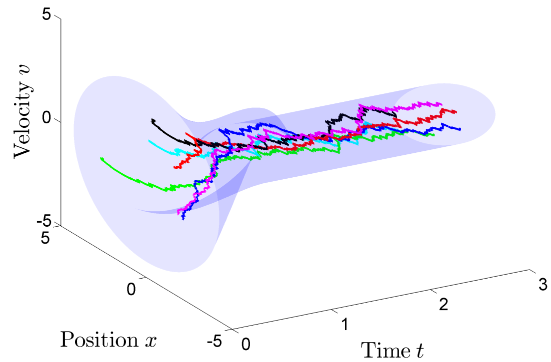

Fig. 1: State trajectories of system in (9)

Here, $u(t)$ is the control input (force) at our disposal, $x(t)$ represents position and $v(t)$ velocity (integral of acceleration due to input forcing), while $w(t)$ represents random displacement due to impulsive accelerations. Alternatively, $\int^t v(\tau)d\tau$ may represent actual position while $x(t)$ may represent noisy measurement of position.

The purpose of the example is to highlight a case where the control is handicapped compared to the effect of noise. Indeed, the displacement $w(t)$ directly affects $x(t)$ while the control effort $u(t)$ needs to be integrated before it mitigates the effect of $w(t)$ on the position $x(t)$ of the particles. We choose

$$\Sigma_1 = \begin{bmatrix} 1 & -1/2\\ -1/2 & 1/2\end{bmatrix} \tag{10}$$

as a candidate stationary state-covariance, which, as it turns out, can be maintained with state-feeback gain $K = [1,\ 1]$. Next, we steer the spread of the particles from an initial Gaussian distribution with $\Sigma_0 = 2I$ at $t = 0$ to the terminal marginal Σ_1 at $t = 1$, and from there on, since Σ_1 is an admissible stationary state-covariance, we maintain the distribution as indicated. Figure 1 depicts typical sample paths in phase space as functions of time.

B. Output feedback

It turns out that when the state is not directly accessible and the control is a function of the observation process

$$dy(t) = Cx(t)dt + Ddv(t), \tag{11}$$

the achievable state covariances are constrained by the optimal estimation error-covariance $P(\cdot)$ provided by Kalman filter:

$$\begin{aligned}
d\hat{x}(t) &= A\hat{x}(t)dt + Bu(t)dt\\
&\quad + P(t)C'(DD')^{-1}(dy - C\hat{x}dt), \quad \hat{x}(0) = 0,\\
\dot{P}(t) &= AP(t) + P(t)A' + B_1B_1'\\
&\quad - P(t)C'(DD')^{-1}CP(t), \quad P(0) = \Sigma_0.
\end{aligned}$$

Indeed, it is shown in [5] that $\Sigma_T \ge P(T)$ is a *necessary condition* for a terminal state covariance to be “reachable” through steering of the ensemble dynamics. It is also shown that strict inequality is $\Sigma_T > P(T)$ is in fact sufficient.

Similarly, it is a consequence of the optimality of the Kalman filter that if P is the optimal stationary error covariance, which of course satisfies the Algebraic Riccati Equation (ARE)

$$AP + PA' + B_1B_1' - PC'(DD')^{-1}CP = 0,$$

any stationary covariance Σ of the state vector must satisfy $\Sigma \geq P$. It is also shown in [5] that the algebraic condition (8) characterizing stationary state covariances, together with the positivity constraint $\Sigma > P$, are in effect *sufficient* in an approximate sense. More specifically, that in this case it is always possible to maintain the ensemble at a stationary condition with covariance equal or arbitrarily close to Σ.

C. Steering diffusion processes

A considerably more general scenario is being discussed in [3] where ensembles of particles/systems were sought to obey

$$dx(t) = f(x(t), t)dt + \sigma(x(t), t)dw(t), \tag{12}$$

while, at the same time, are allowed to be absorbed or created by the medium in which they travel at some rate [11, p 272]. Their total mass/density $\rho(x, t)$, $x \in \mathbb{R}^n$, evolves according to the transport-diffusion equation

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (f(x,t)\rho) + V(x,t)\rho = \frac{1}{2}\sum_{i,j=1}^{n} \frac{\partial^2(a_{ij}(x,t)\rho)}{\partial x_i \partial x_j},$$

with $\rho(x, 0) = \rho_0(x)$. The presence of $V(x, t)$ allows precisely for the possibility of loss or gain of mass, so that the integral of $\rho(x, t)$ over $\mathbb{R}^n$ is not necessarily constant. The formulation in [3] aims at steering the ensemble through a control drift term $\sigma(x(t), t)u(x(t), t)dt$ to a final target distribution. Necessary conditions are obtained and the special case of linear dynamics is dealt in detail.

D. Zero-noise limit and optimal mass transport

It is of interest to study the evolution of an ensemble that follows a deterministic law, between end-point marginals, and determine an optimal control law. E.g., for linear dynamics, to determine

$$\inf_u \mathbb{E}\int_0^T \frac{1}{2}\|u(t,x)\|^2 dt, \tag{13a}$$

$$\dot{x}(t) = Ax(t) + Bu(t, x) \tag{13b}$$

$$x(0) \sim \rho_0, \quad x(T) \sim \rho_T. \tag{13c}$$

The formulation has been studied as the zero-noise limit of the diffusion process

$$dx(t) = Ax(t)dt + Bu(t)dt + \sqrt{\epsilon}Bdw(t), \tag{14}$$

and connections to optimal mass transport drawn in [12], [13], [7]. This approach allows obtaining approximate solutions to general optimal mass transport problems, via solutions to suitable relaxations that in fact, allow for fast and efficient computation [14], [6].

REFERENCES

[1] Y. Chen, T. T. Georgiou, and M. Pavon, “Optimal steering of a linear stochastic system to a final probability distribution, Part I,” *arXiv:1408.2222; Aut. Control, IEEE Transactions on, DOI: 10.1109/TAC.2015.2457784*, 2015.

[2] ——, “Optimal steering of a linear stochastic system to a final probability distribution, Part II,” *arXiv:1410.3447; Automatic Control, IEEE Transactions on, DOI: 10.1109/TAC.2015.2457791*, 2015.

[3] ——, “Optimal steering of inertial particles diffusing anisotropically with losses,” *arXiv:1410.1605 and Proceedings of American Control Conference, pp. 1252-1257, DOI: 10.1109/ACC.2015.7170905*, 2015.

[4] Y. Chen, T. Georgiou, and M. Pavon, “Optimal control of the state statistics for a linear stochastic system,” *arXiv:1503.04885 and Proceedings of the 54th IEEE Conference on Decision and Control*, 2015.

[5] Y. Chen, T. T. Georgiou, and M. Pavon, “Steering state statistics with output feedback,” *arXiv:1504.00874 and Proceedings of the 54th IEEE Conference on Decision and Control*, 2015.

[6] ——, “Entropic and displacement interpolation: a computational approach using the Hilbert metric,” *arXiv:1506.04255*, 2015.

[7] ——, “On the relation between optimal transport and Schrödinger bridges: A stochastic control viewpoint,” *arXiv:1412.4430v1, J. Optim. Theory Appl*, 2015.

[8] ——, “Optimal transport over a linear dynamical system,” *arXiv:1502.01265v1*, 2015.

[9] R. W. Brockett, “Optimal control of the Liouville equation,” *AMS IP Studies in Advanced Mathematics*, vol. 39, p. 23, 2007.

[10] R. Brockett, “Notes on the control of the liouville equation,” in *Control of Partial Differential Equations*. Springer, 2012, pp. 101–129.

[11] S. Karlin and H. Taylor, “A second course in stochastic processes, 1981.”

[12] C. Léonard, “From the Schrödinger problem to the Monge–Kantorovich problem,” *Journal of Functional Analysis*, vol. 262, no. 4, pp. 1879–1920, 2012.

[13] T. Mikami and M. Thieullen, “Optimal transportation problem by stochastic optimal control,” *SIAM Journal on Control and Optimization*, vol. 47, no. 3, pp. 1127–1139, 2008.

[14] M. Cuturi, “Sinkhorn distances: Lightspeed computation of optimal transport,” in *Advances in Neural Information Processing Systems*, 2013, pp. 2292–2300.