IMPERIAL

SPECIAL SESSION 18 · OPTIMAL CONTROL, OPTIMAL TRANSPORT AND RELATED FIELDS

Feedback Control and Local Convexification of Wasserstein Gradient Flows

Local stabilisation and local convexification via an LQR problem

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

Department of Mathematics, Imperial College London

September 2026

Why Control Probability Laws?

Opinion dynamics: preventing (or accelerating?) polarisation

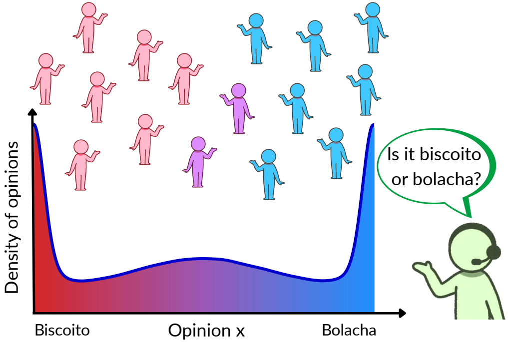

Local interactions can form opinion clusters.

- 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** (λ_1 small):

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

or converge to the **undesired** equilibrium.

- **Goal:** modify the dynamics so, for $\mathrm{dist}(\mu_0, \bar{\mu}) < \varepsilon$,

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

Control viewpoint in feedback form

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

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

Controlling Opinion Dynamics
Local interactions with and without a feedback control

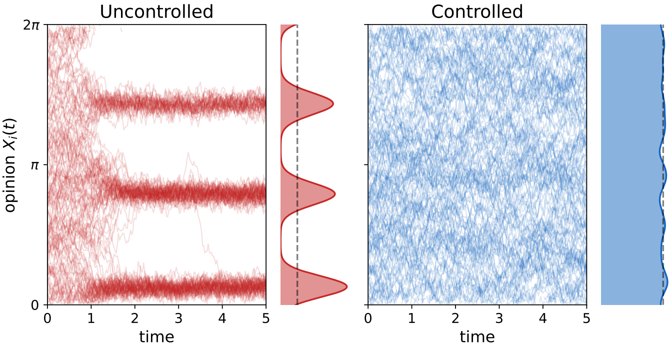

Figure: Agents' opinions following a noisy **Hegselmann–Krause** model on the circle. Without control, the population forms **clusters**. With control, it **remains dispersed**.

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

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.
$$W_2(\mu_t,\bar\mu)\ \le Ce^{-\lambda t}$$

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: PDE feedback stabilisation + Wasserstein Hessian + local convexification.

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 entropy-regularized free energy $(\sigma > 0)$

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

- 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, W is even, and $V, W \in W^{2,\infty}(\mathbb{T}^d)$.

- 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); Chizat (2022); 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,

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

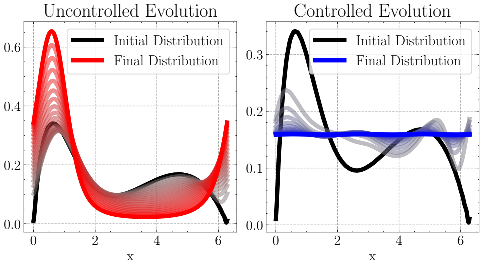

Noisy Kuramoto test case: $V = 0$, $W = -K\cos x$ and the equilibrium $\bar{\mu} = 1/(2\pi)$ is unstable when $K > 2\sigma$.

- 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 directions with slow convergence

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.

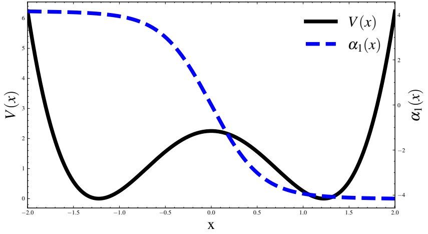

One spatial control profile.

Using the Riccati Equation to Define the Control

Linear-quadratic optimal control formulation

- Let Q be a bounded and nonnegative operator. Penalise the weighted linear dynamics by

$$J(\xi_0, u) := \frac{1}{2}\int_0^\infty e^{2\delta t}\left(\langle Q\xi(t), \xi(t)\rangle_X + |u(t)|^2\right) dt, \quad Qe_k = q_k e_k, q_k > 0 \text{ for } k \leq m.$$

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

$$-(A - \delta I)\Pi - \Pi(A - \delta I) - \Pi P_\delta \Pi + Q = 0 \quad \Rightarrow \quad \Pi e_k = \pi_k e_k.$$

- The closed-loop operator is

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

Outcome: Linear Exponential Stabilisation

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

Convexification Interpretation

The feedback changes the local curvature at the target

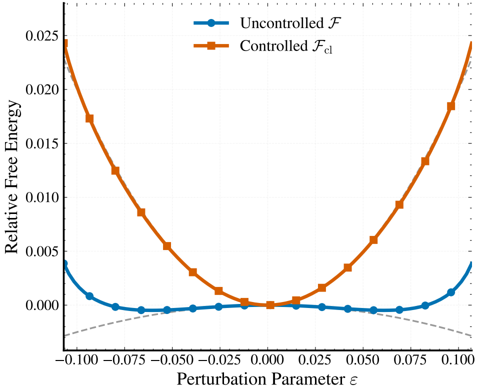

$\mathcal{F}(\mu_\varepsilon)$ and $\mathcal{F}_{\mathrm{cl}}(\mu_\varepsilon)$ for $\mu_\varepsilon = (\mathrm{Id} + \varepsilon \nabla \phi)_\sharp \bar{\mu}$.

- **Closed-loop free energy:**

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

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

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

- **Spectral shift:**

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

Hence $A_\Pi \geq \lambda_\Pi I$, with $\lambda_\Pi := \min_k \lambda_k^\Pi > \delta$.

Local Displacement Convexity

From a pointwise Hessian bound to a local property

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

Fix $0 < \lambda < \lambda_\Pi$ and $\mathcal{V}$ a set of densities where $\mathrm{Hess}_{\mathrm{W}_2}\mathcal{E}$ is well-defined. Set

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

Under local continuity of $\mathrm{Hess}_{\mathrm{W}_2}\mathcal{E}$ and $e_j \in W^{2,\infty}(\mathbb{T}^d)$, relative entropy-form bounds give, for small $r_\star$,

$$\mathrm{Hess}_{\mathrm{W}_2}\mathcal{F}_{\mathrm{cl}}(\mu)(\nabla\phi, \nabla\phi) \geq \lambda \|\nabla\phi\|^2_{L^2(\mu)},$$

for $\mu \in \mathcal{V}_{r_\star}$ and $\phi \in W^{2,\infty}(\mathbb{T}^d)$.

2. Convexity along nearby geodesics

Fix $\alpha \in (0,1)$ and set

$$\mathcal{N}_\varepsilon^\alpha := \left\{ \mu \in \mathcal{P}_2(\mathbb{T}^d) : \|\mu - \bar{\mu}\|_{C^{0,\alpha}} \leq \varepsilon \right\}.$$

Theorem (Local λ-displacement convexity)

For small ε, the geodesic joining any $\mu_0, \mu_1 \in \mathcal{N}_\varepsilon^\alpha$ satisfies, for $s \in [0,1]$,

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

Feedback Control $\rightarrow$ Pointwise Convexification $\rightarrow$
Local Convexification by continuity + regularity

Local Exponential Stabilisation

Guarantee for the fully nonlinear closed-loop PDE

Theorem (Local exponential stabilisation)

Assume **local existence** *and a quadratic local* **Lipschitz estimate** *for the nonlinear remainder. For every rate $\gamma \in (0, \lambda_\Pi)$, there exist $\varepsilon, C_\gamma > 0$ such that if $\nu_0 = \mu_0 - \bar{\mu} \in L^2(\mathbb{T}^d)$ and $\|\nu_0\|_{L^2(\mathbb{T}^d)} \leq \varepsilon$, the closed-loop solution is unique, global, and*

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

Corollary (Wasserstein stabilisation)

By the comparison inequality (Peyre, 2018) $\mathrm{W}_2(\mu_t, \bar{\mu}) \leq 2\|\mu_t - \bar{\mu}\|_{H^{-1}_{\bar{\mu}}}$, we have

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

Kuramoto under Feedback

Feedback accelerates stable equilibria and stabilises unstable ones

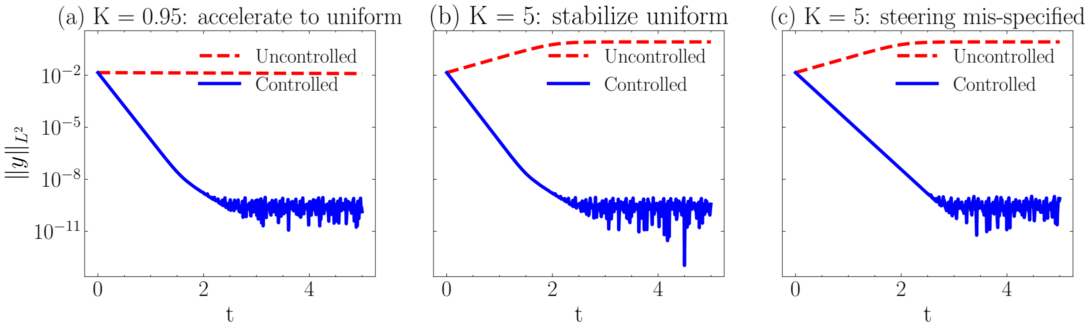

Figure: Control of the noisy Kuramoto model ($\sigma = 0.5, W(x) = -K \cos x$). Panels (a) and (b) target the uniform distribution $\bar{\mu} = (2\pi)^{-1}$, while panel (c) targets a selected translated synchronised equilibrium.

Two Populations under Feedback

Controlling one population can stabilise both

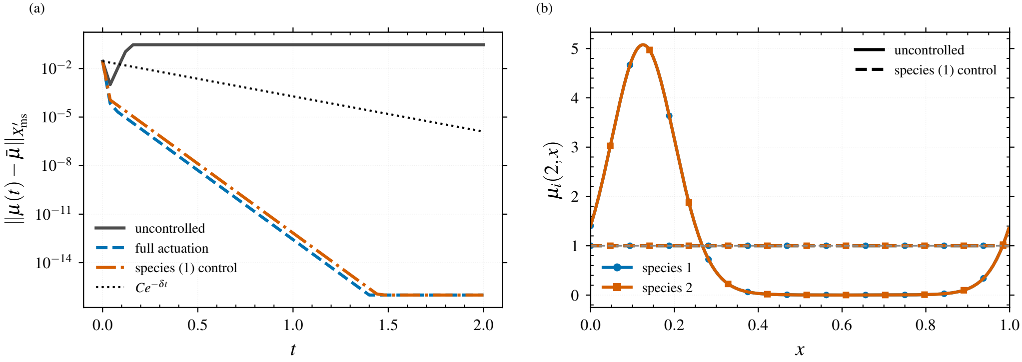

Figure: Two-population Kuramoto model: (a) full versus species 1 actuation; (b) both populations become uniform with species 1 control.

Summary and Open Questions

What I showed you

- A **feedback control design** for Wasserstein gradient flows via an LQR problem.
- A Riccati operator locally **stabilises** the nonlinear dynamics.
- The same operator preserves the gradient-flow structure and makes the modified energy **locally displacement convex**.

Questions for control and optimisation

- Can feedback help **globally**, not only near $\bar{\mu}$?
- Extension to kinetic models?
- How to efficiently compute the eigenfunctions of A in high dimensions?
- Feedback control strategy at the SDE level $\rightarrow$ check session 27.

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

Local Exponential Stabilisation

Guarantee for the fully nonlinear closed-loop PDE

Theorem (Local exponential stabilisation)

Assume **local existence** *and a local* **Lipschitz estimate** *for the nonlinear remainder. For every rate $\gamma \in (0, \lambda_\Pi)$, there exist $\varepsilon, C_\gamma > 0$ such that if $\nu_0 = \mu_0 - \bar{\mu} \in L^2(\mathbb{T}^d)$ and $\|\nu_0\|_{L^2(\mathbb{T}^d)} \leq \varepsilon$, the closed-loop solution is unique, global, and*

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

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 $\gamma < \lambda_\Pi$, 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 $\bar{\mu}$, $\|\mathfrak{R}_\gamma(w)\| \leq C_{\text{rem}} \|w\|^2_{\mathcal{Y}_T}$.
- **Bootstrap:** Thus $\|w\|_{\mathcal{Y}_T} \leq C_{\text{lin}} \|\nu_0\|_{L^2} + C\|w\|^2_{\mathcal{Y}_T}$. For small ν_0, the solution remains below the continuation radius, so $T_{\max} = \infty$ and $\nu_t = O(e^{-\gamma t})$.