IMPERIAL

Steering Probability Distributions via PDE-Constrained Optimisation

Lucas M. Moschen

Dante Kalise, Greg A. Pavliotis

Department of Mathematics, Imperial College London

Overview

✓ **Objective**: Steer the evolution of a probability density toward a desired target by applying optimal control to Fokker–Planck PDEs.

✓ **Motivation**:
- Accelerate **mixing** and **sampling**.
- Stabilise or transition between **metastable** states.
- Influence **agent-based** systems in mean-field settings.

✓ **Method**:
- **Spectral** discretisation of the PDE.
- Use **Pontryagin-based** open-loop control and **Riccati-based** feedback.

Background & Motivation

We study overdamped **Langevin** dynamics

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

where the blue term is the **interaction potential** with the distribution $\mu(\cdot, t)$ of X_t. Such dynamics appear in molecular dynamics, Bayesian sampling, and statistical physics.

Fokker–Planck and McKean–Vlasov

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

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

✓ $\nabla W * \mu$ is the **non-linear and non-local** term.

✓ Under appropriate conditions, the dynamics converge to a **unique equilibrium**, but convergence can be slow.

✓ Inclusion of W can produce multiple equilibria, including **unstable stationary states**.

Connections to OT & Learning

✓ FP equations are **gradient flows** in Wasserstein space: links our work to OT.

✓ We design **control laws** that define transport maps from the initial to the desired density.

✓ The final goal is to **accelerate sampling**.

Optimal Control Approach

We aim to steer $\mu(\cdot, t)$ toward the target μ_∞ by minimizing

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

subject to the Fokker–Planck (or McKean–Vlasov) PDE

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

where $u(t)$ is the control and $\mathcal{N}\phi := \nabla \cdot (\phi \nabla \alpha)$ for a given shape function α.

Schematic Overview of Our Control Framework

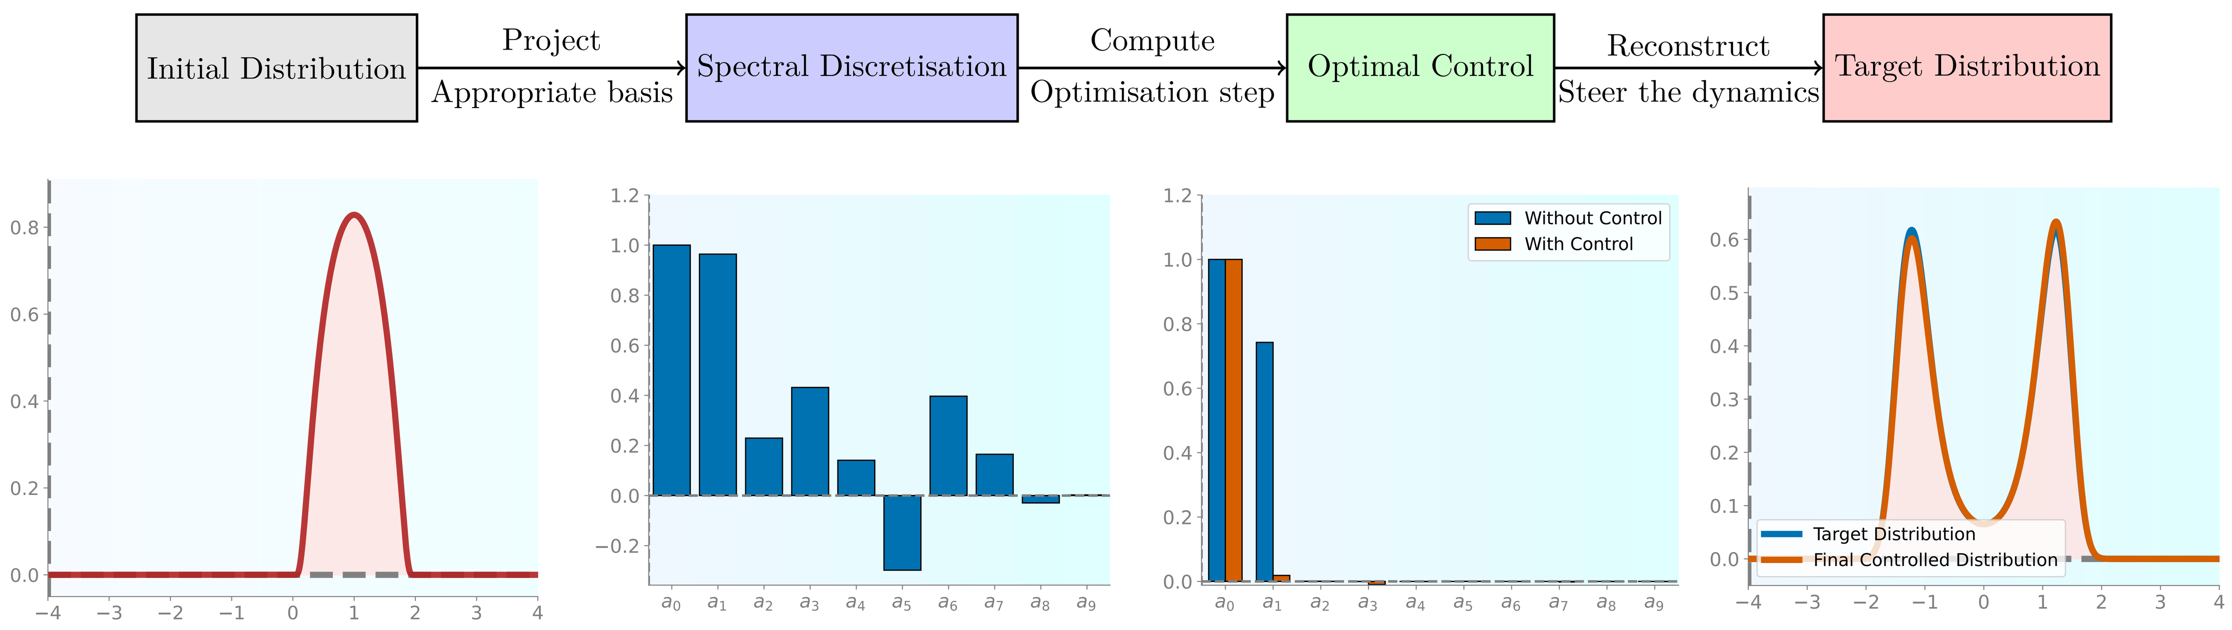

Spectral Discretization

✓ Select an L^2 **basis** that reflects the structure of the problem (e.g., Fourier for periodicity, Hermite for fast decay).

✓ Expand $\mu(x, t)$ in this basis to reduce the PDE to a **finite-dimensional** ODE system.

✓ Apply **Pontryagin-** or **Riccati-based methods** on this reduced system to compute the optimal control law.

Open-Loop Approach (Pontryagin)

✓ Derive optimality conditions via the **Pontryagin Maximum Principle**, yielding a forward PDE for μ and a backward adjoint PDE for ϕ.

✓ The **optimality condition** produces a time-dependent control $u^*(t)$ on $[0, T]$.

✓ Compute $u^*(t)$ using **gradient-based** methods and by solving the combined ODE system for μ and ϕ.

Closed-Loop (Feedback) via a Riccati Equation

✓ **Linearize** the dynamics around μ_∞ by setting $\delta\mu = \mu - \mu_\infty$, and approximate $\mathcal{W}$ and $\mathcal{N}$.

✓ Solve the **Riccati Equation** for the infinite-horizon ($T = \infty$) linear-quadratic problem.

✓ The solution, denoted by Π, yields the **feedback control**

$$u(t) = -\nu^{-1}\mathcal{N}^*\Pi\,\delta\mu,$$

ensuring **robust** convergence.

✓ Apply this control in the original non-linear dynamics.

Results

✓ The **ill-conditioned Gaussian** is a diffusion-based model with a small spectral gap, leading to **slow convergence**.

✓ The **Kuramoto model** describes coupled oscillator dynamics and exhibits **bistability**, making its behaviour highly sensitive to initial conditions.

Ill-conditioned 2D Gaussian

With four control functions, time to reach equilibrium can be reduced significantly.

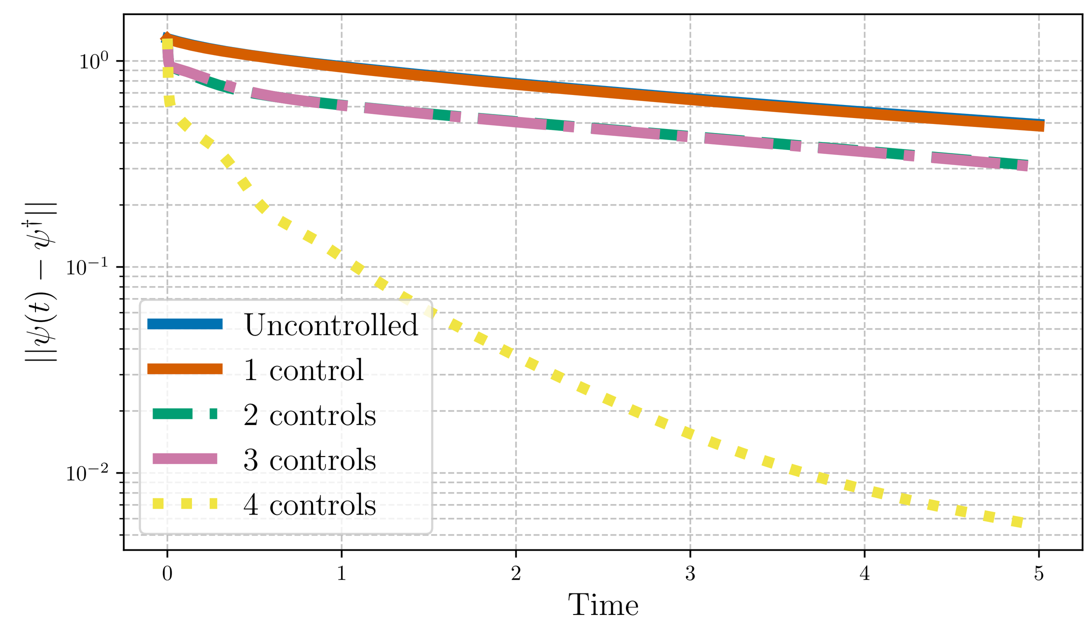

Figure 1: L^2-norm of the difference to the equilibrium with potential $V(x,y) = \frac{1}{2}(x^2 + 0.1y^2)$ and $W(x,y) = 0$. One to four control functions were used with pre-determined α functions.

1D Kuramoto Model

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

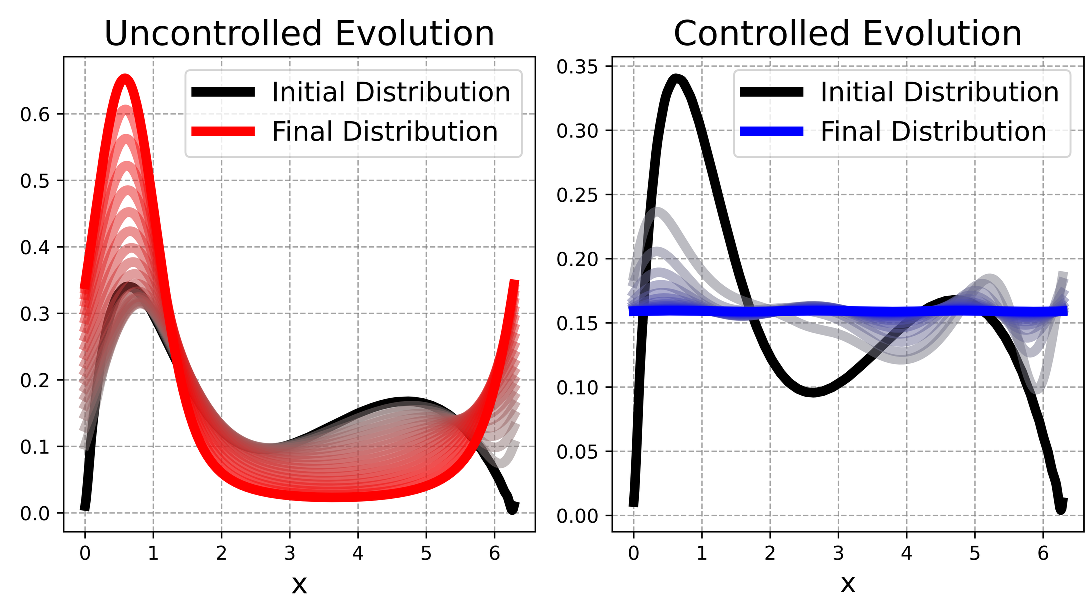

Figure 2: Distribution evolution in the Kuramoto model with $(W(x) = -5\cos(x)$ and $G(x) = 0$. Twenty snapshots show the initial state (black) transitioning to red (uncontrolled) and blue (controlled) equilibria.

Conclusions & Outlook

- **Effective PDE-based control** can increase the spectral gap and accelerate convergence.
- Simulations show that the PDE-based control **stabilizes metastable** states and **accelerates convergence** in challenging regimes.
- **Future**: integration with ML pipelines and OT for sampling applications.

Funders

UK Research and Innovation

cnrs