


A Spectral Approach to Optimal Control of the Fokker-Planck Equation

Dante Kalise¹, Lucas M. Moschen¹, Grigorios A. Pavliotis¹, and Urbain Vaes²

¹Imperial College London (UK) ²Inria (France)

Sampling Viewpoint

Why control Fokker-Planck equations?



Particles in the energy landscape

$$V(x, y) = (x^2 - 1)^2 + y^2.$$

Target (Gibbs) law: $\pi(dx) \propto e^{-V(x)/\sigma}$

- **Overdamped Langevin dynamics**

$$dX_t = -\nabla V(X_t) dt + \sqrt{2\sigma} dB_t, \quad \mu_t = \text{Law}(X_t).$$

- If π satisfies a Poincaré inequality, then

$$\|\mu_t - \pi\|_{L^2(\pi^{-1})} \leq e^{-\lambda t} \|\mu_0 - \pi\|_{L^2(\pi^{-1})}.$$

For nonconvex V , λ can be small $\Rightarrow$ **slow convergence**.

- **Goal today:** design controls acting on V to “increase λ ” and accelerate convergence¹.
- **Applications:** molecular dynamics, optics, sampling algorithms via Langevin dynamics.

¹T. Lelièvre, F. Nier, and G. Pavliotis (2013). “Optimal non-reversible linear drift for the convergence to equilibrium of a diffusion.”

Fokker-Planck Equation and Spectral Gap

From sampling dynamics to an operator viewpoint

- Recall the overdamped Langevin dynamics from the previous slide

$$dX_t = -\nabla V(X_t) dt + \sqrt{2\sigma} dB_t, \quad \mu_t = \text{Law}(X_t).$$

- The law μ_t solves the **Fokker-Planck equation**

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

where $\mathcal{L}^*$ is the adjoint of the Langevin generator $\mathcal{L}$.


- Then $\pi(dx) \propto e^{-V(x)/\sigma} dx$ is the stationary solution with $\|\mu_t - \pi\|_{L^2(\pi^{-1})} \leq e^{-\lambda t} \|\mu_0 - \pi\|_{L^2(\pi^{-1})}$. The parameter λ is the **spectral gap** of $-\mathcal{L}$ in $L^2(\mathbb{R}^d, \pi)$ and “ $\mu_\infty = \pi$.”
- Assumptions:** $V \in C^2(\mathbb{R}^d)$, $\sigma > 0$, $e^{-V/\sigma}$ is integrable in $\mathbb{R}^d$ and

$$W(x) := \frac{1}{4\sigma} |\nabla V(x)|^2 - \frac{1}{2} \Delta V(x), \quad \lim_{|x| \rightarrow \infty} W(x) = \infty,$$

so that V grows fast enough outside compact sets.

Confining but Nonconvex Potentials

Several local minima, quadratic growth at infinity



- V can have **several local minima** inside a bounded region, but grows sufficiently fast as $|x| \rightarrow \infty$.
- In such landscapes, the spectral gap λ can be small: particles spend a long time trapped in wells.

Ground-State Transform: Working in a Self-Adjoint Basis

From Fokker-Planck to Schrödinger

- $\mathcal{L}^*$ 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 ground-state transform

$$\mathcal{U}: L^2(\mathbb{R}^d, \mu_\infty^{-1}) \rightarrow L^2(\mathbb{R}^d), \quad \mathcal{U}[\varphi](x) = \frac{\varphi(x)}{\sqrt{\mu_\infty(x)}}.$$

- The transformed operator

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

is a Schrödinger operator with potential W .

- $\mathcal{H}$ is **self-adjoint** with **compact resolvent**²: orthonormal basis of eigenfunctions $(e_k)_{k \geq 1}$ in $L^2(\mathbb{R}^d)$ and eigenvalues $\lambda_k \rightarrow +\infty$.
- The spectral gap of $\mathcal{L}^*$ is $\lambda_1 > 0$.

²M. Reed and B. Simon (1978). "IV: Analysis of Operators."

Control by Shaping the Potential

Targeting slow modes of the Schrödinger operator

- We modify the potential as³


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

where $\alpha_j(x)$ are **shape functions** in space and $u_j(t)$ are time-dependent controls (to be optimised).

- In Schrödinger variables $\psi = \mathcal{U}\mu$, the dynamics become

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

with $\mathcal{N}_j = \mathcal{U} \mathcal{A}_j^* \mathcal{U}^{-1}$ and $\mathcal{A}_j^* \phi = \nabla \cdot (\phi \nabla \alpha_j)$.



Schematic control shaping a double-well potential.

We design the α_j so that $\mathcal{N}_j$ strongly acts on the **slow eigenmodes** of $\mathcal{H}$.

³T. Breiten, K. Kunisch, and L. Pfeiffer (2018). "Control strategies for the Fokker-Planck equation."

Optimal Control Formulation

Finite-horizon tracking in Schrödinger variables

- **Goal:** drive $\psi(t)$ quickly towards $\psi_\infty = \mathcal{U}\mu_\infty$.
- We consider a **finite-horizon quadratic tracking problem** in Schrödinger variables:

$$J(u) = \frac{\kappa}{2} \int_0^T \|\psi(t) - \psi_\infty\|_{L^2}^2 dt + \frac{\nu}{2} \sum_{j=1}^m \int_0^T u_j(t)^2 dt + \frac{1}{2} \|\psi(T) - \psi_\infty\|_{L^2}^2,$$

subject to

$$\partial_t \psi_t = -\left(\mathcal{H} - \sum_{j=1}^m u_j(t) \mathcal{N}_j\right) \psi_t, \quad \psi_0 = \mathcal{U}\mu_0.$$

- Pontryagin's maximum principle⁴ gives an **adjoint equation** for ϕ and a **gradient** with respect to u . In particular,

$$u_j^*(t) = \frac{1}{\nu} \langle \phi(t), \mathcal{N}_j \psi(t) \rangle_{L^2},$$

where ϕ solves the adjoint equation backward in time.

⁴M. S. Aronna and F. Troeltzsch (2021). "First and second order optimality conditions for the control of Fokker-Planck equations."

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

Adjoint-based gradient and optimisation

- The operator $\mathcal{H}$ is **self-adjoint**, has **compact resolvent**, and has a **complete orthonormal** set of the eigenfunctions. Let $\mathcal{H}e_k = \lambda_k e_k$ with $\lambda_k \rightarrow \infty$.
- We use a **spectral Galerkin discretisation**⁵ in this basis:
 1. Initialize $u^{(0)}(t)$ (optionally from a Riccati feedback on the linearised system, which is cheap).
 2. Forward solve for $\psi^{(k)}$ (state equation).
 3. Backward solve for $\phi^{(k)}$ (adjoint equation).
 4. Update $u^{(k+1)} = u^{(k)} - \gamma_k \nabla_u J$, γ_k computed through Barzilai-Borwein step.
 5. Stop when $\|\nabla_u J\| < \text{tol}$.
- The Barzilai-Borwein step⁶ (with safeguard line search) acts as a quasi-Newton update, significantly accelerating convergence in our gradient iterations.


⁵M. Mohammadi and A. Borzi (2015). "A Hermite spectral method for a Fokker-Planck optimal control problem in an unbounded domain."

⁶B. Azmi and K. Kunisch (2020). "Analysis of the Barzilai-Borwein step-sizes for problems in Hilbert spaces"

Numerical Results on $\mathbb{R}^2$

Enlarging the spectral gap in practice

- Test cases in dimension $d = 2$:
 - ill-conditioned Gaussian
 $V(x, y) = ax^2 + by^2$ with $b \ll a$,
 - double-well potential with four minima.
- Controls are built to act on slow modes of $\mathcal{H}$.
- We observe:
 - faster decay of low-order spectral coefficients,
 - enlargement of the spectral gap,
 - robust performance over different initial conditions and misplacement of α_j



Time evolution of spectral coefficients in a double-well potential.

Remark: relies on eigenvalue computations which become challenging in very high dimension.

Recent Work: Feedback Control of Interacting Particle Systems

McKean-Vlasov PDEs

- Interacting particle systems lead to the **McKean-Vlasov equation**

$$\partial_t \mu_t = \sigma \Delta \mu_t + \nabla \cdot (\mu_t \nabla V) + \nabla \cdot (\mu_t \nabla W * \mu_t)$$


with interaction potential W .

- Multiple equilibria!* The linearised operator around a steady state $\bar{\mu}$ has discrete spectrum.
- By solving a Riccati equation for Π , we obtain a **feedback law**⁷

$$u(t) = -\mathcal{B}^* \Pi(\mu_t - \bar{\mu}),$$

$$\text{leading to } \|\mu_t - \bar{\mu}\|_{L^2(\bar{\mu}^{-1})} \lesssim e^{-\delta t} \|\mu_0 - \bar{\mu}\|_{L^2(\bar{\mu}^{-1})}.$$

⁷D. Kalise, L. M., G. Pavliotis (2025). "Linearization-Based Feedback Stabilization of McKean-Vlasov PDEs."



Noisy Kuramoto model on $[0, 2\pi)$: $V(x) = 0$,
 $W(x) = -K \cos x$, $\sigma > 0$.

Conclusions and Outlook

Take-away messages

- We transform a Fokker-Planck control into a **Schrödinger spectral shaping** problem.
- An adjoint-based algorithm in the **spectral Galerkin basis** using the eigenfunctions enlarges the effective spectral gap and accelerates convergence.
- **Feedback control** is a different approach that leads to δ -stabilisation around any steady state.

Ongoing work: sampling and molecular dynamics

- High-dimensional **molecular dynamics sampling**: build effective dynamics along reaction coordinates and design controls there, then lift them back to the full Langevin / particle system.
- **Gradient-flow connection**: view Fokker-Planck / McKean-Vlasov as gradient flows in the **Wasserstein probability space** to analyse how the feedback law changes the geometry.

IMPERIAL

Thank you for your attention!

A Spectral Approach to Optimal Control of the Fokker-Planck Equation
December 10th, 2025

Optimality System

- State equation (Schrödinger variables):

$$\partial_t \psi = -\left(\mathcal{H} - \sum_{j=1}^m u_j(t) \mathcal{N}_j\right) \psi, \quad \psi(0) = \psi_0.$$

- Adjoint equation:

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


- Gradient of the cost functional:

$$\frac{\delta J}{\delta u_j}(t) = \nu u_j(t) - \langle \phi(t), \mathcal{N}_j \psi(t) \rangle_{L^2}, \quad j = 1, \dots, m.$$

- Optimal controls satisfy

$$u_j(t) = \frac{1}{\nu} \langle \phi(t), \mathcal{N}_j \psi(t) \rangle_{L^2}.$$

Quadratic Potential: Error Norm Decay



Time evolution of the error norm for the quadratic potential with $a = 1$. Line styles distinguish the values of b (solid for $b = 0.1$, dashed for $b = 0.05$), while colours indicate the control strategy (black: uncontrolled, blue: LQR feedback control, red: optimal control).

Wasserstein Gradient-Flow Viewpoint

Riemannian structure on probability measures

- $(\mathcal{P}_2(\Omega), W_2)$: space of probability measures on $\Omega \subset \mathbb{R}^d$ with finite second moment, endowed with the 2-Wasserstein distance. Tangent vectors at μ are velocity fields v_t such that (continuity equation)

$$\partial_t \mu_t + \nabla \cdot (\mu_t v_t) = 0.$$

- Free energy functional for McKean-Vlasov dynamics:

$$\mathcal{F}(\mu) = \int_{\Omega} V(x) \mu(x) dx + \frac{1}{2} \iint_{\Omega \times \Omega} W(x-y) \mu(x) \mu(y) dx dy + \sigma \int_{\Omega} \mu(x) \log \mu(x) dx.$$

- The McKean-Vlasov Fokker-Planck equation

$$\partial_t \mu = \sigma \Delta \mu + \nabla \cdot (\mu \nabla V) + \nabla \cdot (\mu \nabla W * \mu)$$

can be written formally as the **gradient flow** of $\mathcal{F}$ in $(\mathcal{P}_2(\Omega), W_2)$:

$$\partial_t \mu_t = \nabla \cdot \left(\mu_t \nabla \frac{\delta \mathcal{F}}{\delta \mu}(\mu_t) \right).$$