Steering Probability Distributions with Optimal Control

Lucas M. Moschen
Supervised by Greg Pavliotis and Dante Kalise
Department of Mathematics, Imperial College London

- ✓ **Goal:** Analyse the evolution of stochastic systems and **steer the probability distributions** to desired states.
- ✓ **Why it matters**
  - Guide **agent-based** systems in mean-field settings.
  - Speed up **sampling** from simulations.
  - Stabilise or switch **metastable** states.
- ✓ **How we do it**
  - **Spectral** discretisation turns a PDE into an ODE system.
  - Design an **open-loop** time-varying control or a **closed-loop** feedback law.

Background & Motivation

- ✓ **Why mean-field models?** Large systems are best described by their **density** rather than by tracking each individual:
  - Pedestrians in a plaza → crowd flow
  - Ideas in a network → density of viewpoints
  - Molecules in a solvent → probability clouds
- ✓ **Control challenges:**
  - Use an external controller to reshape the mean-field **distribution** as desired.
  - Achieve fast, low-cost influence over collective behaviour.

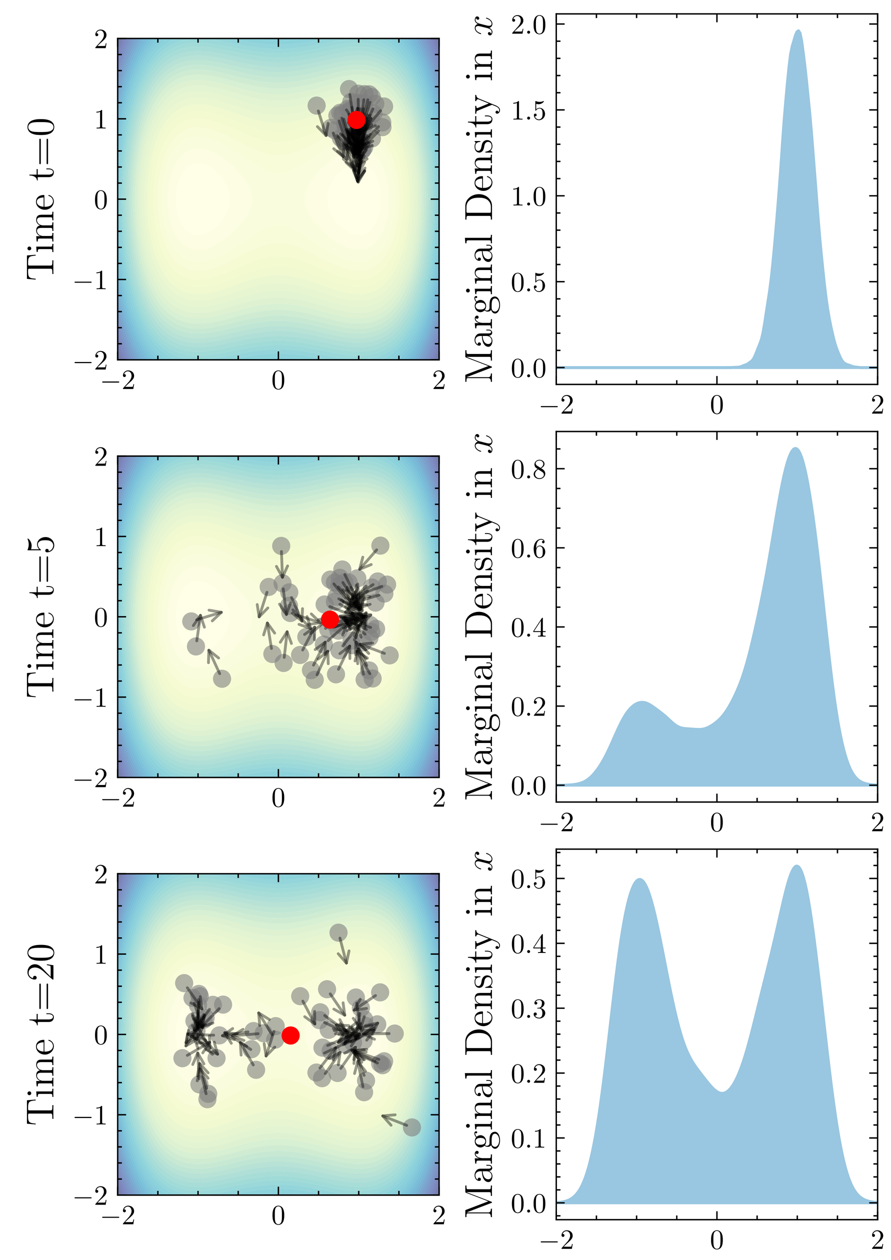

Figure 1: Evolution of particles with potential $V(x,y) = (x^2-1)/4^2 + y^2/2$ at times $t = 0, 5, 20$ (left), and their x-marginal densities (right). The red dot marks the average position.

Mathematical Background

We study overdamped **Langevin** dynamics

$$dX_t = -\left(\nabla V(X_t) + \nabla W * \mu(X_t, t)\right) dt + \sqrt{2\sigma}\, dW_t,$$

where the blue term is the **interaction potential** with the distribution $\mu(\cdot, t)$ of X_t and

$$\nabla W * \mu(x,t) := \int_{\mathbb{R}^d} \nabla W(x-y)\mu(y,t)\, dy.$$

Schematic Overview of Our Control Framework

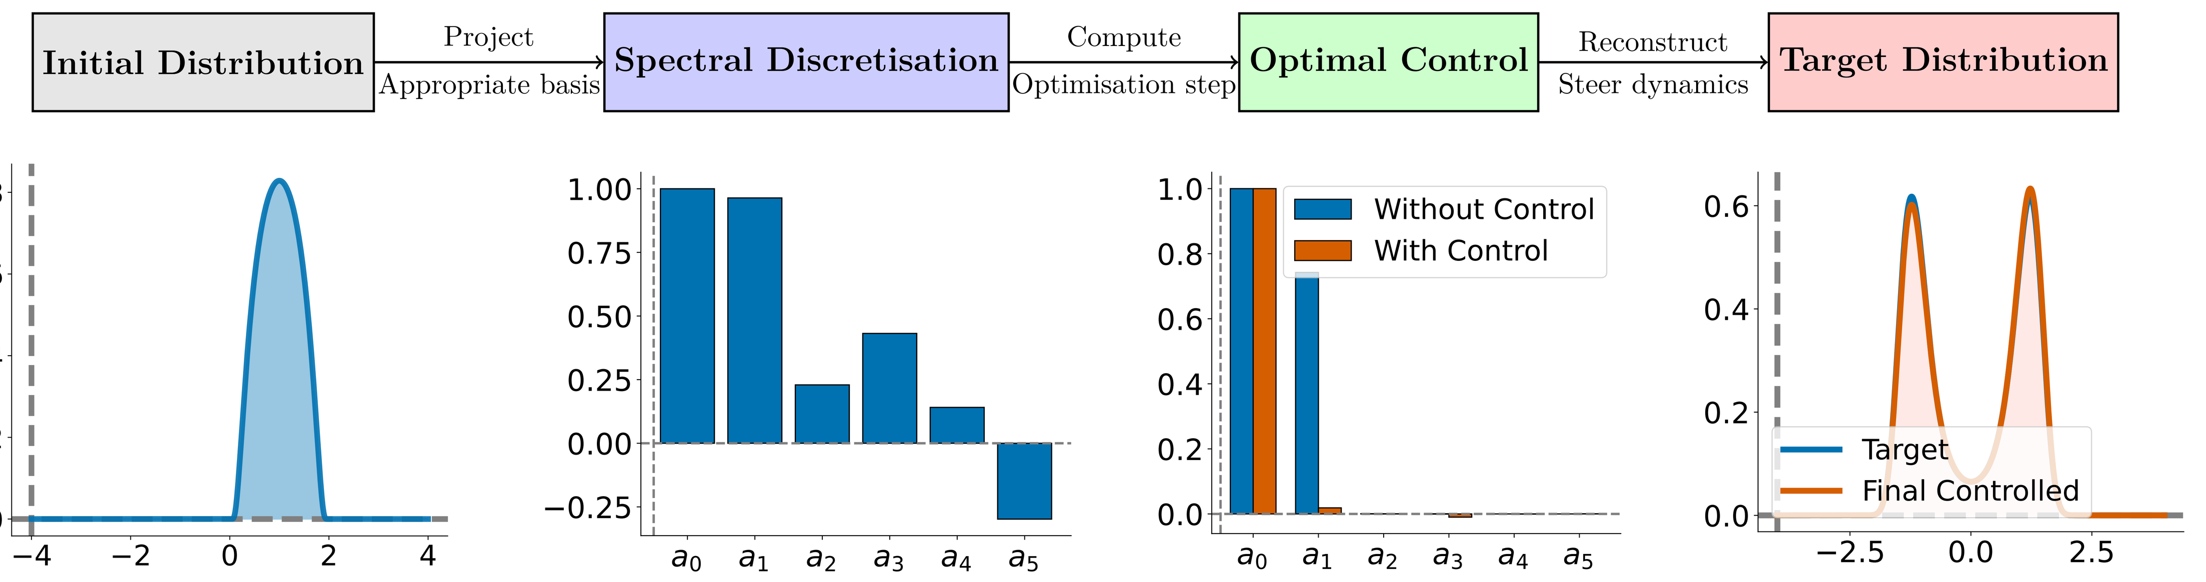

Fokker-Planck and McKean-Vlasov

The associated density $\mu(x,t)$ evolves as

$$\partial_t \mu = \nabla \cdot \Big(\mu \nabla V + \mu \left(\nabla W * \mu\right)\Big) + \sigma \Delta \mu.$$

- ✓ **Non-linear/non-local** term $\mu(\nabla W * \mu)$.
- ✓ **Slow** convergence, especially when multiple equilibria exist.
- ✓ Interaction potential can produce **unstable stationary states**.

Optimal Control Approach

Steer $\mu(\cdot, t)$ toward the target $\mu^\dagger$ by minimising

$$J(u) := \frac{1}{2}\int_0^T \left(\|\mu(\cdot,t) - \mu^\dagger\|_{L^2}^2 + \nu\|u(t)\|^2\right) dt,$$

subject to the PDE constraint

$$\partial_t \mu = \mathcal{A}\mu + \mathcal{W}(\mu) + u\mathcal{N}\mu,$$

where $\mathcal{A}$ includes the diffusion and drift, $\mathcal{W}$ the interaction, and the control $u(t)$ acts via $\mathcal{N}\phi := \nabla \cdot (\phi \nabla \alpha)$ for a chosen function α. In other words, adding $u(t)\,\alpha(x)$ "pushes" V into

$$V(x) + u(t)\,\alpha(x).$$

Method Summary

- ✓ Choose a suitable spectral **basis of functions** (e.g., L^2-orthonormal, periodic).
- ✓ Expand $\mu(x,t)$ in this basis to reduce the PDE to a **finite-dimensional** ODE system.
- ✓ Apply the optimal control via **Pontryagin–** or **Riccati–**based methods on the reduced system to **steer** μ to $\mu^\dagger$.
- ✓ Use the control in the original dynamics.

Numerical Results

- ✓ The **Kuramoto model** describes coupled oscillator dynamics and exhibits **bistability**, making its behaviour highly sensitive to initial conditions.
- ✓ The **ill-conditioned Gaussian** is a diffusion-based model with a small spectral gap, leading to **slow convergence**.

1D Kuramoto Model

Our feedback control efficiently switches the system between two stable equilibria.

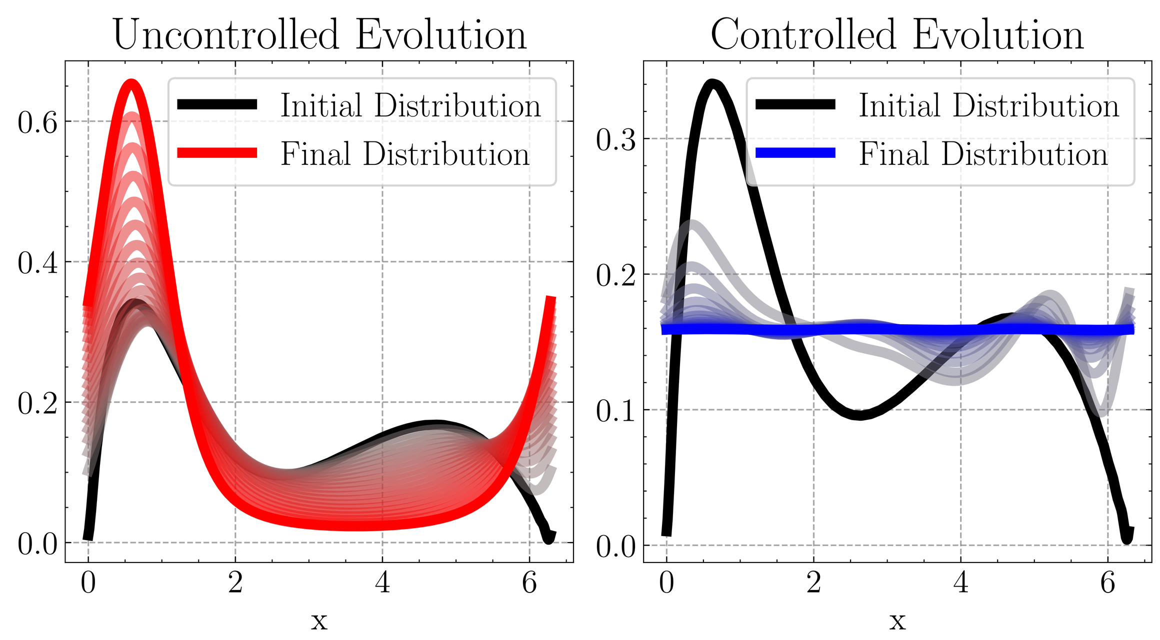

Figure 2: Distribution evolution in the Kuramoto model with $W(x) = -5\cos(x)$ and $G(x) = 0$. Twenty snapshots show the initial state (black) transitioning to red (uncontrolled) and blue (controlled) equilibria.

Ill-conditioned 2D Gaussian

Time to reach equilibrium can be reduced significantly by applying an open-loop control.

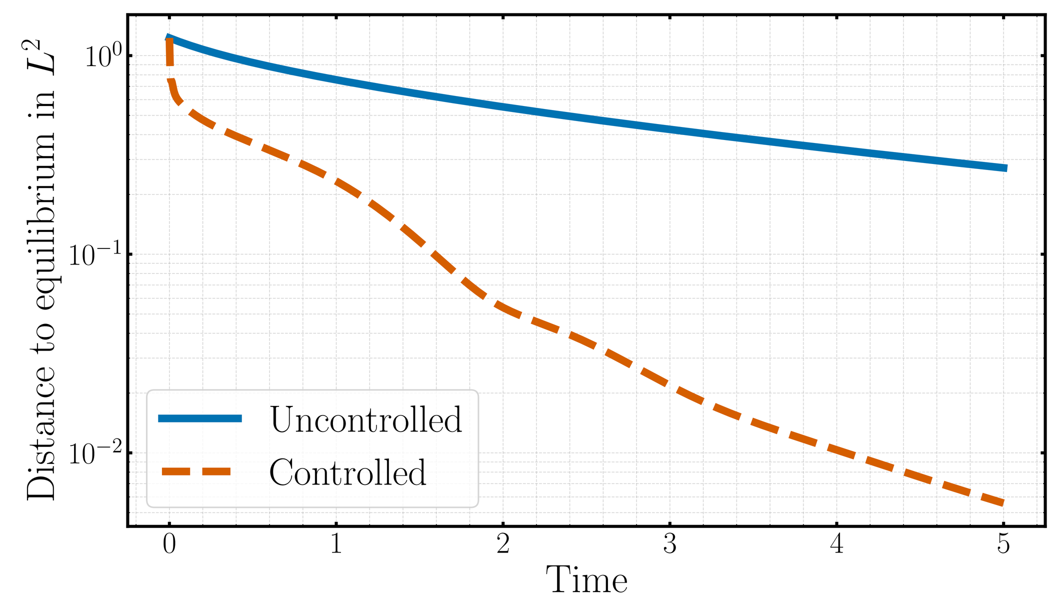

Figure 3: L^2-norm of the difference to the equilibrium with potential $V(x,y) = \frac{1}{2}(x^2 + 0.1y^2)$ and $W(x,y) = 0$. One to four control functions were used with pre-determined α functions.

Conclusions & Take-Home

- ✓ **PDE-based control** can accelerate convergence and shape distributions.
- ✓ Simulations show the method **stabilises metastable** states in challenging regimes.
- ✓ **Future**: solve high-dimensional problems.

Key References

- ✓ Albi, G., Choi, Y.-P., Fornasier, M., Kalise, D. ``Mean-Field Control Hierarchy''. 2017.
- ✓ Breiten, T., Kunisch, K., Pfeiffer, L. ``Control Strategies for the Fokker-Planck Equation''. 2018.
- ✓ Kalise, D., Moschen, L., Pavliotis, G., Vaes, U. ``A Spectral Approach to Optimal Control of the Fokker-Planck Equation''. 2025.

Funders

UK Research and Innovation
cnrs