IMPERIAL

COML 2026 · NextGen Synergy: Control Theory & Machine Learning

On the Local 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

29 April 2026

Why Control Probability Laws?
Controlling the behaviour of particle systems

Feedback metaphor: in a Scottish ceilidh, a caller guides the group's choreography. Photo © Dave Conner (CC BY 4.0); clip-art adaptation by L. Moschen.

- Sampling and **many-particle systems** can be modelled microscopically via SDEs, leading to PDEs for the law $\mu_t \in \mathcal{P}_2(\Omega)$.
- The dynamics may converge **slowly**, or may settle at an equilibrium **different** from a desired target $\bar{\mu}$.
- **Goal:** change the dynamics so, 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

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

- **Question:** what is the impact of the control on the energy landscape?

Sampling Viewpoint
The long-term behaviour of particle systems can be slow

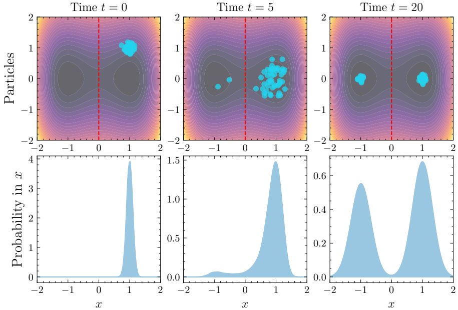

Particles moving in a nonconvex energy landscape.

Sampling case $W = 0$: Langevin dynamics target $\nu \propto e^{-V/sigma}$.

- **Overdamped Langevin dynamics:**

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

- In the mean-field limit[1] $\mu_t = \mathrm{Law}(X_t)$ solves

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

- Under a **functional inequality** such as a Poincaré or log-Sobolev inequality,

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

For nonconvex V, λ can be small $\Rightarrow$ **slow convergence**; with W, several macroscopic equilibria can coexist.

[1] McKean (1969), Sznitman (1991); Carrillo, McCann and Villani (2003).

Wasserstein Gradient Flows for Mean-Field Models

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[2] 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 $\delta \mathcal{F} / \delta \mu$ is the first variation.

- Stationary states solve

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

  Slow convergence and multiple equilibria depend on the free-energy $\mathcal{F}$.[3]

[2] Jordan, Kinderlehrer, Otto (1998), Villani (2009)

[3] Carrillo, Gvalani, Pavliotis and Schlichting (2020); Monmarché and Reygner (2023)

Where This Fits
Three neighbouring literatures

Control of SDEs

Mean-field type control: a central planner acts on McKean-Vlasov SDEs or particle systems.

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

Control of PDEs

Direct control at the level of the law: Fokker-Planck, McKean-Vlasov, or continuity equations.

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

Sampling

Long-time behaviour, slow relaxation, metastability, propagation of chaos.

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

This talk: PDE-level feedback control for local stabilisation, connected to the geometry of the free energy.

The Local Stabilisation Problem

Relating the Hessian with the linearisation

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

$$\|\mu_t - \bar{\mu}\|_{\mathcal{T}} \leq Ce^{-\delta t}\|\mu_0 - \bar{\mu}\|_{\mathcal{T}}, \quad \|\mu_t - \bar{\mu}\|_{\mathcal{T}} := \|\mathcal{I}_{\bar{\mu}}^{-1}(\mu_t - \bar{\mu})\|_X, \quad \langle \xi, \eta \rangle_X = \int_\Omega \nabla\xi \cdot \nabla\eta \, d\bar{\mu}.$$

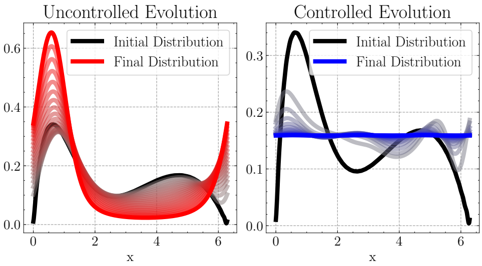

Noisy Kuramoto test case: $V = 0$, $W = -K\cos x$ and the equilibrium $\bar{\mu} = 1/(2\pi)$ is unstable when $K > 1$.

- Let $\rho = \mu - \bar{\mu} = -\nabla \cdot (\bar{\mu}\nabla\xi) =: \mathcal{I}_{\bar{\mu}}\xi$.
- The Wasserstein Hessian is represented by a self-adjoint operator with **compact resolvent**

$$A = \mathrm{Hess}_{\mathrm{W}_2}\mathcal{F}(\bar{\mu}) \text{ on } X = H^1_{\mathrm{zm}}(\bar{\mu}),$$

and measures the local curvature of $\mathcal{F}$.

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

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

Spectral Analysis of the Wasserstein Hessian
What prevents fast convergence?

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

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

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

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

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

Idea[4]: add m profiles α_j for the slow modes, $\mathcal{X}_\delta \subseteq \mathrm{span}\{\alpha_1, \ldots, \alpha_m\}$,

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

[4] Breiten, Kunisch and Pfeiffer (2018), Kalise, Moschen, Pavliotis (2025).

Control on the Slow Spectral Window
Actuation on the slow modes

The controlled linearised dynamics takes the form

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

where $Bu := \sum_{j=1}^{m} u_j \alpha_j$. We set $\alpha_j = \phi_j$ so that

$$BB^* = P_\delta.$$

Under this setting, the system satisfies the infinite-dimensional δ-**Hautus** condition

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

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

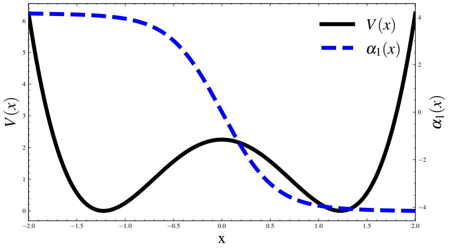

One spatial control profile.

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

- Penalise the weighted linear dynamics by

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

- The optimal control is

$$u^*(t) = B^*\Pi\xi(t),$$

where Π is the stabilising nonnegative solution to the Riccati equation

$$-(A-\delta I)\Pi - \Pi(A-\delta I) - \Pi\overbrace{BB^*}^{P_\delta}\Pi + M = 0.$$

- If $M\phi_k = m_k\phi_k$, then[5] $\Pi\phi_k = \pi_k\phi_k$ and

$$\pi_k = -(\lambda_k - \delta) + \sqrt{(\lambda_k-\delta)^2 + m_k} \quad (1 \le k \le m).$$

Closed-loop dynamics

$$\partial_t\xi = -\underbrace{(A + BB^*\Pi)}_{A_\Pi}\xi.$$

Closed-loop eigenvalues

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

Outcome

$$\|\xi(t)\|_X \le Ce^{-\delta t}\|\xi(0)\|_X.$$

[5] Kalise, Moschen, Pavliotis (2026).

Energy Convexification

Geometric interpretation of the spectral shift

Closed-loop free energy (locally around $\bar{\mu}$)

$$\mathcal{F}_{\mathsf{cl}}(\mu) = \mathcal{F}(\mu) + \frac{1}{2}\langle \mathcal{I}_{\bar{\mu}} P_\delta \Pi \mathcal{I}_{\bar{\mu}}^{-1}(\mu - \bar{\mu}), \mu - \bar{\mu}\rangle_{\mathcal{T}}$$

- The first variation of $\mathcal{F}_{\mathsf{cl}}$ vanishes at $\bar{\mu}$.
- The Wasserstein Hessian of $\mathcal{F}_{\mathsf{cl}}$ at $\bar{\mu}$ coincides with the closed-loop operator:

$$\mathsf{Hess}_{\mathsf{W}_2}\mathcal{F}_{\mathsf{cl}}(\bar{\mu}) = A_\Pi \succeq \lambda_\Pi I > \delta I.$$

- The feedback acts as a **local convexifier** via the relation $L_\Pi := -\mathcal{I}_{\bar{\mu}} A_\Pi \mathcal{I}_{\bar{\mu}}^{-1}$.

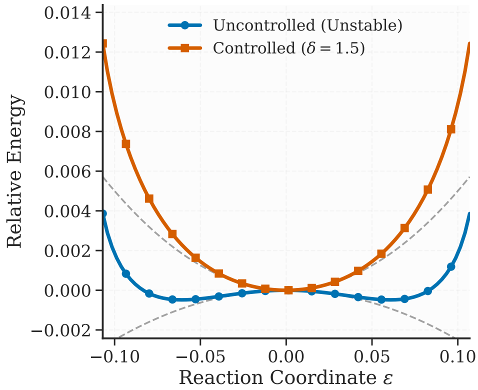

Free energy along a fixed direction ξ (Kuramoto model).

Local Exponential Stabilisation

Guarantee for the fully nonlinear closed-loop PDE

Local exponential stabilisation

For sufficiently small perturbations $\mu_0 = \bar{\mu} + \nu_0, \nu_0 \in \mathcal{I}_{\bar{\mu}}(H^2(\Omega)/\mathbb{R})$, the feedback law $u(t) = B^* \Pi \xi(t)$ yields a unique global solution satisfying

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

Equivalently, $\xi(t) = \mathcal{I}_{\bar{\mu}}^{-1}(\mu_t - \bar{\mu})$ decays exponentially in X with rate at least δ.

How the proof works

- **Linear stability:** The operator A_Π provides an exponentially stable semigroup $\|e^{-A_\Pi t}\| \leq e^{-\delta t}$.
- This extends to the density level by the isomorphism $\mathcal{I}_{\bar{\mu}}$.
- **Nonlinear remainder:** The fully nonlinear terms are assumed to be locally Lipschitz (strictly quadratic near $\bar{\mu}$). For McKean–Vlasov, these are verified under $V \in W^{1,\infty}(\Omega) \cap W^{2,p}(\Omega)$ and $W \in W^{2,\infty}(\Omega)$.
- **Contraction mapping:** A Banach fixed-point argument in an exponentially weighted space closes the local estimate.

Kuramoto under Feedback

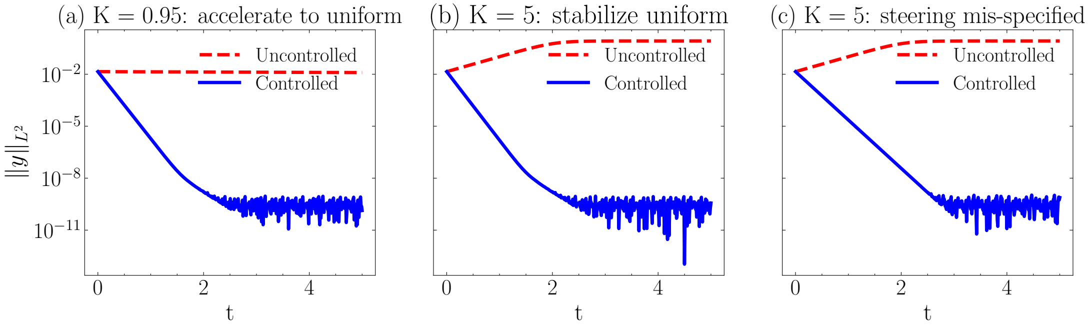

Figure: Control of the noisy Kuramoto model ($\sigma = 0.5$, $W(x) = -K\cos x$). Panels (a) and (b) target the uniform distribution $\bar{\mu} = (2\pi)^{-1}$, while panel (c) targets a selected translated synchronised equilibrium.

Using the Designed Feedback Control at the Particle Level

With actuators $\alpha_j = \phi_j$ and $\Pi\phi_j = \pi_j\phi_j$, the feedback control is given by

$$u_j(t) = (B^*\Pi\xi_t)_j = \pi_j \langle \xi_t, \phi_j \rangle_X = \pi_j \int_\Omega \phi_j(x)\, d(\mu_t - \bar{\mu})(x) = \pi_j \int_\Omega \phi_j(x)\, d\mu_t$$

since $\phi_j \in X$ so $\int_\Omega \phi_j\, d\bar{\mu} = 0$. Replacing μ_t by $\mu_t^N = \frac{1}{N}\sum_{\ell=1}^N \delta_{X_t^\ell}$ yields $u_j^N(t) = \frac{\pi_j}{N}\sum_{\ell=1}^N \phi_j(X_t^\ell)$.

The PDE feedback is implemented from particle observables, without reconstructing the density.

Controlled N-particle system

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

Combining[8] with the local PDE result gives, for $t \le T$ and $\beta_d(N) \to 0$,

$$\mathbb{E}\, \mathrm{W}_2^2(\mu_t^N, \bar{\mu}) \le C_T\, \beta_d(N) + Ce^{-2\delta t}\|\mu_0 - \bar{\mu}\|_{\mathcal{T}}^2.$$

[8] Fournier and Guillin (2015): on compact domains, one may take $\beta_d(N) = N^{-1/2}$ for $d < 4$, $N^{-1/2}\log(1+N)$ for $d = 4$, and $N^{-2/d}$ for $d > 4$.

Effect of the Feedback Control on Particle Convergence

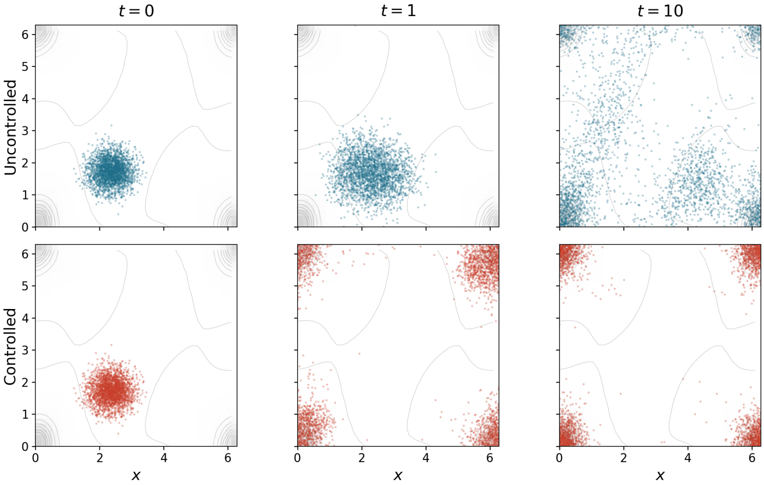

Sample clouds for uncontrolled and feedback-controlled overdamped Langevin dynamics on $\mathbb{T}^2$.

Conclusions

What the framework gives

- This work brings the **Riccati feedback** framework into Wasserstein geometry.
- The feedback **modifies** the local curvature of the energy functional at the target and **stabilises** the nonlinear dynamics exponentially.
- The same design can be implemented at the **particle level**.

Papers behind the talk:

2025: Linearization-Based Feedback Stabilisation of McKean–Vlasov PDEs

2026: Feedback Control and Local Convexification of Wasserstein Gradient Flows

Contact: lucas.moschen22@imperial.ac.uk

Next question: sample from $\bar{\mu}(dx) \propto e^{-V(x)}\,dx$ using feedback without the normalisation constant.

New problem: accelerating sampling becomes an **eigenvalue problem** for the Hessian of the free energy.

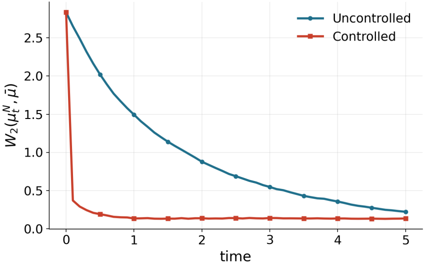

Beyond local stabilisation, can feedback control become a practical mechanism for sampling?

Multi-Species Coupled Dynamics

What matters is the gradient-flow structure with entropy

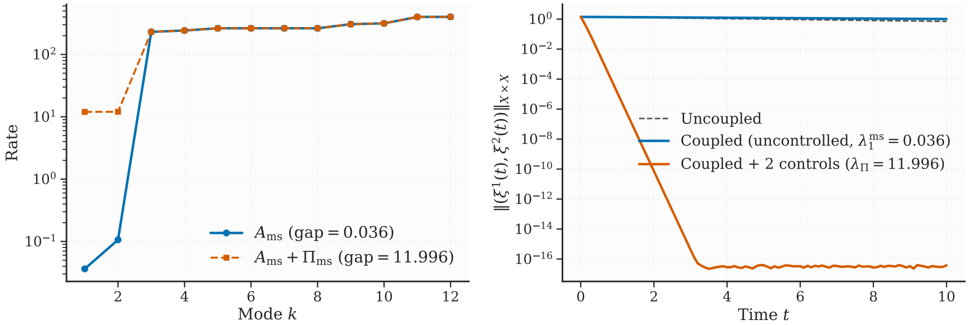

Figure: Feedback stabilisation of a two-species system on $\mathbb{T}$ with $V(x) = -2\cos(4\pi x)$, $\sigma = 0.5$, and cross-interaction $W_{ij}(x) = -0.5\cos(2\pi x)$. A two-control Riccati feedback computed on the product space $X \times X$ shifts the two slow joint modes above the target rate $\delta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