Dynamical fluctuations in systems with and without detailed balance

Robert Jack (University of Cambridge)
ICTS, 2017

Joint work with:
Johannes Zimmer, Marcus Kaiser (University of Bath)

[Kaiser, RLJ and Zimmer, J Stat Phys **168**, 259 (2017);
also arXiv:1708.01453]

Outline

– *Motivation*: Monte Carlo (or MCMC) in statistical physics,
+ is it useful to break detailed balance?

– *Large-deviation analysis* of convergence to equilibrium
+ theory of interacting particles on hydrodynamic scale

– *Geometry* of reversible / irreversible processes,
+ implications for convergence to equilibrium

– Outlook

Motivation

Physical system described by a stochastic model with discrete states (configurations) $\mathcal{C}$

Eg, Ising model, simple exclusion processes, Markov state models of proteins

Transition rates $r(\mathcal{C} \to \mathcal{C}')$.
In the steady state, the probability of state $\mathcal{C}$ is $\pi(\mathcal{C})$.

The stochastic model might be a "realistic" model of the physical system's dynamical properties, or a Monte Carlo algorithm for sampling from π.

Detailed balance: In some special cases we have

$$\pi(\mathcal{C})r(\mathcal{C} \to \mathcal{C}') = \pi(\mathcal{C}')r(\mathcal{C}' \to \mathcal{C})$$

in which case the steady state is time-reversal symmetric

Motivation

Systems with detailed balance:

are typically models for equilibrium systems

are convenient to use in Monte Carlo algorithms

have steady states without persistent currents (time-reversal symmetry)

Systems without detailed balance:

are typically models for non-equilibrium systems

are less often used in Monte Carlo algorithms

have steady states with persistent currents

Comparing these cases, we can learn
(i) how to model non-eqm systems (in contrast with eqm)
(ii) how to design new Monte Carlo algorithms (non-detailed balance)

Motivation

[Bernard and Krauth, PRL 2011]

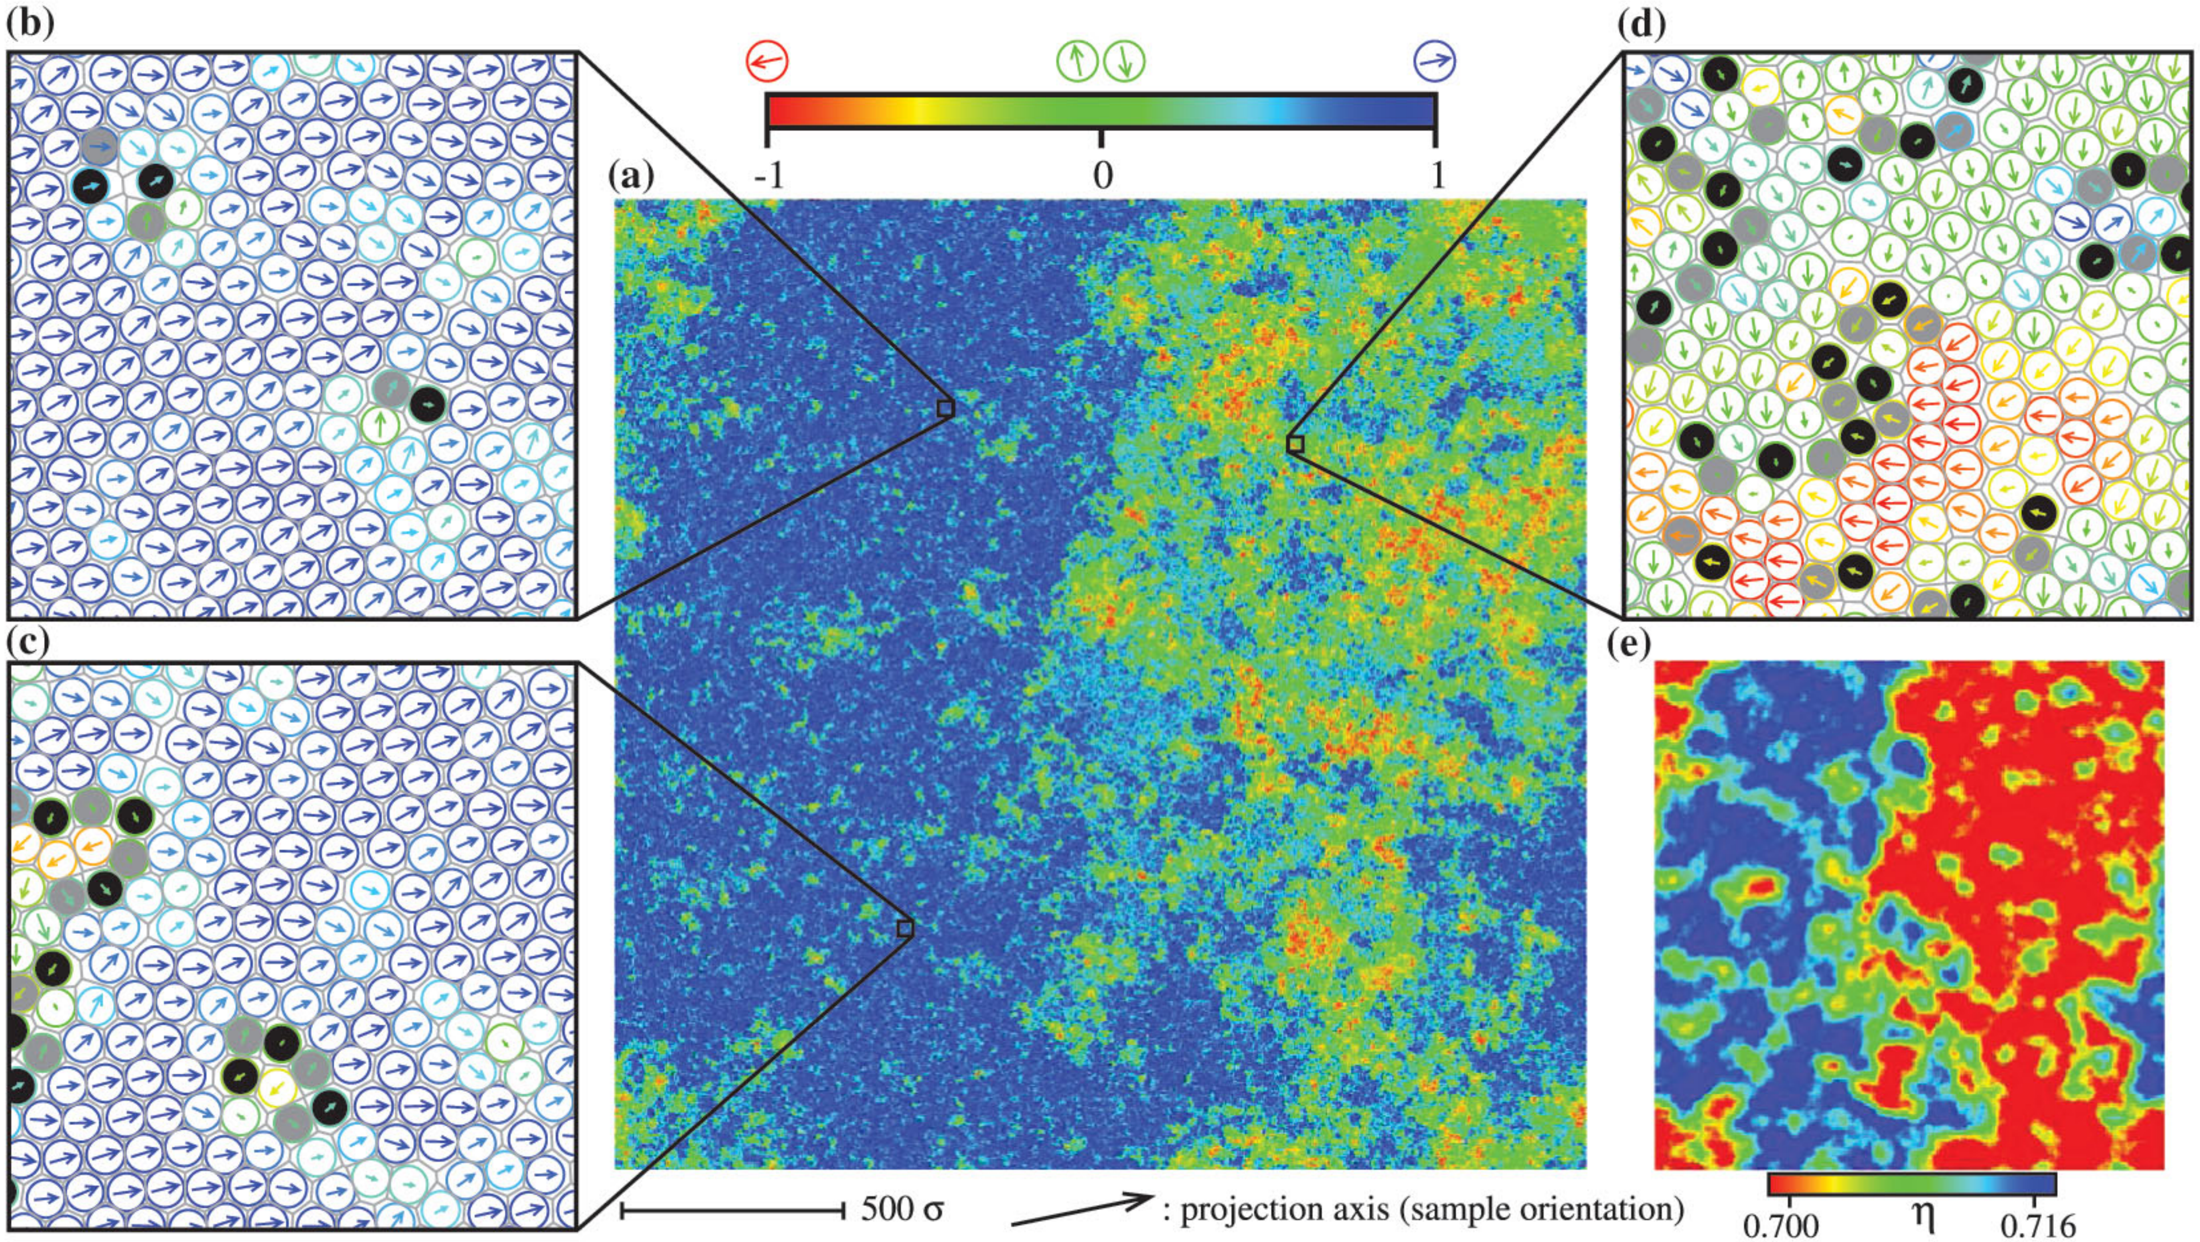

Monte Carlo algorithm *without* detailed balance allowed analysis of 10^6 hard discs, settling outstanding question as to nature of phase transition in this model.

Motivation

Claim: breaking detailed balance *generally* speeds up convergence to the steady state (and relaxation within this state)

[so MC methods without detailed balance tend to be more efficient]

Question 1: how to compare DB with non-DB in a fair way?

Question 2: how to measure convergence to equilibrium?

Questions 3,4,5… how much faster are the non-DB algorithms? what are the mechanisms of acceleration? what does this tell us about differences between eqm and non-eqm steady states?

Q1: compare DB and non-DB

Recall transition rates are $r(\mathcal{C} \to \mathcal{C}')$, steady state probabilities are $\pi(\mathcal{C})$

Define

$$r^S(\mathcal{C} \to \mathcal{C}') = \frac{1}{2}\left[r(\mathcal{C} \to \mathcal{C}') + \pi(\mathcal{C}')r(\mathcal{C}' \to \mathcal{C})\frac{1}{\pi(\mathcal{C})}\right]$$

The system with transition rates r^S obeys DB and has the same steady state as the original system.

It is “fair” to compare the systems with rates r and r^S.

Q2: rates of convergence

To compare rates of convergence...

... can always use the spectral gap

With definitions above, the spectral gap of the non-DB (irreversible) process is no smaller than the gap of the DB (reversible) process

[see eg Spiliopoulos and Rey-Bellet, J Stat Phys 2016
also Diaconis, Pavliotis, Ottobre, Hwang, Bierkens, ...]

... i.e. *non-DB converges faster*

As usual with spectral gaps, the proof is not very intuitive

... also spectral analysis tends to scale badly with N

Q2: rates of convergence

Master equation:

$$\partial_t p(\mathcal{C},t) = \sum_{\mathcal{C}'} [p(\mathcal{C}',t) r(\mathcal{C}' \to \mathcal{C}) - p(\mathcal{C},t) r(\mathcal{C} \to \mathcal{C}')]$$

Write as $\partial_t |p\rangle = \mathbb{W}|p\rangle$ with $\mathbb{W}$ the (forward) generator or master operator.

General solution:

$$p(\mathcal{C},t) = \pi(\mathcal{C}) + \sum_i y_i a_i(\mathcal{C}) \mathrm{e}^{-\lambda_i t}$$

with $\{\lambda_i\}$ the non-zero eigenvalues of $\mathbb{W}$, for which $\mathrm{Re}(\lambda_i) > 0$, also a_i is an eigenvector and y_i determined by the initial condition.

Spectral gap is $\alpha = \min_i \mathrm{Re}(-\lambda_i)$
Large gap implies fast convergence to equilibrium

Comparing DB and non-DB, can show faster convergence without DB:

$$\alpha \geq \alpha^S$$

[see eg Spiliopoulos and Rey-Bellet, J Stat Phys 2016]

... larger spectral gap

Quick proof:

Suppose that $\langle z|\mathbb{W} = -\lambda\langle z|$ with $\lambda \neq 0$.

Then

$$-\lambda = \frac{\langle z|\mathbb{W}\hat{\pi}|z\rangle}{\langle z|\hat{\pi}|z\rangle}, \qquad \hat{\pi} = \sum_{\mathcal{C}} |\mathcal{C}\rangle\pi(\mathcal{C})\langle\mathcal{C}|$$

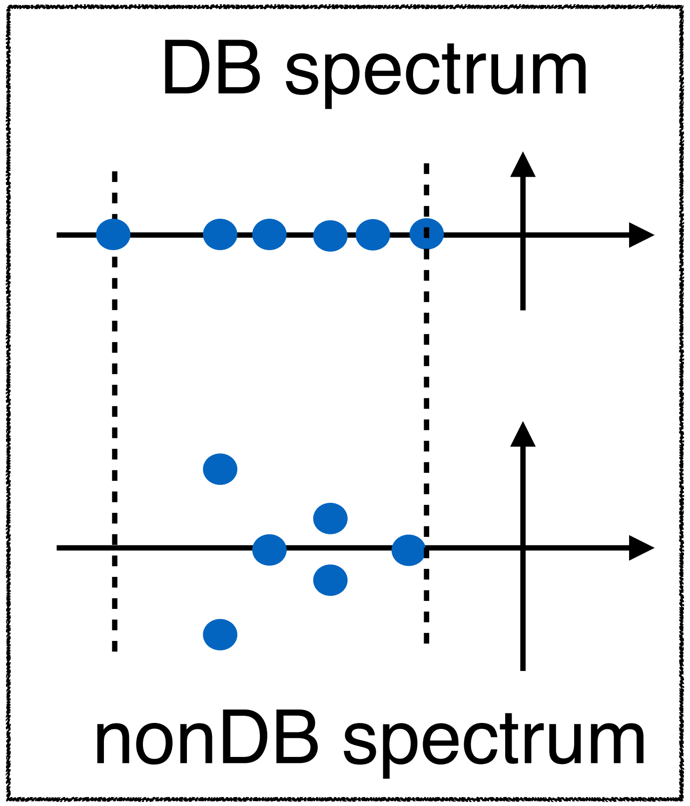

Take real part and use definition of $\mathbb{W}^S$

$$\mathrm{Re}(\lambda) = \frac{-\langle z|\mathbb{W}\hat{\pi} + (\mathbb{W}\hat{\pi})^\dagger|z\rangle}{2\langle z|\hat{\pi}|z\rangle} = \frac{-\langle z|\mathbb{W}^S\hat{\pi}|z\rangle}{\langle z|\hat{\pi}|z\rangle}$$

Then, by a similar argument to Rayleigh-Ritz,

$$\alpha^S \leq \mathrm{Re}(\lambda) \leq \lambda^S_{\mathrm{max}}$$

where we recall α^S is the spectral gap of $\mathbb{W}^S$ (smallest non-zero eigenvalue of $-\mathbb{W}^S$) and λ^S_{max} is the largest eigenvalue of $-\mathbb{W}^S$

So the eigenvalues of the nonDB process have real parts "inside" the spectrum of the DB process.

[see eg Kaiser, RLJ and Zimmer, J Stat Phys 2017]

Q2: rates of convergence

[see eg Spiliopoulos and Rey-Bellet, J Stat Phys 2014 & 2016 ;
Bierkens, Stat Comput, 2015]

Alternative characterisation of convergence to steady state:

Consider a trajectory over a long time $\mathcal{T}$, let $\mu^{\mathcal{T}}(R)$ be the fraction of time spent in configuration R.

As $\mathcal{T} \to \infty$, convergence to equilbrium means that $\mu^{\mathcal{T}}(R) \to \pi(R)$

Large deviations:

$$\mathrm{Prob}(\mu^{\mathcal{T}} \approx \rho) \asymp \mathrm{e}^{-\mathcal{T} I_2(\rho)}, \qquad I_2(\pi) = 0$$

I_2 is (literally) a rate function. . . the larger is $I_2(\rho)$, the more rapid is the convergence

I_2 determines the probability that we see a distribution ρ instead of the expected distribution π.

Comparing DB and non-DB, can show faster convergence without DB:

$$I_2(\rho) \geq I_2^S(\rho)$$

Why does this happen? (Q3)

Breaking detailed balance increases the spectral gap $\alpha \geq \alpha^S$ and reduces the probability of observing non-typical distributions $I_2 \geq I_2^S$.

The proofs of these statements are not very intuitive...

What is going on?

To get an idea, we consider *Macroscopic Fluctuation Theory*

[see eg Bertini, De Sole, Gabrielli, Jona-Lasinio, Landim, Rev Mod Phys 2015]

Consider a system of many particles hopping in real space; zoom out so that the individual particles are not visible any more; expect the density ρ to obey a diffusion-type equation

$$\partial_t \rho = -\nabla \cdot J(\rho), \qquad J(\rho) = -D\nabla\rho + \chi E$$

with D being a diffusion constant, χ the mobility, and E an external field (which need not be conservative)

Macroscopic FT example

Weakly asymmetric exclusion process: particles hopping on a lattice, with at most one particle per site, and periodic boundaries,

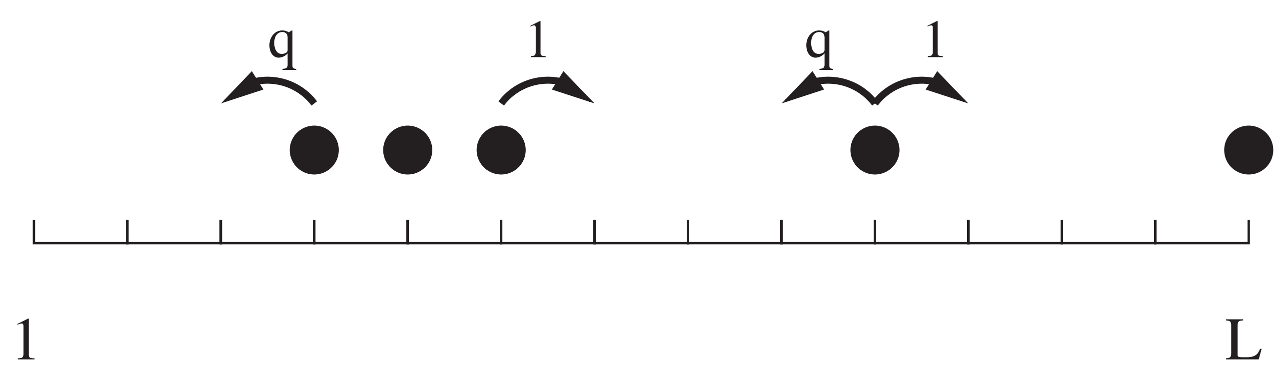

Take $L \to \infty$ at fixed density but $q = 1 - E/L$.

In this case

$$J(\rho) = -\nabla\rho + \chi E, \qquad \chi(x) = 2\rho(x)[1-\rho(x)].$$

For fixed L, the original results for spectral gap and I_2 still apply
(The DB case is $E = 0$)

What happens for large L?

Macroscopic FT

Let $\rho^{\mathcal{T}}$ be the time-averaged density and $j^{\mathcal{T}}$ the time-averaged current.
As $\mathcal{T} \to \infty$, expect $j^{\mathcal{T}} \to J(\rho^{\mathcal{T}})$ and $\rho^{\mathcal{T}} \to \overline{\rho}$, the steady-state density

In fact, for $\nabla \cdot j = 0$ and assuming no dynamical phase transitions,

$$\mathrm{Prob}\left[(\rho^{\mathcal{T}}, j^{\mathcal{T}}) \approx (\rho, j))\right] \sim \exp\left(-\mathcal{T} \underbrace{\int (j - J(\rho)) \frac{1}{4\chi} (j - J(\rho)) \mathrm{d}x}_{\text{"Level-2.5 rate function" } I_{2.5}(\rho, j)}\right)$$

There is an analogue of I_2 in this system, which is

$$I_2(\rho) = \inf_{j : \nabla \cdot j = 0} I_{2.5}(\rho, j)$$

$I_2(\rho)$ sets the probability of observing a non-typical density $\rho^{\mathcal{T}} \neq \overline{\rho}$.

We will show that $I_2(\rho) \geq I_2^S(\rho)$, so breaking detailed balance accelerates convergence (again)

Splitting the current

[see eg Bertini, De Sole, Gabrielli, Jona-Lasinio, Landim, Rev Mod Phys 2015]

Within MFT, it is known that we can decompose

$$J(\rho) = J_S(\rho) + J_A(\rho)$$

where J_S is the current for the reversible (DB) model and J_A is the remainder.

Breaking DB corresponds to introducing a finite J^A;
we compare always models with fixed J^S.

Moreover, the two currents are orthogonal, in the sense that

$$\int J_A \cdot \chi^{-1} J_S \,\mathrm{d}x = 0$$

(Note this is orthogonality in the sense of fields, the currents at any given point in space are not orthogonal. . .)

Also J_S follows the free-energy gradient, that is $J_S = -\chi \nabla \frac{\delta \mathcal{F}}{\delta \rho}$

A picture…

[Kaiser, RLJ and Zimmer, J Stat Phys 2017]

Consider the space of density profiles, and plot contours of the probability

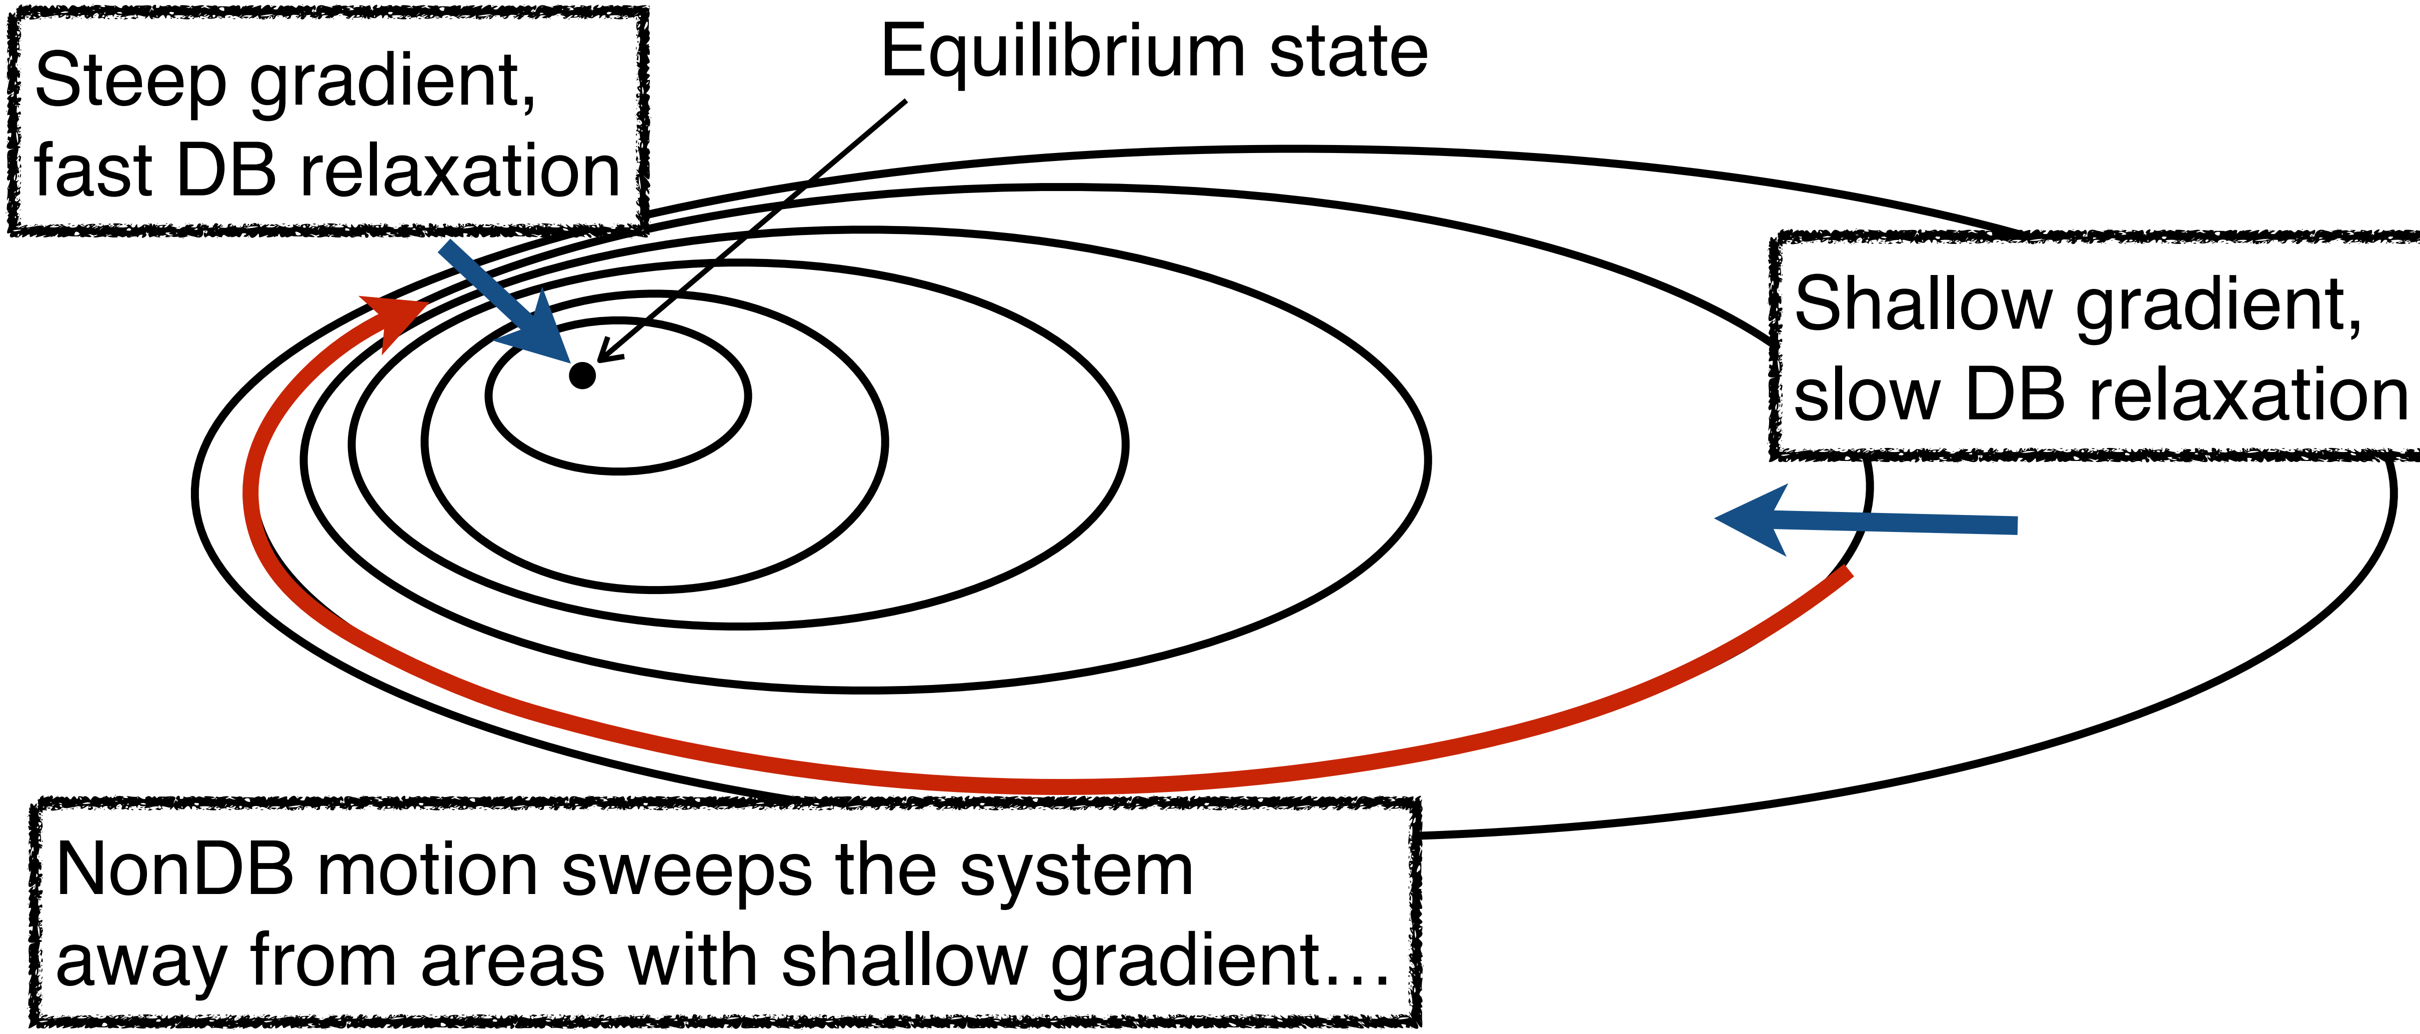

For this to make sense, we need a *metric* structure on the space of of density profiles, this is a (modified) *Wasserstein metric*

[see eg RLJ and Zimmer, J Phys A 2014]

Of course, nonDB motion can also sweep us away from steep regions, but on average, we still go faster towards the steady state

Faster convergence: a bound on I_2

[Kaiser, RLJ and Zimmer, J Stat Phys 2017]

Recall: “level 2.5”

$$\mathrm{Prob}\left[(\rho^{\mathcal{T}}, j^{\mathcal{T}}) \approx (\rho, j))\right] \sim \exp\left(-\mathcal{T} I_{2.5}(\rho, j)\right),$$

$$I_{2.5}(\rho, j) = \int (j - J(\rho)) \frac{1}{4\chi(\rho)} (j - J(\rho)) \mathrm{d}x$$

Currents:

$$J(\rho) = J_S(\rho) + J_A(\rho), \qquad \int J_S \cdot \frac{1}{\chi} J_A \, \mathrm{d}x = 0, \qquad J_S = -\chi \nabla \frac{\delta \mathcal{F}}{\delta \rho}$$

Show that:

$$I_2(\rho) \geq I_2^S(\rho), \qquad I_2(\rho) = \inf_{j : \mathrm{div}\, j = 0} I_{2.5}(\rho, j)$$

Bounds on I_2

[Kaiser, RLJ and Zimmer, J Stat Phys 2017]

$$I^S_{2.5} = \int (j - J_S)\cdot \frac{1}{4\chi}(j - J_S)\,\mathrm{d}x$$

$$= \int j\cdot\frac{1}{4\chi}j + J_S\cdot\frac{1}{4\chi}J_S - \boxed{2j\cdot\frac{1}{4\chi}J_S}\,\mathrm{d}x$$

Vanishes by gradient structure, div j = 0

Hence $I^S_2 = \int J_S\cdot\frac{1}{\chi}J_S\,\mathrm{d}x$

$$I_{2.5} = \int (j - J_S - J_A)\cdot\frac{1}{4\chi}(j - J_S - J_A)\,\mathrm{d}x$$

$$= \int J_S\cdot\frac{1}{4\chi}J_S - \boxed{2j\cdot\frac{1}{4\chi}J_S} - \boxed{2J_A\cdot\frac{1}{4\chi}J_S} + (j - J_A)\cdot\frac{1}{4\chi}(j - J_A)\,\mathrm{d}x$$

vanishes

Vanishes by orthogonality

$$I_{2.5} = I^S_2 + \int (j - J_A)\cdot\frac{1}{4\chi}(j - J_A)\,\mathrm{d}x$$

Hence $I_2 \geq I^S_2$

Can also get an exact formula for $I_2 - I^S_2$, but we need to solve a PDE.

Bounds on I_2 : optimal control

[Kaiser, RLJ and Zimmer, J Stat Phys 2017, also arXiv:1708.01453]

To compute $I_2(\rho)\ldots$
add a force $F_\rho = \nabla\phi$ to the system, that makes the density ρ typical.
(auxiliary process / driven process / optimally-controlled process)

[Doob, Fleming, Evans, Maes, RLJ-Sollich, Touchette-Chetrite, ...]

The current associated to this force is $\chi(\rho)\nabla\phi$, and ϕ solves

$$\nabla \cdot \underbrace{(J^A(\rho) - \chi(\rho)\nabla\phi)}_{\text{(steady-state current } \tilde{J}_{\rm ss} \text{ of the optimally controlled process)}} = 0$$

Then finally...

$$\begin{aligned} I_2(\rho) &= \inf_{j:\mathrm{div}\, j=0} I_{2.5}(\rho, j) \\ &= I_2^S(\rho) + \int_\Lambda (J^A - \tilde{J}_{\rm ss}) \cdot \frac{1}{4\chi}(J^A - \tilde{J}_{\rm ss}) \\ &= I_2^S(\rho) + \int_\Lambda \nabla\phi \cdot \frac{\chi}{4}\nabla\phi \end{aligned}$$

A picture…

[Kaiser, RLJ and Zimmer, J Stat Phys 2017]

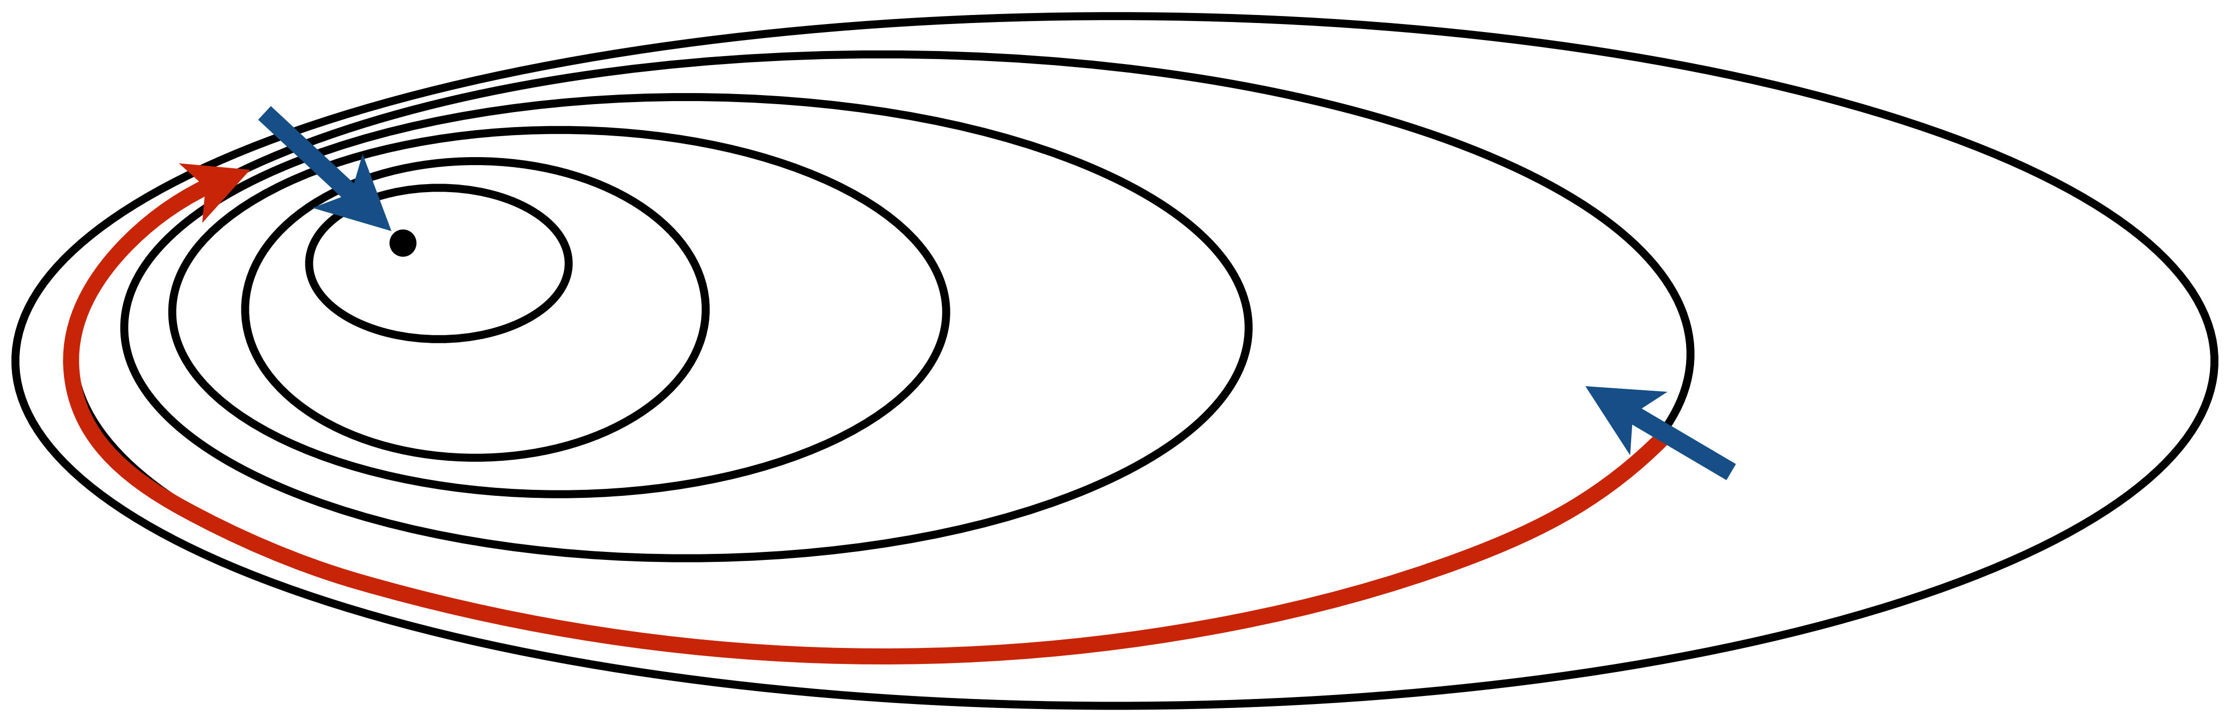

... contours of the free energy / quasipotential

Probability cost to stay at ρ requires that we resist both gradient and orthogonal forces

$$I_2(\rho) = I_2^S(\rho) + \int_\Lambda \nabla\phi \cdot \frac{\chi}{4} \nabla\phi$$

Numerical example

[Kaiser, RLJ and Zimmer, J Stat Phys 2017]

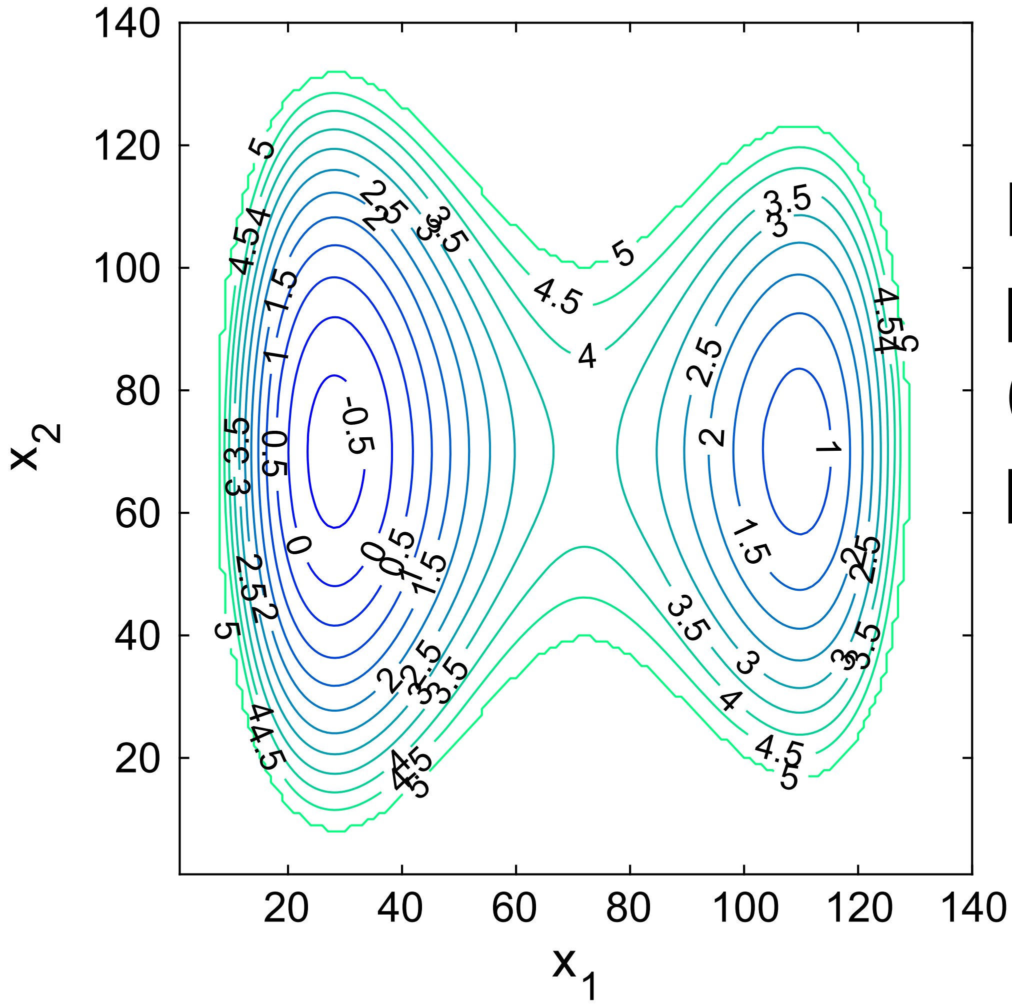

Non-interacting particles in 2d (also zero-range processes)

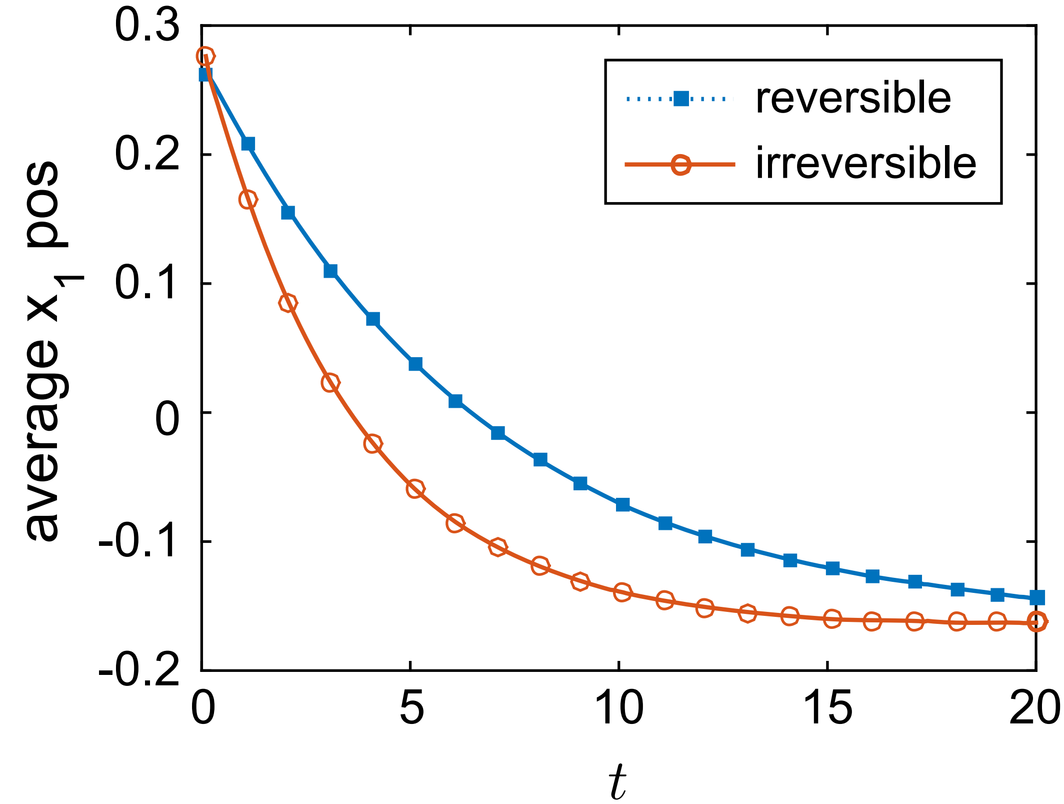

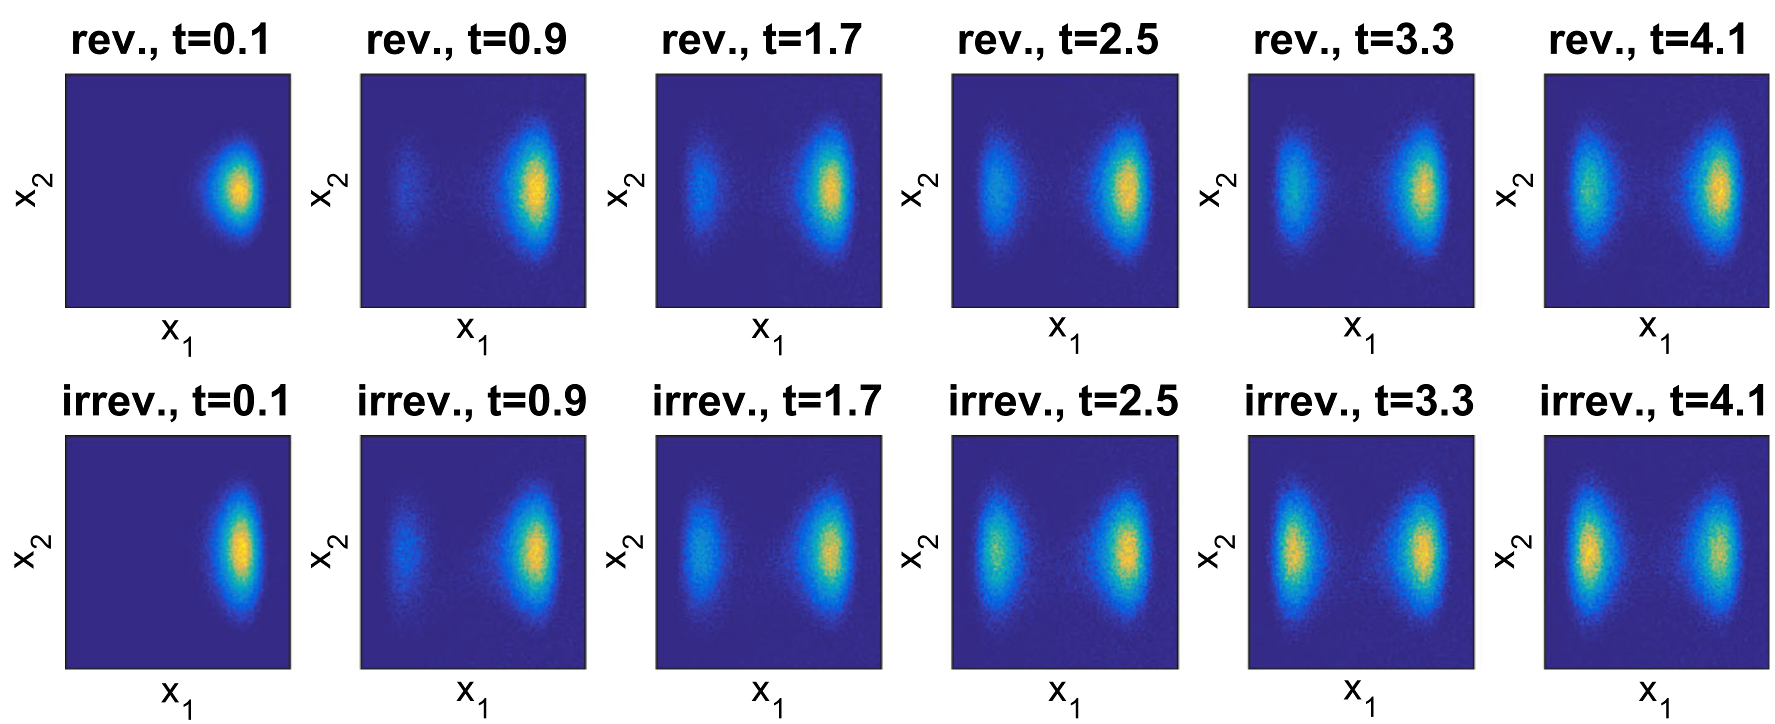

Mild speedup (~ 2x)

General Markov chains

So far we worked on the hydrodynamic scale, but speedup by breaking DB is a general result for Markov chains

Question: Do these "geometrical" results have counterparts for the general case?

Answer: sort of...

Within MFT we have a quadratic rate function (squared norm)

$$I_{2.5}(\rho, j) = \int (j - J(\rho)) \frac{1}{4\chi(\rho)} (j - J(\rho)) \mathrm{d}x = \tfrac{1}{4} \| j - J(\rho) \|^2_{\chi^{-1}}$$

+ gradient + orthogonality

$$J_S = -\chi \nabla \frac{\delta \mathcal{F}}{\delta \rho}, \qquad 0 = \int J_S \cdot \frac{1}{\chi} J_A \, \mathrm{d}x = \langle J_S, J_A \rangle_{\chi^{-1}}$$

For general Markov chains, the algebraic structure remains but the geometry is lost ("generalised gradients", "generalised orthogonality")

[Kaiser, RLJ and Zimmer, arXiv:1708.01453]

Geometry of irreversible Markov chains

[Kaiser, RLJ and Zimmer, arXiv:1708.01453; see Marcus' talk last week]

Write MFT in terms of forces, let $J = \chi F$

$$I_{2.5} = \tfrac{1}{4}\|j - \chi F\|^2_{\chi^{-1}}, \quad F^S = -\nabla \tfrac{\delta \mathcal{F}}{\delta \rho}, \quad \|F^S + F^A\|^2_\chi = \|F^S - F^A\|^2_\chi,$$

$$\Rightarrow \langle F^S, F^A \rangle_\chi = 0$$

$$I_2(\rho) = \tfrac{1}{4}\left\|J^S\right\|^2_{\chi^{-1}} + \tfrac{1}{4}\left\|\tilde{J}_{\mathrm{ss}} - \chi F^A\right\|^2_{\chi^{-1}}$$

Markov chains: forces are $F^{\mathcal{C},\mathcal{C}'}(\rho) = \log \frac{\rho(\mathcal{C}) r(\mathcal{C}\to\mathcal{C}')}{\rho(\mathcal{C}') r(\mathcal{C}'\to\mathcal{C})}$

$I_{2.5} = \Phi(\rho, j, F)$, non-negative, convex in j and F

$F = F^S + F^A$ with $F^S = -\nabla^{\mathcal{C},\mathcal{C}'} \frac{\delta \mathcal{F}}{\delta \rho}$, discrete gradient of the free energy

$\Psi^\star(\rho, F) > 0$ is symmetric and convex in F, analogous to $\|F\|^2_\chi$

$$\Psi^\star(\rho, F^S + F^A) = \Psi^\star(\rho, F^S - F^A), \text{ “generalised orthogonality”}$$

$\ldots I_2(\rho) = \Phi^S(\rho, 0, F^S) + \Phi(\rho, \tilde{J}_{\mathrm{ss}}, F^A)$, $\tilde{J}_{\mathrm{ss}}$ comes from optimal control

Outlook / summary

Irreversible (nonDB) samplers seem to be efficient, they can push us away from areas where the energy gradient is shallow... (big effect?)

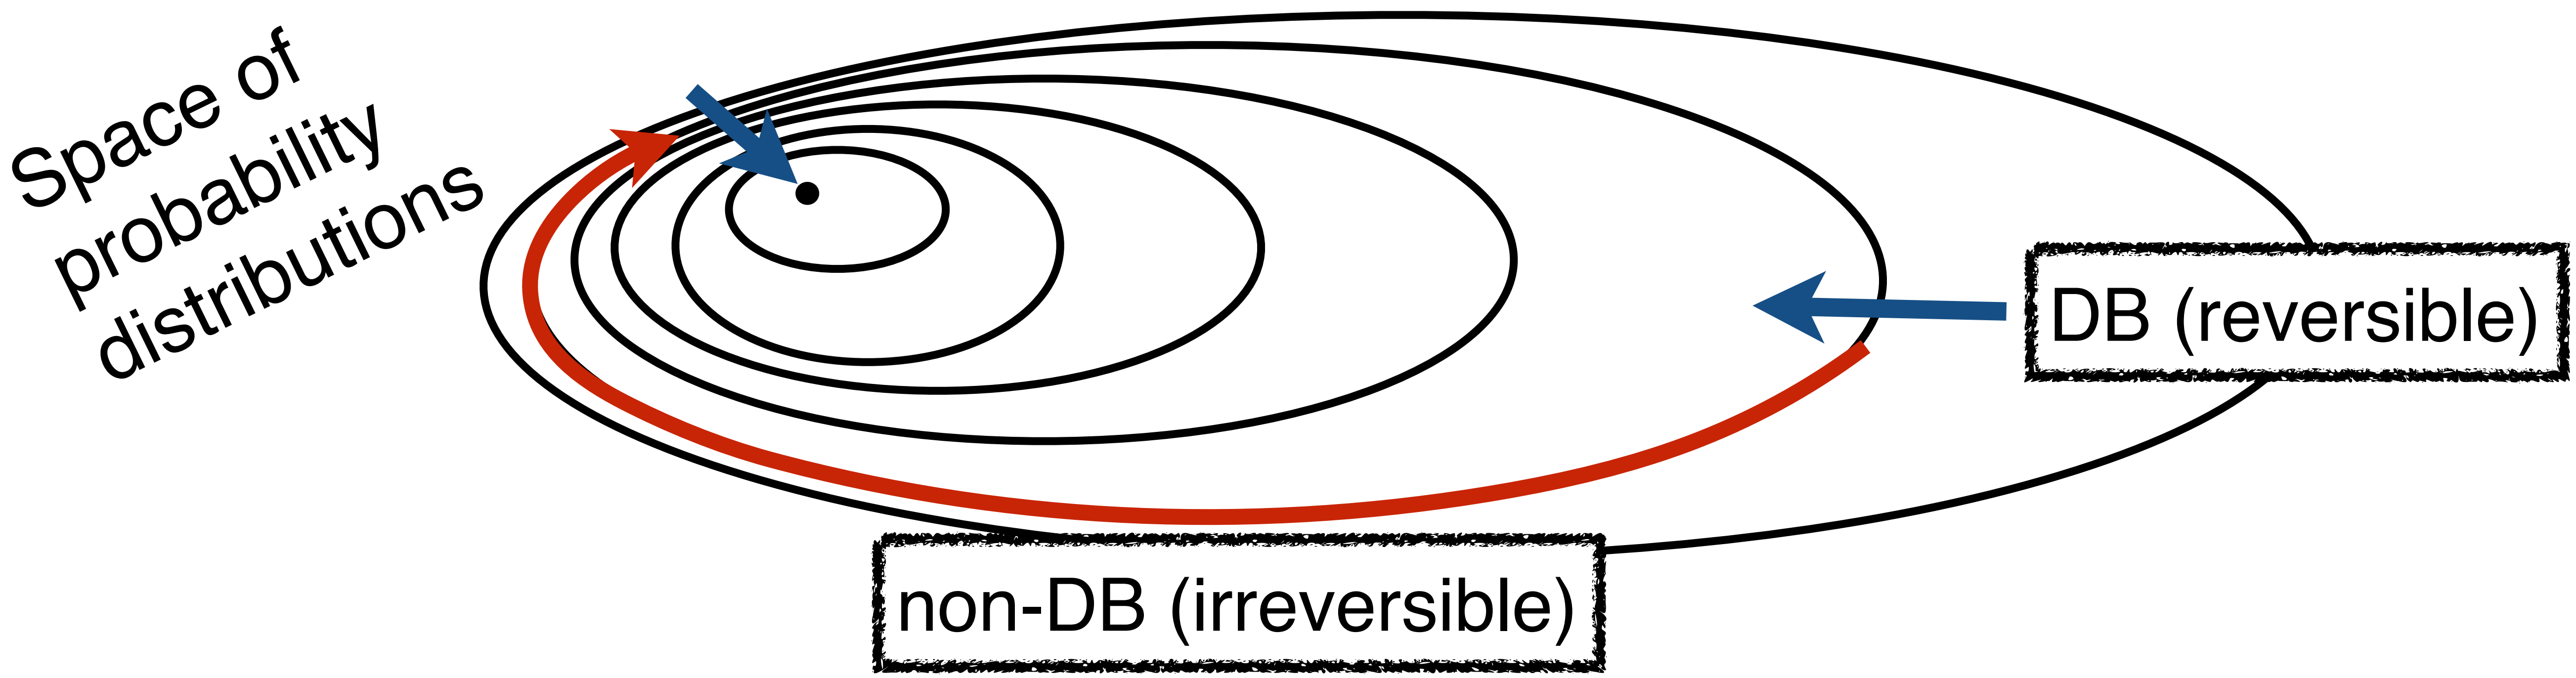

Large deviation rate functions can help to characterise sampling efficiency

The efficiency gain can be computed using ("generalised") orthogonality of forces and/or currents

Are these geometrical ideas useful for more general non-eqm questions?

Final advertisement

[RLJ, Kaiser, and Zimmer, arXiv:1709.04771]

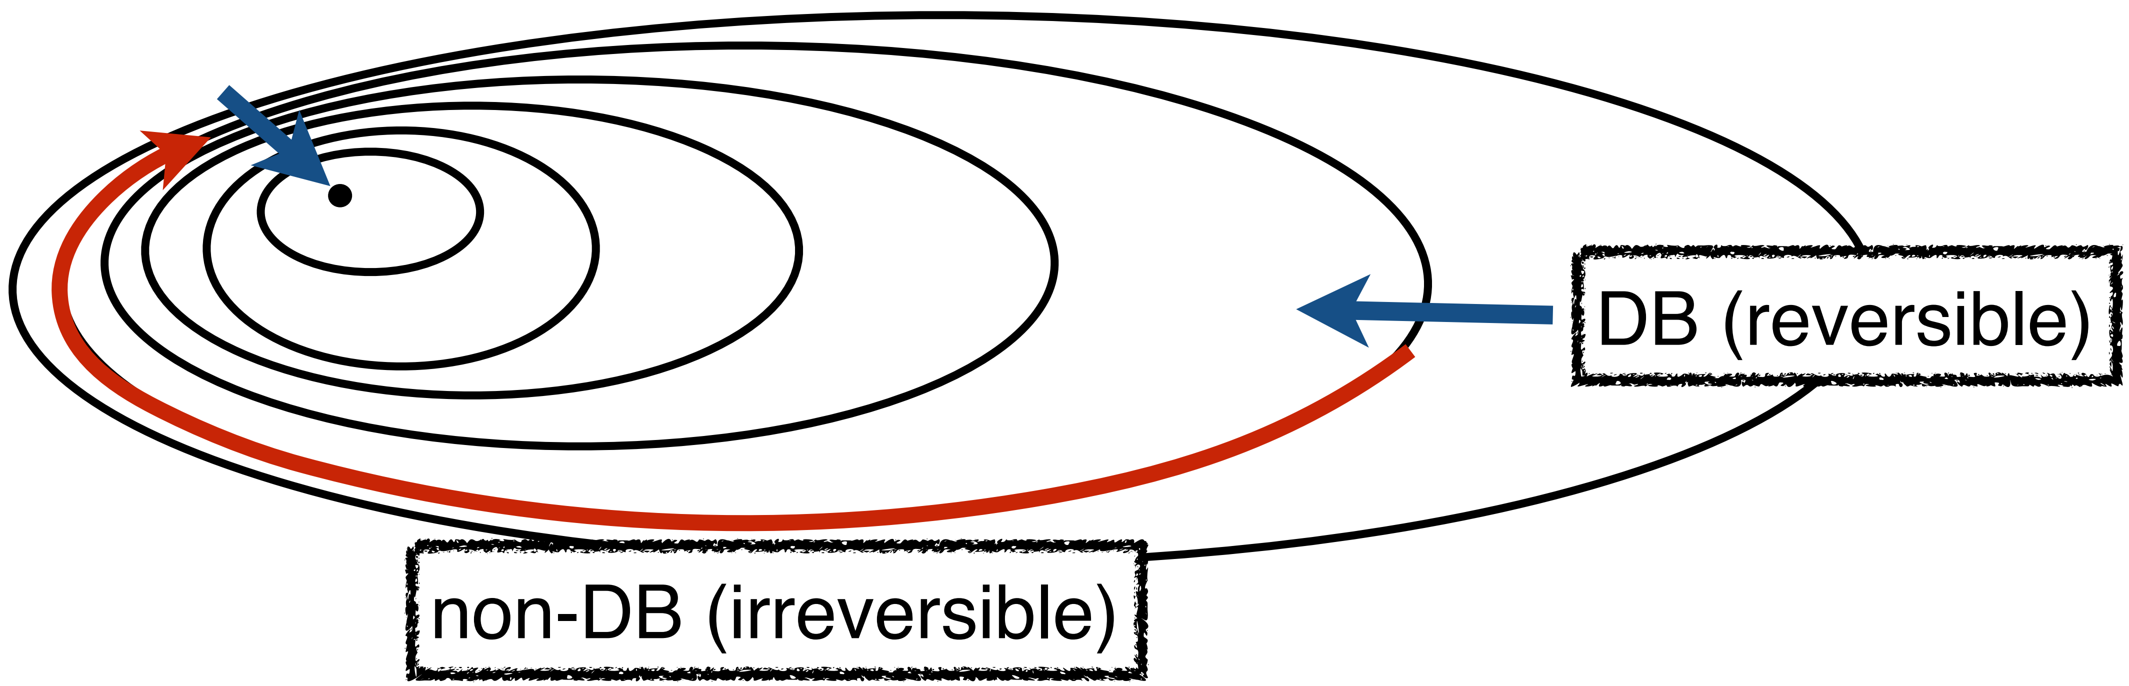

Question: does the gradient / orthogonal structure of MFT survive if we have "second order Langevin" (positions and momenta with thermal noise only for momenta)?

Answer: partially...

(note also, J_A controls steady-state entropy production)