IMPERIAL

LAWOC 2026

Controlling interacting particle systems

On the Stabilisation of Interacting Particle Systems and its Connection to Energy Convexification

Dante Kalise, **Lucas M. Moschen**, and Grigorios A. Pavliotis

Department of Mathematics, Imperial College London

30 July 2026

Why Control Probability Laws?

Opinion dynamics: preventing (or accelerating?) polarisation

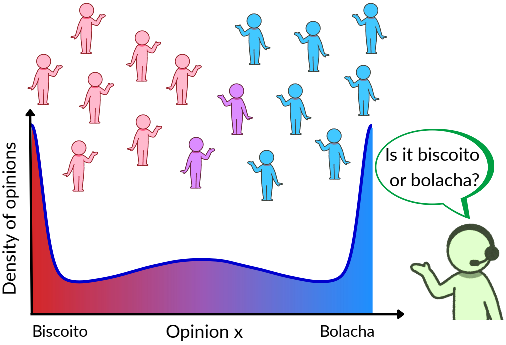

Local interactions can fragment a population into opinion clusters.

- Interacting **particle systems** modelled with SDEs lead to mean-field PDEs for the law $\mu_t \in \mathcal{P}_2(\Omega)$.
- The dynamics can be **slow** or converge to the **undesired** equilibrium.
- **Goal:** modify the dynamics so that, for $\mathsf{dist}(\mu_0, \bar{\mu})$ small,

$$\mathsf{dist}(\mu_t, \bar{\mu}) \leq Ce^{-\delta t}\mathsf{dist}(\mu_0, \bar{\mu}).$$

Control viewpoint in feedback form

$$\partial_t \mu_t = \mathcal{G}(\mu_t) + \mathcal{B}(\mu_t)u(t), \qquad u(t) = \mathcal{K}(\mu_t - \bar{\mu}).$$

Controlling Opinion Dynamics

Local interactions with and without a feedback control

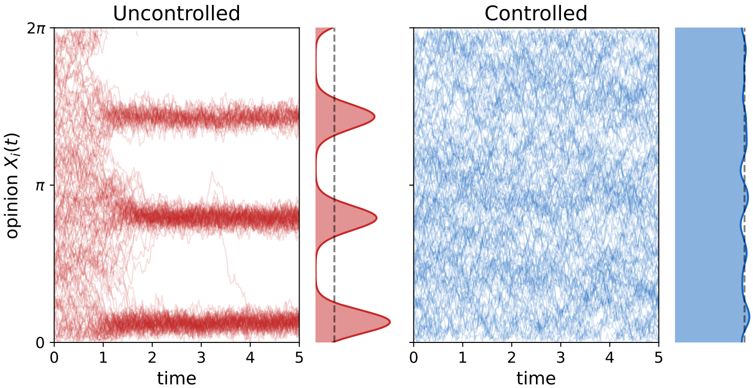

Figure: Agents' opinions following a noisy **Hegselmann–Krause** model on the circle. Without control, the population forms **clusters**. With control, it **remains dispersed**.

The Sampling Problem

McKean–Vlasov dynamics: slow convergence

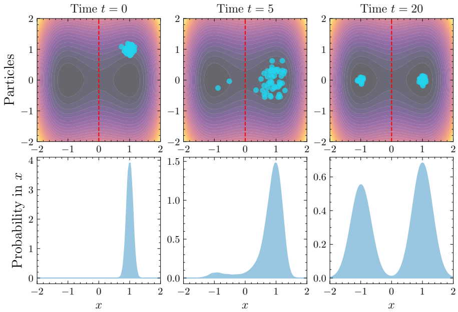

Sampling target $\bar{\mu} \propto e^{-(x^4 - 2x^2)/4}$

$$\mu_t^N = \frac{1}{N} \sum_{i=1}^{N} \delta_{X_t^i}$$

- **Overdamped Langevin:**

$$dX_t^i = -\nabla \big(V + W * \mu_t^N \big)(X_t^i)\, dt + \sqrt{2\sigma}\, dB_t^i.$$

- λ_1 = spectral gap of the linearised Fokker–Planck operator around $\bar{\mu}$ controls the decay rate

$$\|\mu_t - \bar{\mu}\|_{L^2(\bar{\mu}^{-1})} \leq C e^{-\lambda_1 t} \|\mu_0 - \bar{\mu}\|_{L^2(\bar{\mu}^{-1})}.$$

  Nonconvexity or strong W can lead to $0 < \lambda_1 \ll 1$.

- **Feedback controlled dynamics ($\nabla U[\bar{\mu}] = 0$):**

$$dX_t^i = -\nabla \big(V + W * \mu_t^N + U[\mu_t^N] \big)(X_t^i)\, dt + \sqrt{2\sigma}\, dB_t^i.$$

* McKean (1969), Sznitman (1991); Carrillo, McCann and Villani (2003).

This Talk Connects Three Viewpoints
Feedback stabilisation at the distribution level

Control of SDEs

A social planner controls agents at the particle level minimising $\int_0^T \int \ell(x, \mu_t, u)\, d\mu_t\, dt$.

Bensoussan, Frehse & Yam (2013); Carmona & Delarue (2018); Andersson & Djehiche (2011); Pham & Wei (2017); Fornasier & Solombrino (2014); Albi, Choi, Fornasier & Kalise (2017); Djetté, Possamai & Tan (2022); Cardaliaguet, Daudin, Jackson & Souganidis (2023); Barbu (2023).

Control of PDEs

Direct control at the distribution level using FP equations.

$$\partial_t \mu = \mathcal{G}(\mu) + \mathcal{B}(\mu)\, u$$

Blaquère (1992); Fleig & Guglielmi (2017); Aronna & Tröltzsch (2021); Breiten, Kunisch & Pfeiffer (2018); Breiten & Kunisch (2023); Duprez, Morancey & Rossi (2019); Pogodaev & Rossi (2024); Bicego, Kalise & Pavliotis (2025); Albi, Bicego & Kalise (2025).

Sampling & long-time

Exponential mixing, propagation of chaos, metastability.

$$\mathrm{W}_2(\mu_t, \bar{\mu}) \le Ce^{-\lambda t}$$

Bolley, Gentil & Guillin (2012); Carrillo, Gvalani, Pavliotis & Schlichting (2020); Raginsky, Rakhlin & Telgarsky (2017); Chizat (2022); Lelièvre, Nier & Pavliotis (2013); Duncan, Nüsken & Pavliotis (2017); Wang & Chizat (2026); Monmarché (2025); Monmarché & Reygner (2025).

This talk: feedback control $u(t) = \mathcal{K}(\mu_t - \bar{\mu})$ at the PDE level $+\ \delta$-stabilisation + particle implementation $u_j^N(t) = \frac{\pi_j}{N} \sum_\ell \alpha_j(X_t^\ell)$.

The McKean–Vlasov Equation as a Gradient Flow

Free energy $\mathcal{F} = \mathcal{E} + \sigma\,$Ent and descent dynamics

- On $\Omega = \mathbb{T}^d$, define the free energy

$$\mathcal{F}(\mu) := \mathcal{E}(\mu) + \sigma \mathrm{Ent}(\mu) = \int_\Omega V\,d\mu + \frac{1}{2}\iint_{\Omega\times\Omega} W(x-y)\,d\mu(x)\,d\mu(y) + \sigma\int_\Omega \mu \log \mu\,dx.$$

- The mean-field PDE is the W_2-gradient flow of $\mathcal{F}$:

$$\partial_t \mu_t = \nabla\cdot\left(\mu_t \nabla \frac{\delta\mathcal{F}}{\delta\mu}(\mu_t)\right) = \sigma\Delta\mu_t + \nabla\cdot(\mu_t\nabla V) + \nabla\cdot(\mu_t\nabla(W*\mu_t)).$$

  Here $\delta\mathcal{F}/\delta\mu$ denotes the first variation.

- Stationary states solve $\dfrac{\delta\mathcal{F}}{\delta\mu}(\bar\mu) = \mathrm{const} \quad\Longleftrightarrow\quad \bar\mu \propto \exp\left\{-\dfrac{V + W*\bar\mu}{\sigma}\right\}$.

  Convergence depends on the curvature of $\mathcal{F}$ around $\bar\mu$.

* Jordan, Kinderlehrer, Otto (1998), Villani (2009); Carrillo, Gvalani, Pavliotis and Schlichting (2020); Monmarché and Reygner (2023)

The Local Stabilisation Problem
The Wasserstein Hessian generates the linearised dynamics

Given a target equilibrium $\bar{\mu}$ and a rate $\delta > 0$, design **feedback control** so that, locally,

$$\|\mu_t - \bar{\mu}\|_{H^{-1}_{\bar{\mu}}} \leq Ce^{-\delta t}\|\mu_0 - \bar{\mu}\|_{H^{-1}_{\bar{\mu}}}.$$

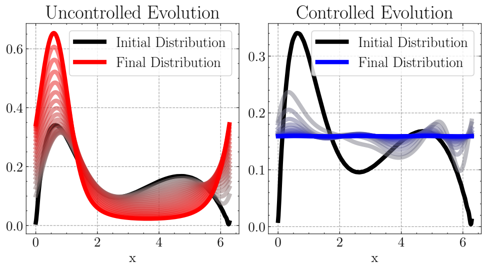

Noisy Kuramoto test case: $V = 0$, $W = -K\cos x$ and the equilibrium $\bar{\mu} = 1/(2\pi)$ is unstable when $K > 2\sigma$.

- Let $\langle \xi, \eta \rangle_X = \int_\Omega \nabla\xi \cdot \nabla\eta \, d\bar{\mu}$, $\quad \|y\|_{H^{-1}_{\bar{\mu}}} := \|\mathcal{I}_{\bar{\mu}}^{-1} y\|_X$.
- Let $y = \mu - \bar{\mu} = -\nabla \cdot (\bar{\mu}\nabla\xi) =: \mathcal{I}_{\bar{\mu}}\xi$.
- The Wasserstein Hessian is represented by a self-adjoint operator with **compact resolvent**

$$A = \mathrm{Hess}_{W_2}\mathcal{F}(\bar{\mu}) \text{ on } X = \dot{H}^1_{\bar{\mu}}.$$

Linearising $\partial_t \mu_t = \mathcal{G}(\mu_t)$ around $\bar{\mu}$

$$L = -\mathcal{I}_{\bar{\mu}} A \mathcal{I}_{\bar{\mu}}^{-1}, \qquad \partial_t y = Ly, \qquad \partial_t \xi = -A\xi.$$

Spectral Analysis of the Wasserstein Hessian
What prevents fast convergence?

Since A has compact resolvent, choose an X-orthonormal eigenbasis

$$Ae_k = \lambda_k e_k, \qquad \xi(t) = e^{-At}\xi_0 = \sum_{k\geq 1} e^{-\lambda_k t}\langle \xi_0, e_k\rangle_X e_k, \qquad \lambda_1 \leq \cdots \leq \lambda_k \to +\infty.$$

For a desired rate $\delta > 0$, set the **slow spectral subspace**

$$\mathcal{X}_\delta = \mathrm{span}\{e_k : \lambda_k \leq \delta\}, \qquad P_\delta \xi = \sum_{\lambda_k \leq \delta} \langle \xi, e_k\rangle_X e_k, \qquad m := \dim \mathcal{X}_\delta < \infty.$$

Only $P_\delta \xi$ can obstruct the prescribed decay rate.

Idea: add m profiles α_j for the slow modes, $\mathcal{X}_\delta \subseteq \mathrm{span}\{\alpha_1, \ldots, \alpha_m\}$,

$$\partial_t \mu_t = \nabla \cdot \left[\mu_t \nabla \left(\frac{\delta \mathcal{F}}{\delta \mu}(\mu_t) + \sum_{j=1}^{m} u_j(t)\alpha_j\right)\right].$$

* Breiten, Kunisch and Pfeiffer (2018), Kalise, Moschen, Pavliotis (2025).

Designing the Feedback Control
Actuation on the slow modes

The controlled linearised dynamics takes the form

$$\partial_t y = Ly + \sum_{j=1}^{m} u_j(t)\,\nabla\cdot(\bar{\mu}\nabla\alpha_j), \quad \partial_t \xi = -A\xi - Bu,$$

where $Bu := \sum_{j=1}^{m} u_j\alpha_j$. Choosing $\alpha_j = e_j$ gives $BB^* = P_\delta$ and yields the infinite-dimensional δ-**Hautus** criterion

$$\ker(\lambda I + A) \cap \ker(B^*) = \{0\} \quad \text{for all } \Re\lambda \geq -\delta,$$

guaranteeing the existence of a **stabilising Riccati feedback** under an additional detectability assumption.

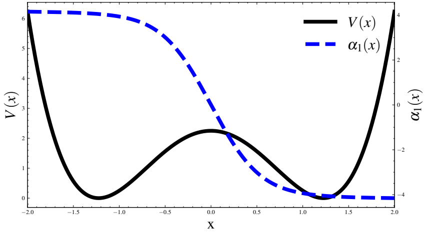

One spatial control profile.

Using the Riccati Equation to Define the Control
Linear-quadratic optimal control formulation

- Penalise the weighted linear dynamics by

$$J(\xi_0, u) := \frac{1}{2}\int_0^\infty e^{2\delta t}\left(\langle M\xi(t), \xi(t)\rangle_X + |u(t)|^2\right) dt.$$

- **Feedback Law:** The optimal control is $u^*(t) = B^*\Pi\xi(t)$, where Π is the stabilising non-negative solution to the Riccati equation

$$-(A-\delta I)\Pi - \Pi(A-\delta I) - \Pi P_\delta \Pi + M = 0.$$

- The closed-loop operator is

$$A_\Pi := A + BB^*\Pi = A + P_\delta\Pi \geq \delta\,\mathsf{Id}\,.$$

Outcome: Linear Exponential Stabilisation

$$\partial_t\xi = -A_\Pi\xi \quad \implies \quad \|\xi(t)\|_X \leq Ce^{-\delta t}\|\xi(0)\|_X.$$

Convexification Interpretation

The feedback changes the local curvature at the target

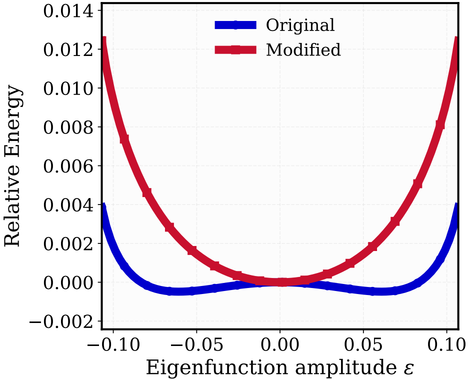

- **Closed-loop free energy:**

$$\mathcal{F}_{\mathsf{cl}}(\mu) = \mathcal{F}(\mu) + \frac{1}{2}\sum_{j=1}^{m} \pi_j \left(\int_{\mathbb{T}^d} e_j \, d(\mu - \bar{\mu}) \right)^2 .$$

- **Modified curvature:** $\partial_t \mu = \nabla \cdot \left(\mu \nabla \frac{\delta \mathcal{F}_{\mathsf{cl}}}{\delta \mu} \right)$ and

$$\mathsf{Hess}_{W_2} \mathcal{F}_{\mathsf{cl}}(\bar{\mu}) = A + P_\delta \Pi = A_\Pi .$$

- **Diagonal construction:** if $Me_k = m_k e_k$, then $\Pi e_k = \pi_k e_k$ and the eigenvalues of A_Π are:

$$\lambda_k^\Pi = \begin{cases} \delta + \sqrt{(\lambda_k - \delta)^2 + m_k} & k \leq m, \\ \lambda_k & k > m. \end{cases}$$

Local Displacement Convexity

From a pointwise Hessian bound to a local property

1. Hessian around $\bar{\mu}$

Let $\mathcal{V}$ be a set of densities where the Hessian is well-defined and

$$\mathcal{V}_r := \left\{ \mu \in \mathcal{V} : \left\| \frac{\mu}{\bar{\mu}} - 1 \right\|_{L^\infty(\mathbb{T}^d)} \leq r \right\}.$$

Under local continuity of the Hessian,

$$\mathrm{Hess}_{\mathrm{W}_2} \mathcal{F}_{\mathrm{cl}}(\mu) \geq \mathrm{Hess}_{\mathrm{W}_2} \mathcal{F}_{\mathrm{cl}}(\bar{\mu}) - \varepsilon(r)\,\mathrm{Id},$$

for $\varepsilon(r) \to 0$. Hence, for $r_\star \ll 1$,

$$\mathrm{Hess}_{\mathrm{W}_2} \mathcal{F}_{\mathrm{cl}}(\mu) \geq \lambda\,\mathrm{Id}, \qquad \mu \in \mathcal{V}_{r_\star}.$$

2. Convexity along nearby geodesics

Let $\mathcal{V}_\star \subseteq \mathcal{V}_{r_\star}$ contain $\bar{\mu}$ and every constant-speed W_2-geodesic joining its elements.

Theorem (Local λ-displacement convexity)

For r_* sufficiently small, every $\mu_0, \mu_1 \in \mathcal{V}_\star$, and $s \in [0,1]$,

$$\mathcal{F}_{\mathrm{cl}}(\mu_s) \leq (1-s)\mathcal{F}_{\mathrm{cl}}(\mu_0) + s\mathcal{F}_{\mathrm{cl}}(\mu_1) - \frac{\lambda}{2} s(1-s)\mathrm{W}_2^2(\mu_0, \mu_1).$$

Feedback Control → Pointwise Convexification → Local Convexification by continuity

Local Exponential Stabilisation

Guarantee for the fully nonlinear closed-loop PDE

Theorem (Local exponential stabilisation)

For every rate $\gamma \in (0, \delta]$*, there exist* $\varepsilon, C_\gamma > 0$ *such that if* $\nu_0 = \mu_0 - \bar{\mu}$ *satisfies* $\mathcal{I}_{\bar{\mu}}^{-1}\nu_0 \in \dot{H}^2(\mathbb{T}^d)$ *and* $\|\nu_0\|_{L^2(\mathbb{T}^d)} \leq \varepsilon$*, the closed-loop solution is global and*

$$\|\mu_t - \bar{\mu}\|_{H_{\bar{\mu}}^{-1}} \leq C_\gamma e^{-\gamma t} \|\mu_0 - \bar{\mu}\|_{L^2(\mathbb{T}^d)}.$$

How the proof works

- Setting $w_t := e^{\gamma t}\nu_t$ gives $\partial_t w_t - (L_\Pi + \gamma I) w_t = \mathfrak{R}_\gamma(w)$.
- **Weighted Linear estimate:** Since $\gamma \leq \delta$, maximal regularity yields $\|w\|_{\mathcal{Y}_T} \leq C_{\text{lin}}\big(\|\nu_0\|_{L^2} + \|\mathfrak{R}_\gamma(w)\|\big)$, uniformly in T.
- **Quadratic remainder:** Near $\bar{\mu}$, $\|\mathfrak{R}_\gamma(w)\| \leq C_{\text{rem}}\|w\|_{\mathcal{Y}_T}^2$.
- **Bootstrap:** Thus $\|w\|_{\mathcal{Y}_T} \leq C_{\text{lin}}\|\nu_0\|_{L^2} + C\|w\|_{\mathcal{Y}_T}^2$. For small ν_0, the solution remains below the continuation radius, so $T_{\max} = \infty$ and $\nu_t = O(e^{-\gamma t})$.

Local Exponential Stabilisation

Guarantee for the fully nonlinear closed-loop PDE

Theorem (Local exponential stabilisation)

For every rate $\gamma \in (0, \delta]$, *there exist* $\varepsilon, C_\gamma > 0$ *such that if* $\nu_0 = \mu_0 - \bar{\mu}$ *satisfies* $\mathcal{I}_{\bar{\mu}}^{-1}\nu_0 \in \dot{H}^2(\mathbb{T}^d)$ *and* $\|\nu_0\|_{L^2(\mathbb{T}^d)} \leq \varepsilon$, *the closed-loop solution is global and*

$$\|\mu_t - \bar{\mu}\|_{H_{\bar{\mu}}^{-1}} \leq C_\gamma e^{-\gamma t}\|\mu_0 - \bar{\mu}\|_{L^2(\mathbb{T}^d)}.$$

Corollary (Wasserstein stabilisation)

By the comparison inequality $\mathrm{W}_2(\mu_t, \bar{\mu}) \leq 2\|\mu_t - \bar{\mu}\|_{H_{\bar{\mu}}^{-1}}$, *we have*

$$\mathrm{W}_2(\mu_t, \bar{\mu}) \leq 2C_\gamma e^{-\gamma t}\|\nu_0\|_{L^2(\mathbb{T}^d)}.$$

Kuramoto under Feedback

Feedback accelerates stable equilibria and stabilises unstable ones

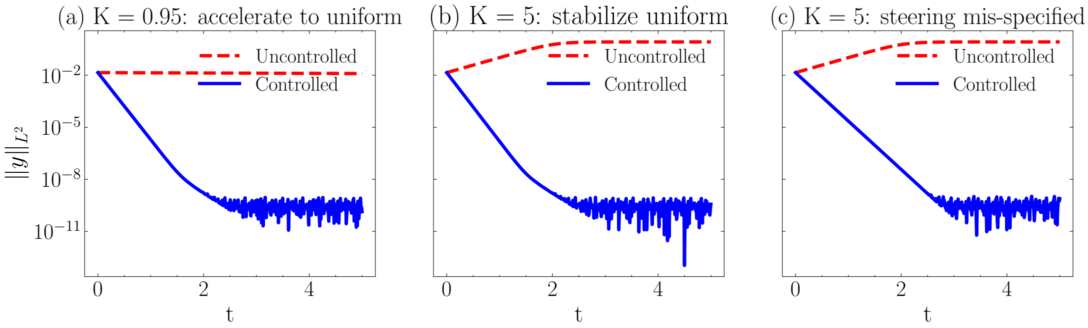

Figure: Control of the noisy Kuramoto model ($\sigma = 0.5, W(x) = -K \cos x$). Panels (a) and (b) target the uniform distribution $\bar{\mu} = (2\pi)^{-1}$, while panel (c) targets a selected translated synchronised equilibrium.

Feedback on the Sphere

The same spectral design extends to closed manifolds

On a closed Riemannian manifold (Ω, g),

$$\mathcal{F}_g(\mu) = \sigma \mathsf{Ent}_g(\mu) + \int_\Omega V\, d\mu, \qquad \bar{\mu} \propto e^{-V/\sigma}.$$

The Wasserstein Hessian at $\bar{\mu}$ is

$$\mathsf{Hess}_{\mathsf{W}_2} \mathcal{F}_g(\bar{\mu})[\nabla_g \varphi, \nabla_g \varphi] = \sigma \int_\Omega |\nabla_g^2 \varphi|_g^2 \, d\bar{\mu} + \int_\Omega \big(\sigma \mathsf{Ric}_g + \nabla_g^2 V\big)[\nabla_g \varphi, \nabla_g \varphi]\, d\bar{\mu}.$$

- The **Ricci contribution** is a bounded lower-order perturbation, so the Hessian remains **self-adjoint with compact resolvent**.
- The same finite-rank **Riccati feedback** applies.

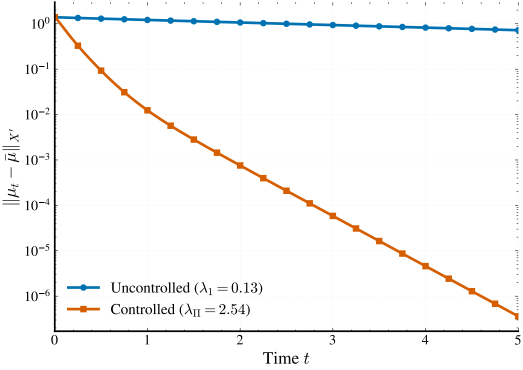

On $\mathbb{S}^2$, with $V(\theta) = -2\cos^2\theta$ and $\sigma = 0.5$.

From the PDE to the Particle System

How do we apply the PDE control in the SDEs?

The PDE control is computed with $\alpha_j = e_j$. By plugging in the empirical measure, we have

$$u_j(t) = \pi_j \int_\Omega e_j \, d\mu_t \quad \Rightarrow \quad u_j^N(t) = \frac{\pi_j}{N} \sum_{\ell=1}^{N} e_j(X_t^\ell).$$

Advantage: The control is evaluated directly from particle positions. **No density reconstruction.**

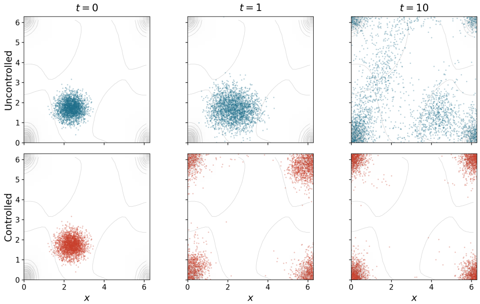

Uncontrolled (top) vs. feedback-controlled (bottom).

Controlled N-particle system

$$dX_t^i = -\left(\nabla V(X_t^i) + \frac{1}{N} \sum_{k=1}^{N} \nabla W(X_t^i - X_t^k) + \sum_{j=1}^{m} u_j^N(t) \nabla e_j(X_t^i) \right) dt + \sqrt{2\sigma} \, dB_t^i.$$

The First Estimate

Finite-time propagation of chaos

Theorem

For every $T > 0$ and $t \in [0, T]$, the mean squared Wasserstein distance satisfies

$$\mathbb{E}\mathrm{W}_2^2(\mu_t^N, \bar{\mu}) \le \frac{C_0}{N}(e^{C_1 t} - 1) + C_1 \beta_d(N) + C_2 e^{-2\delta t} \|\nu_0\|_{L^2(\mathbb{T}^d)}^2.$$

- $\beta_d(N)$ is the empirical measure error and comes from Fournier–Guillin (2015).
- **Issue:** The Gronwall error $e^{C_1 t}/N$, due to the coupling error $\mathrm{W}_2(\mu_t^N, \bar{\mu}_t^N)$, blows up for large t.
- The estimate ceases to be useful beyond the timescale $t \approx (\log N)/C_1$.
- **Idea:** The PDE provides strict **local contraction** $\mathrm{W}_2(S_T^{\Pi}\mu, \bar{\mu}) \le q\mathrm{W}_2(\mu, \bar{\mu})$ for $q < 1$. We can avoid global Gronwall by restarting the estimate over time blocks of length T.

Long-Time Confinement

PDE contraction is stable under empirical approximation

Theorem

For large N and a given initial configuration $\mathbf{X}_0 = \mathbf{x}$, if $\mathrm{W}_2(\mu_0^N, \bar{\mu})$ is small, the empirical measure satisfies:

$$\mathbb{E}_{\mathbf{x}} \mathrm{W}_2^2(\mu_t^N, \bar{\mu}) \le Ce^{-2\gamma t} \mathrm{W}_2^2(\mu_0^N, \bar{\mu}) + C\beta_d(N) + C \min\{1, te^{-cN}\}.$$

- The distance decays exponentially until it hits the statistical floor $\beta_d(N)$.
- The error is bounded by $C\beta_d(N)$ for **exponentially long times** $t \le e^{c' N}$.
- The feedback provides a particle-level burn-in accelerator.
- But rare exits from the local attraction region **remain possible**.

Summary and Open Questions

What I showed you

- A **feedback control design** for Wasserstein gradient flows via an LQR problem.
- The same Riccati operator locally **stabilises** the dynamics and **convexifies** the energy.
- **Particle implementation:** m empirical observables and confinement in expectation for times $e^{c'N}$.

Questions for control and optimisation

- Can feedback help **globally**, not only near $\bar{\mu}$?
- Computing eigenfunctions in high dimensions is a complicated task.
- What happens if the normalising constant Z is unknown? $\rightarrow$ work in progress
- Extension to kinetic models?

Happy to discuss any of these!

Papers:

2026: Linearization-Based Feedback Stabilization of McKean–Vlasov PDEs (to appear in SICON)

2026: Feedback Control and Local Convexification of Wasserstein Gradient Flows

lucas.moschen22@imperial.ac.uk

Multi-Species Coupled Dynamics

What matters is the gradient-flow structure with entropy

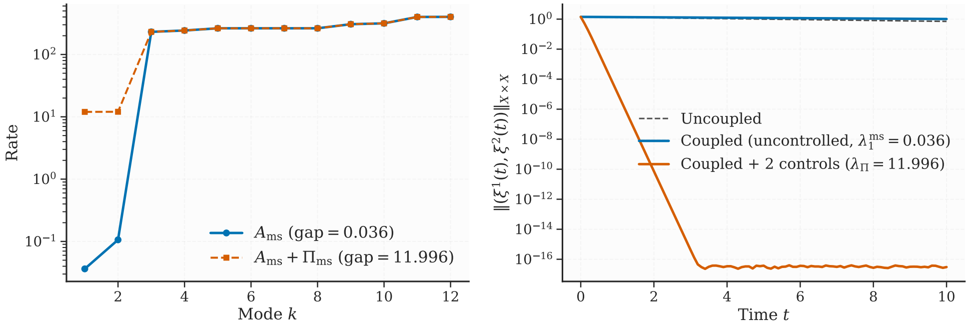

Figure: Feedback stabilisation of a two-species system on $\mathbb{T}$ with $V(x) = -2\cos(4\pi x)$, $\sigma = 0.5$, and cross-interaction $W_{ij}(x) = -0.5\cos(2\pi x)$. A two-control Riccati feedback computed on the product space $X \times X$ shifts the two slow joint modes above the target rate $\delta = 5$ (left, red).

Ric_g

Definition

For the Riemann curvature tensor R_g and an orthonormal basis $\{e_i\}_{i=1}^d$ of $T_x\Omega$,

$$\mathrm{Ric}_g(v, w) := \sum_{i=1}^{d} \langle R_g(e_i, v)w, e_i \rangle_g \, .$$

For unit v, $\mathrm{Ric}_g(v, v)$ is the sum of the sectional curvatures in the directions orthogonal to v.

Examples

$$\begin{aligned} \mathrm{Ric}_g &= 0 && \text{on the flat } \mathbb{T}^d, \\ \mathrm{Ric}_g &= (d-1)g && \text{on the unit } \mathbb{S}^d. \end{aligned}$$

In particular, $\mathrm{Ric}_g = g$ on the unit sphere $\mathbb{S}^2$.

Why does it appear?

The Bochner identity is

$$\begin{aligned} \frac{1}{2}\Delta_g |\nabla_g \phi|_g^2 = |\nabla_g^2 \phi|_g^2 + \langle \nabla_g \phi, \nabla_g \Delta_g \phi \rangle_g \\ + \mathrm{Ric}_g[\nabla_g \phi, \nabla_g \phi]. \end{aligned}$$

Applied along W_2-geodesics, it gives

$$\begin{aligned} &\mathrm{Hess}_{W_2} \mathrm{Ent}_g(\mu)[\nabla_g \phi, \nabla_g \phi] \\ &\quad = \int_\Omega |\nabla_g^2 \phi|_g^2 \, d\mu + \int_\Omega \mathrm{Ric}_g[\nabla_g \phi, \nabla_g \phi] \, d\mu. \end{aligned}$$