IMPERIAL

PGMODAYS 2025

Optimal Control and Stabilisation of Fokker-Planck and McKean-Vlasov Equations

Lucas Moschen

Supervised by Greg Pavliotis and Dante Kalise

This project is supported by the CNRS-ICL PhD program.

November 18th, 2025

Controlling the Behaviour of Particle Systems

- Systems with many interacting particles arise in physics, chemistry, and machine learning.
- Their collective dynamics can be described by stochastic differential equations and PDEs.
- We use control theory to steer the system to desired states. **But why?**

How to steer particles?

The long-term behaviour of particle systems can be slow.

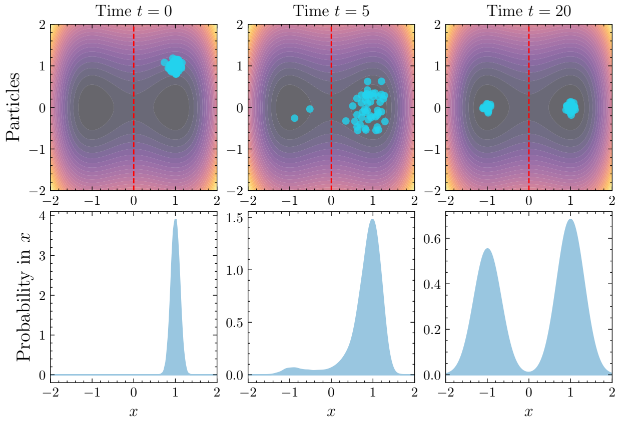

Figure: Particles in the energy landscape $V(x, y) = (x^2 - 1)^2 + y^2$.

- **Stochastic particle systems**: models in molecular dynamics, Bayesian sampling, and collective behaviour.
- Often converges *slowly*, limiting efficiency.
- **Our aim:** develop control strategies to steer distributions towards desired targets, mainly changing the *long-term* behaviour.

From Particles to the Mean-Field Limit and Control

Particles $\xrightarrow{N \to \infty}$ Density $\mu(x, t)$ $\xrightarrow{\text{External control}}$ Controlled dynamics

- N interacting particles moving under noise and forces.
- Approximate the particles by a density $\mu(t, x)$ for large N.
- μ evolves as a **limit PDE (mean-field limit)**

$$\partial_t \mu = \underbrace{\sigma \Delta \mu}_{\text{Diffusion}} + \underbrace{\nabla \cdot (\mu \nabla V)}_{\text{Energy drift}} + \underbrace{\nabla \cdot (\mu \nabla W * \mu)}_{\text{Interaction drift}}.$$

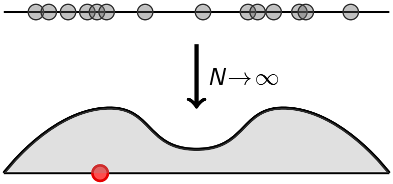

We use optimal control and LQR feedback to accelerate and stabilise mean-field particle systems at the PDE level.

Spectral Optimal Control of Fokker-Planck[1]

- **Overdamped Langevin dynamics:**

$$dX_t = -\nabla V(X_t)\, dt + \sqrt{2\sigma}\, dW_t, \quad X_t \in \mathbb{R}^d,$$

with $\mu(t,\cdot) = \mathrm{Law}(X_t)$ satisfying the Fokker-Planck equation

$$\partial_t \mu = \mathcal{L}^* \mu := \nabla \cdot (\mu \nabla V) + \sigma \Delta \mu.$$

- **Stationary density:** $\mu_\infty \propto e^{-V/\sigma}$.

- Under suitable conditions on V,

$$\|\mu(t,\cdot) - \mu_\infty\|_{L^2(\mathbb{R}^d)} \le \|\mu(0,\cdot) - \mu_\infty\|_{L^2(\mathbb{R}^d)} e^{-\lambda t},$$

where the rate $\lambda > 0$ is the **spectral gap** of the Fokker-Planck operator $\mathcal{L}^*$. Even in this case, **convergence may be slow**.

[1] *D. Kalise, L.M., G. Pavliotis, U. Vaes (2025). A Spectral Approach to Optimal Control of the Fokker-Planck Equation*

Working in the Correct Space
From the Fokker-Planck Operator to a Schrödinger Operator

- The Fokker-Planck operator $\mathcal{L}^*$ is not self-adjoint in $L^2(\mathbb{R}^d)$, but it is in

$$L^2(\mathbb{R}^d, \mu_\infty^{-1}) := \{f : f/\sqrt{\mu_\infty} \in L^2(\mathbb{R}^d)\}.$$

- Define the unitary map $\mathcal{U} : L^2(\mathbb{R}^d, \mu_\infty^{-1}) \to L^2(\mathbb{R}^d)$, such that $\mathcal{U}[\varphi](x) = \dfrac{\varphi(x)}{\sqrt{\mu_\infty(x)}}$.

- Ground-state transform: $\mathcal{H} := -\mathcal{U}\mathcal{L}^*\mathcal{U}^{-1} = -\sigma\Delta + W(x)$, with

$$W(x) := \frac{1}{4\sigma}|\nabla V(x)|^2 - \frac{1}{2}\Delta V(x).$$

- If $\psi = \mathcal{U}\mu$: $\partial_t \psi = -\mathcal{H}\psi$ (imaginary-time Schrödinger form).

Control Formulation

- Modify the potential as[2]

$$V(x) \mapsto V(x) + \overbrace{\sum_{j=1}^{m} u_j(t)\, \alpha_j(x)}^{\text{Control input}},$$

  for **chosen** spatial shape controls α_j and **to be optimized** time-dependent controls u_j.

- In Schrödinger variables:

$$\partial_t \psi = -\Big(\mathcal{H} - \sum_{j=1}^{m} u_j(t)\mathcal{N}_j\Big)\psi,$$

  where $\mathcal{N}_j[\varphi] = \nabla \cdot (\varphi \nabla \alpha_j) - \nabla \alpha_j \cdot \nabla V \varphi / 2\sigma$.

- α_j **target slow modes** to enlarge the spectral gap.

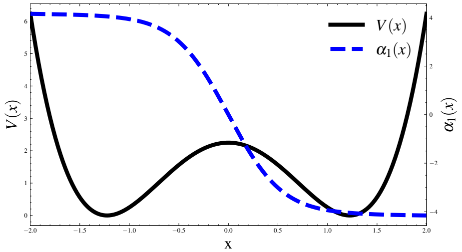

Schematic of α_1 targeting the wells.

[2] *T. Breiten, K. Kunisch, and L. Pfeiffer (2018). Control strategies for the Fokker-Planck equation.*

Open-loop Optimal Control

Pontryagin Maximum Principle + Barzilai-Borwein method

- Finite-horizon tracking:

$$J(u) = \frac{\kappa}{2}\int_0^T \|\psi - \psi_\infty\|^2_{L^2(\mathbb{R}^d)} dt + \frac{\nu}{2}\sum_{i=1}^m \int_0^T u_i(t)^2 dt + \frac{1}{2}\|\psi(T) - \psi_\infty\|^2_{L^2(\mathbb{R}^d)}.$$

- Adjoint system (first-order conditions):

$$\partial_t \phi = \mathcal{H}\phi - \sum_{j=1}^m u_j \mathcal{N}_j^* \phi + \kappa(\psi - \psi_\infty), \quad \phi(T,\cdot) = \psi_\infty - \psi(T,\cdot).$$

- Gradient and optimal control:

$$\nabla_{u_j} J = \nu u_j - \langle \mathcal{N}_j^* \phi, \psi \rangle_{L^2(\mathbb{R}^d)} \quad \Rightarrow \quad u_j(t) = \frac{1}{\nu}\langle \phi, \mathcal{N}_j \psi \rangle_{L^2(\mathbb{R}^d)}.$$

- Shape functions α_j are chosen via $\nabla \cdot (\mu_\infty \nabla \alpha_j) = \sqrt{\mu_\infty} e_j$.

Spectral Galerkin on $\mathbb{R}^d$ and Algorithm

- The operator $\mathcal{H}$ is **self-adjoint**, has **compact resolvent**, and **completeness of the eigenfunctions**. Let $\mathcal{H}e_k = \lambda_k e_k$ with $\lambda_k \to \infty$.

- **Algorithm 1 (sketch):**
  1. Initialize $u^{(0)}(t)$ (optionally from a Riccati feedback on the linearized system, which is cheap).
  2. Forward solve for $\psi^{(k)}$, the state equation.
  3. Backward solve for $\phi^{(k)}$, the adjoint equation.
  4. Update $u^{(k+1)} = u^{(k)} - \gamma_k \nabla_u J$, γ_k computed through Barzilai-Borwein step.
  5. Stop when the norm of the gradient is smaller than a tolerance.

- The Barzilai-Borwein step acts as a quasi-Newton update, significantly accelerating convergence in our gradient iterations.

Numerical Experiments

- $\mathbb{R}^2$ tests:
  - Ill-conditioned Gaussian $V(x, y) = ax^2 + by^2$ with $b \ll a$ (spectral gap b).
  - Double-well potential with four minima.
- Controls **target slow modes** $\Rightarrow$ faster decay of modal coefficients and effective spectral-gap enlargement.
- **Bottom line:** spectral open-loop control on $\mathbb{R}^d$ is efficient and extends[3].
- **Issue:** depends on solving an eigenvalue problem $\Rightarrow$ complicated in high dimensions.

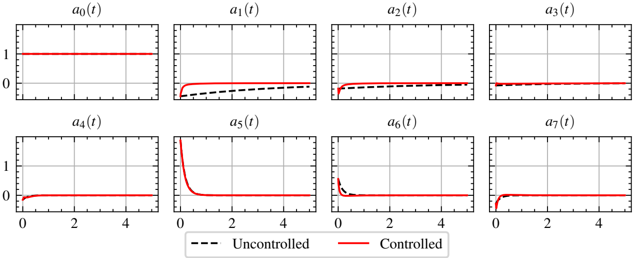

Time evolution of the coefficients of the spectral decomposition of the double-well potential.

[3] *M. Annunziato, A. Borzì (2018). A Fokker-Planck control framework for stochastic systems*

The McKean-Vlasov Equation[5][6]

- **What if the particles interact?**[4]

$$\partial_t \mu = \underbrace{\sigma \Delta \mu}_{\text{Diffusion}} + \underbrace{\nabla \cdot (\mu \nabla V)}_{\text{External drift}} + \underbrace{\nabla \cdot (\mu \nabla W * \mu)}_{\text{Interaction drift}}$$

- $\mu(t, x)$: evolution of the density of agents.
- $W(x)$: introduces **nonlinearity** and **nonlocality** through $(\nabla W * y)(x) := \int_\Omega \nabla W(x - x')\, y(x')\, dx'$.
- Fokker-Planck PDE for the McKean-Vlasov SDE where the drift depends on the law μ.

[4] *D. Kalise, L.M., G. Pavliotis (2025). Linearization-Based Feedback Stabilization of McKean-Vlasov PDEs*
[5] *A.-S. Sznitman (1991). Topics in propagation of chaos*
[6] *J.A. Carrillo, R.J. McCann, and C. Villani (2003). Kinetic equilibration rates for granular media and related equations*

Why Control Interacting Systems?
Motivations for Feedback Design

- Uncontrolled systems:
  - Possible **slow convergence** to the long-term (small spectral gap).
  - Converge to an **undesirable equilibrium**.
- Our goals:
  - **Accelerate convergence** to a desired steady state.
  - **Steer the system** toward or away from specific modes.
- Solution approach:
  - Introduce **time-dependent feedback potentials** into the PDE dynamics.[7]
  - Linearisation + Spectral analysis + Linear-Quadratic control + Nonlinear Estimates $\Rightarrow$ Stabilisation

[7] *T. Breiten, K. Kunisch, and L. Pfeiffer (2018). Control strategies for the Fokker-Planck equation.*

Noisy Kuramoto Model
What kind of problem?

- Coupled phase oscillators with sine interaction that synchronize.[8]
- Let $\Omega = [0, 2\pi)$, $V(x) = 0$, $W(x) = -K\cos(x)$
- If $K \leq 1$: $\bar{\mu}(x) = \frac{1}{2\pi}$ is the unique steady state.
- If $K > 1$: $\bar{\mu} = \frac{1}{2\pi}$ becomes unstable and infinite new steady states appear.

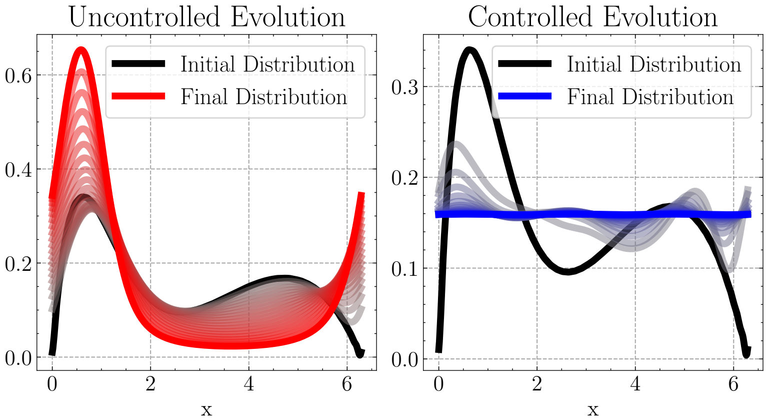

[8] *L. Bertini, G. Giacomin, K. Pakdaman (2010). Dynamical Aspects of Mean Field Plane Rotators and the Kuramoto Model*

Numerical Implementation
Noisy Kuramoto under Feedback Control

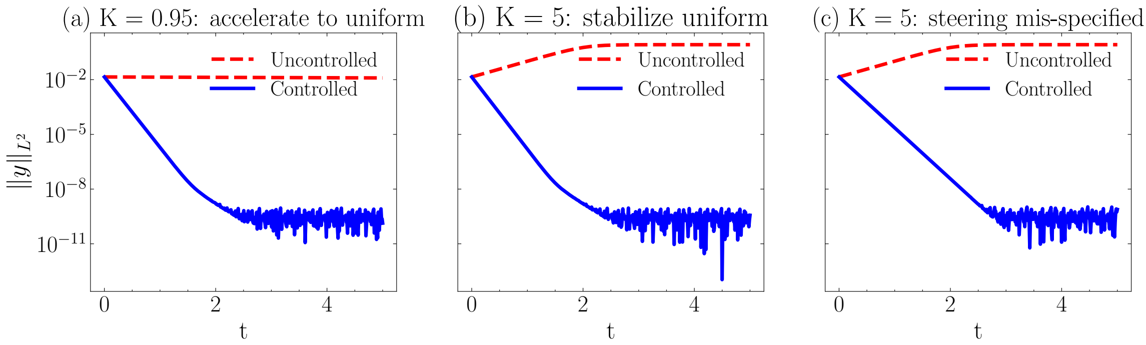

Figure: Set $V(x) = 0$, $W(x) = -K\cos(x)$ and $\sigma = 0.5$.

Long-Time Behaviour of the McKean-Vlasov Equation

- **Steady states:** time-independent solutions $\bar{\mu}$ satisfying

$$\nabla \cdot (\sigma \nabla \bar{\mu} + \bar{\mu}(\nabla V + \nabla W * \bar{\mu})) = 0 \quad \Rightarrow \quad \bar{\mu} \propto \exp\left\{-\frac{1}{\sigma}(V + W * \bar{\mu})\right\}.$$

- **Gradient flow:** the dynamics decrease a free energy functional. Steady states are critical points of this energy.

- **Convex potentials:** convexity of V and W implies uniqueness of $\bar{\mu}$ and $\mu(t, \cdot) \xrightarrow{L_1} \bar{\mu}$ exponentially.[9]

- **Nonconvex potentials:** multiple steady states may emerge.

[9] *F. Malrieu (2001). Logarithmic Sobolev inequalities for some nonlinear PDEs*

Operator-Theoretic Reformulation and Linearisation

- **Perturbation variable:** let $y = \mu - \bar{\mu}$. Then

$$\dot{y} = \overbrace{(\mathcal{A} + \mathcal{D}_W)y}^{\text{linear part}} + \sum_{j=1}^{m} \overbrace{u_j(t)\,\mathcal{N}_j y}^{\text{bilinear control}} + \sum_{j=1}^{m} \overbrace{\mathcal{B}_j\, u_j(t)}^{\text{pure input term}} + \overbrace{\mathcal{W}(y),}^{\text{nonlinear remainder}} \qquad \mathcal{B}_j = \mathcal{N}_j \bar{\mu}.$$

- **Operators:** $\mathcal{A}\mu = \nabla \cdot (\sigma \nabla \mu + \mu \nabla V)$, $\mathcal{N}_j \mu = \nabla \cdot (\mu \nabla \alpha_j)$, $\mathcal{W}(\mu) = \nabla \cdot (\mu \nabla W * \mu)$, $\mathcal{D}_W$ is the Fréchet derivative of $\mathcal{W}$ at $\bar{\mu}$.

- **State space:** $\mathcal{X}_0 = \left\{ y : \int_\Omega |y(x)|^2\, \bar{\mu}^{-1}(x)\, dx < \infty \text{ and } \int_\Omega y(x)\, dx = 0 \right\}$.

- **Zero-mean subspace:** $\mathcal{A} + \mathcal{D}_W$ preserves mass, so 0 is an eigenvalue we do not need to control.

How do we plan to control?

- **Goal:** stabilise an equilibrium $\bar{\mu}$ of the McKean-Vlasov equation via **feedback control**.

- **Linearised dynamics around** $\bar{\mu}$:

$$\dot{y} = (\mathcal{A} + \mathcal{D}_W)y + \cancel{\textstyle\sum_{j=1}^{m} u_j(t)\,\mathcal{N}_j y} + \sum_{j=1}^{m} \mathcal{B}_j\, u_j(t) + \cancel{\mathcal{W}(y)}, \quad , \quad \mathcal{B}_j = \nabla \cdot (\bar{\mu} \nabla \alpha_j).$$

- **Key idea:** fix $\delta > 0$. Control the **directions** related to eigenvalues with real part $\geq -\delta$.

- **Strategy:**

Spectral analysis of $\mathcal{A} + \mathcal{D}_W$ → Identify unstable modes ($\Re\lambda \geq -\delta$) → Use LQR to build a finite-dim feedback → Closed-loop stability Contraction argument

Spectral Reformulation via Ground-State Transform

- **Ground-state transform:** define the unitary map

$$\mathcal{U} : L^2(\Omega, \bar{\mu}^{-1}) \to L^2(\Omega), \qquad \mathcal{U}y = \frac{y}{\sqrt{\bar{\mu}}}.$$

- **Transformed operator:** the linearised operator becomes

$$\mathcal{H} := -\mathcal{U}(\mathcal{A} + \mathcal{D}_W)\mathcal{U}^{-1} = -\sigma\Delta + \Psi(x) + \mathcal{K},$$

  where $-\sigma\Delta + \Psi$ is a Schrödinger operator and $\mathcal{K}$ is Hilbert-Schmidt.

- **Spectral structure:** $\mathcal{H}$ has **compact resolvent** and a pure point spectrum[10], with eigenvalues accumulating at $+\infty$ ($\sigma(\mathcal{H}) = -\sigma(\mathcal{A} + \mathcal{D}_W)$) and only **finitely many** eigenvalues with $\Re\lambda \leq \delta$. These correspond to the slow/unstable modes we wish to control.

[10] *M. Reed and B. Simon (1978). Methods of Modern Mathematical Physics IV: Analysis of Operators*

Linear-Quadratic Control of the Linearised Problem

- **Infinite-horizon LQ problem:** let $\mathcal{L} = \mathcal{A} + \mathcal{D}_W$ and $\mathcal{B} = [\mathcal{B}_1, \ldots, \mathcal{B}_m]$. We solve

$$\min_{u(\cdot)} \frac{1}{2} \int_0^\infty e^{2\delta t} \Big(\langle y(t), \mathcal{M} y(t) \rangle + |u(t)|^2 \Big) \, dt \quad \text{s.t.} \quad \dot{y}(t) = \mathcal{L} y(t) + \mathcal{B} u(t).$$

- **Riccati equation and feedback:** find Π self-adjoint solving

$$(\mathcal{L}^* + \delta I)\Pi + \Pi(\mathcal{L} + \delta I) - \Pi \mathcal{B} \mathcal{B}^* \Pi + \mathcal{M} = 0$$

  and $u(t) = -\mathcal{B}^* \Pi y(t)$ yields exponential decay with rate δ (infinite-dimensional Hautus condition[11]).
- **Choosing the actuators $\mathcal{B}_j$:** expand

$$y(t,x) = \sum_{n \geq 1} c_n(t) \, \varphi_n(x), \qquad \mathcal{L} \varphi_n = \lambda_n \varphi_n,$$

  and choose $\mathcal{B}_j$ so that they are aligned with the eigenfunctions corresponding to the m eigenvalues with $\Re \lambda_n \geq -\delta$, i.e. $\mathcal{B}_j = \nabla \cdot (\bar{\mu} \nabla \alpha_j) \approx \varphi_j$.

[11] *R. Curtain and H. Zwart (2005). An Introduction to Infinite-Dimensional Linear Systems Theory*

Local Exponential Stabilisation: Closed-Loop View

The equation for $\psi = e^{\delta t}\mathcal{U}y$ under the feedback $u = -\mathcal{B}^*\Pi y$

$$\dot{\psi} = -\mathcal{H}_\Pi\psi - \big(\hat{\mathcal{B}}^*\hat{\Pi}\psi\big)\hat{\mathcal{N}}^\delta(t)\psi + \hat{\mathcal{W}}^\delta(\psi), \quad \psi(0,\cdot) = \mathcal{U}(\mu_0 - \bar{\mu}),$$

where $\mathcal{H}_\Pi = \mathcal{H} - \delta I + \hat{\mathcal{B}}\hat{\mathcal{B}}^*\hat{\Pi}$.

- **Linear decay:** $\mathcal{H}_\Pi$ generates an exponentially stable semigroup with rate δ.
- **Nonlinear remainder:** the term $\hat{\mathcal{W}}^\delta(\psi)$ admits a Lipschitz-type estimate.
- **Contraction argument:** for $\|\psi_0\|$ sufficiently small, the full closed-loop system defines a **contraction mapping** in a suitable function space[12]. Therefore,

$$\|\psi(t)\|_{L^2(\Omega)} \leq C\|\psi_0\|_{L^2(\Omega)} \Rightarrow \|y(t)\|_{L^2(\Omega,\bar{\mu}^{-1})} \leq Ce^{-\delta t}\|\mu_0 - \bar{\mu}\|_{L^2(\Omega,\bar{\mu}^{-1})}.$$

[12] *D. Kalise, L. M, G. Pavliotis (2025). Linearization-Based Feedback Stabilization of McKean-Vlasov PDEs*

Kuramoto model + Potential V

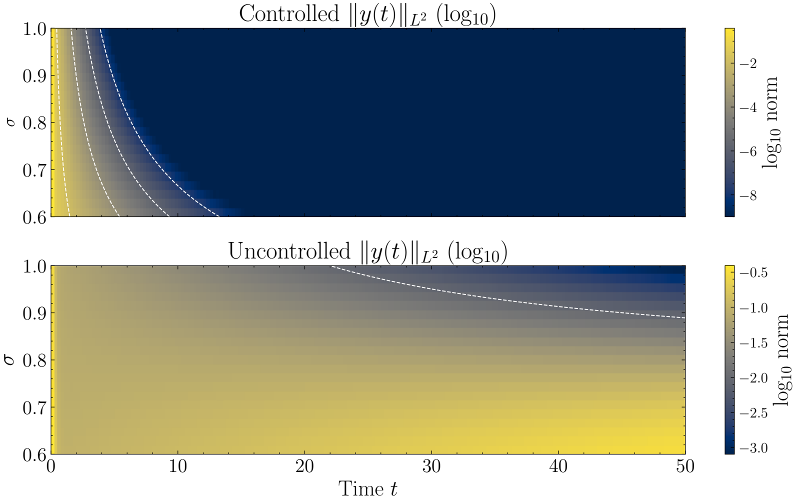

Figure: Set $V(x) = 0.05\cos(x)$ and $W(x) = -\cos(x)$.

Conclusions

Wrap-up and Perspectives

- Developed a **spectral framework for open-loop optimal control** of Fokker-Planck equations
- Established a **feedback stabilisation theory** for McKean-Vlasov PDEs on the torus.
- Demonstrated **numerical robustness and acceleration** on several examples.
- **Ongoing work:**
  - How to use this control in the particle level, the SDEs? Can we adapt for high-dimensional sampling problems?
  - Results in unbounded domain $\mathbb{R}^d$.
  - Control of gradient flows in probability spaces.

IMPERIAL

Thank you for your attention!