Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

The colors of water

A tutorial on the Generalized Langevin Thermostat

Michele Ceriotti

Lugano, 09/03/2010

Outline

- Canonical sampling:
  - Sampling efficiency: why, when and how
  - NVT molecular dynamics local and global schemes
  - The (generalized) Langevin equation
- A case study: liquid, flexible water
  - Optimal sampling
  - Preserving dynamical properties
  - The making of a GLE thermostat
- Other applications of GLE thermo: call for implementers...

Molecular Dynamics

- Modelling of the dynamics of a system by reproducing the motion of the atoms

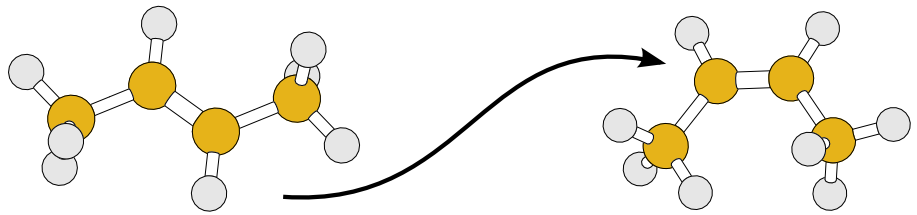

- Numerical integration of Hamilton's equations:

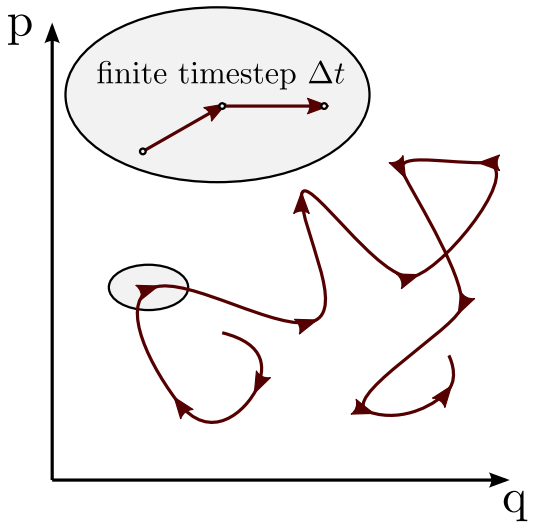

$$\frac{\partial q}{\partial t} = \frac{\partial H}{\partial p} = \frac{p}{m}$$

$$\frac{\partial p}{\partial t} = -\frac{\partial H}{\partial q} = -\frac{\partial V}{\partial q}$$

- Must modify to sample **canonical** ensemble

Ergodic sampling

- Ergodic hypothesis: equivalence between ensemble averages and time averages along a trajectory

$$\langle A \rangle = \int \mathrm{d}\boldsymbol{p}\mathrm{d}\boldsymbol{q} A\left(\mathbf{q},\mathbf{p}\right) e^{-\beta H(\mathbf{q},\mathbf{p})} = \lim_{T\to\infty} \frac{1}{T} \int_0^T A\left(\mathbf{q}\left(t\right),\mathbf{p}\left(t\right)\right) \mathrm{d}t$$

  - Points along the trajectory must be distributed based on $e^{-\beta H(\mathbf{q},\mathbf{p})}$ $\Leftrightarrow$ fluctuation-dissipation theorem/detailed balance

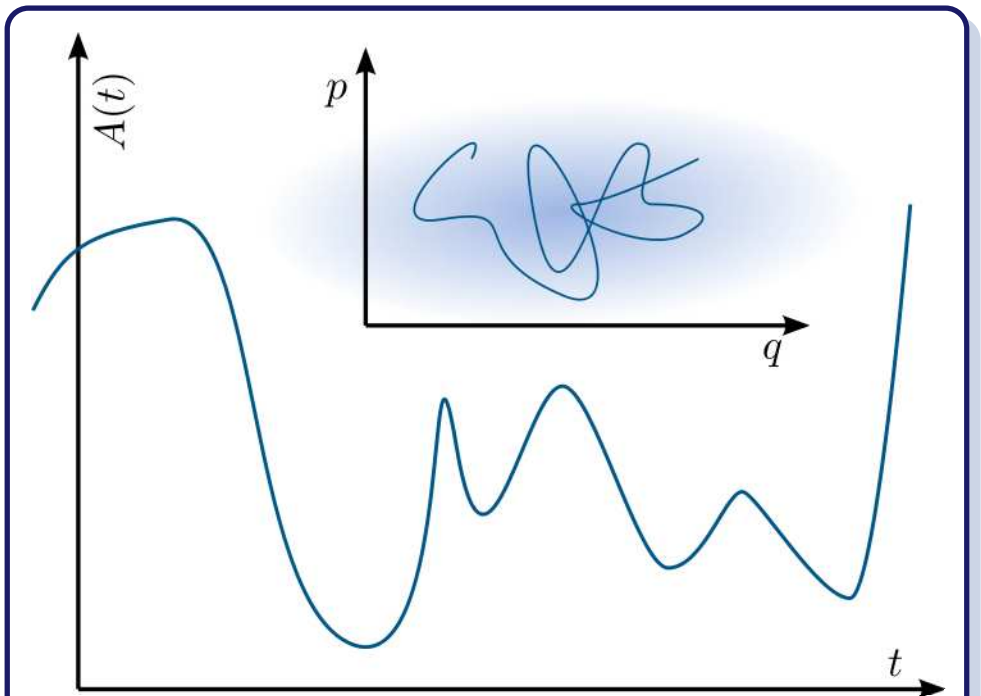

Measuring ergodicity

- The error on averages decreases with the square root of the number of **uncorrelated** samples

- Sampling more often than the correlation time does not improve convergency

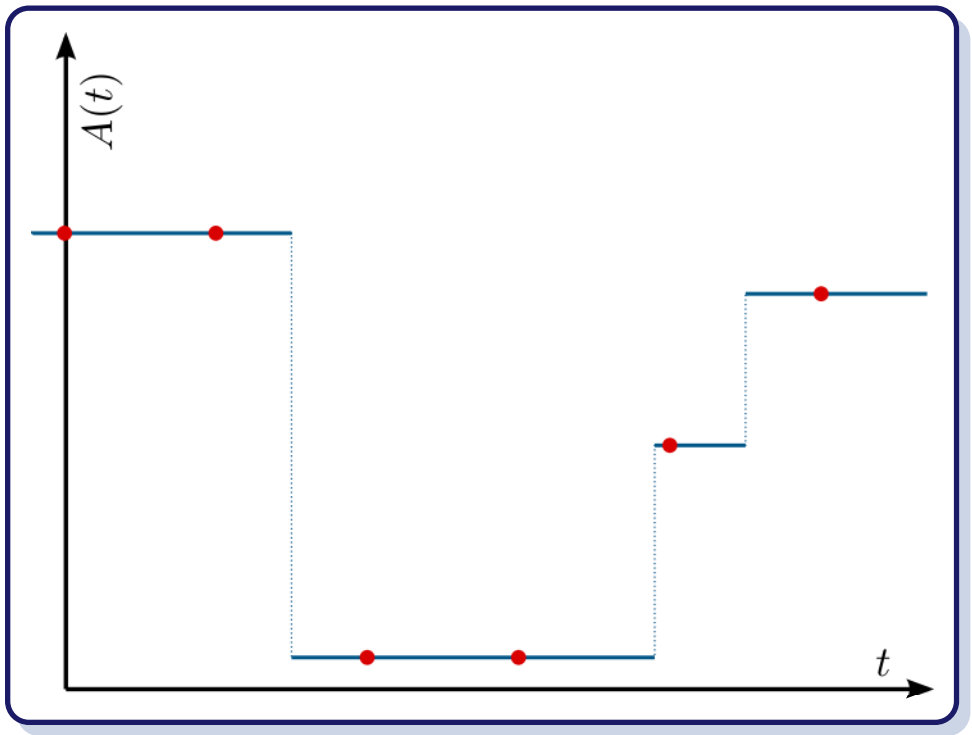

Measuring ergodicity

- The error on averages decreases with the square root of the number of **uncorrelated** samples

- Sampling more often than the correlation time does not improve convergency

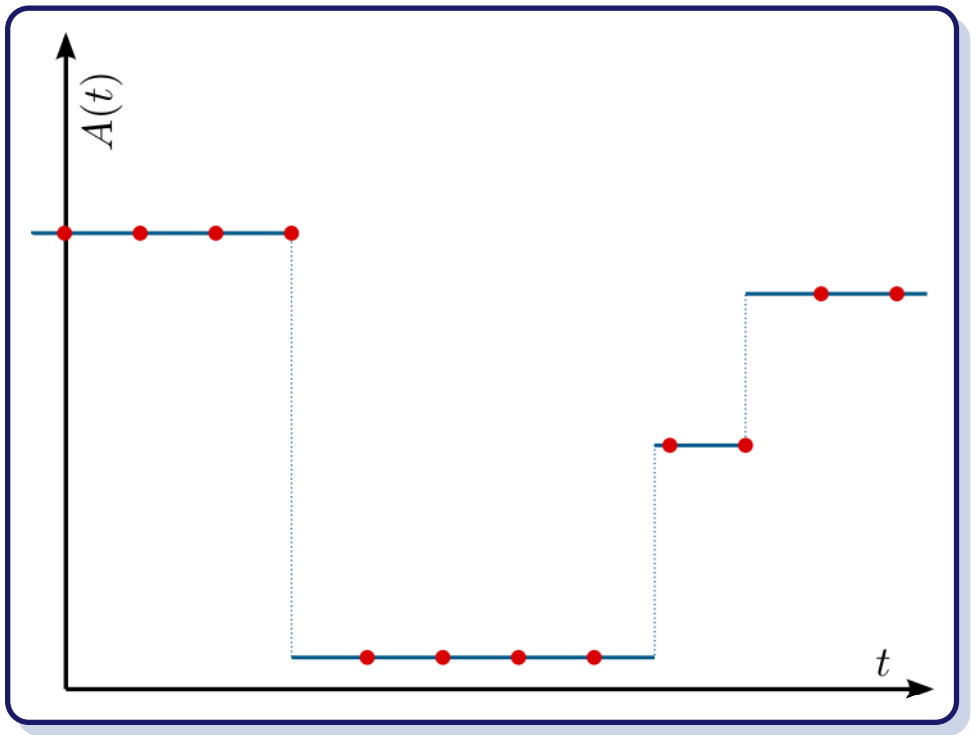

Measuring ergodicity

- The error on averages decreases with the square root of the number of **uncorrelated** samples

- Sampling more often than the correlation time does not improve convergency

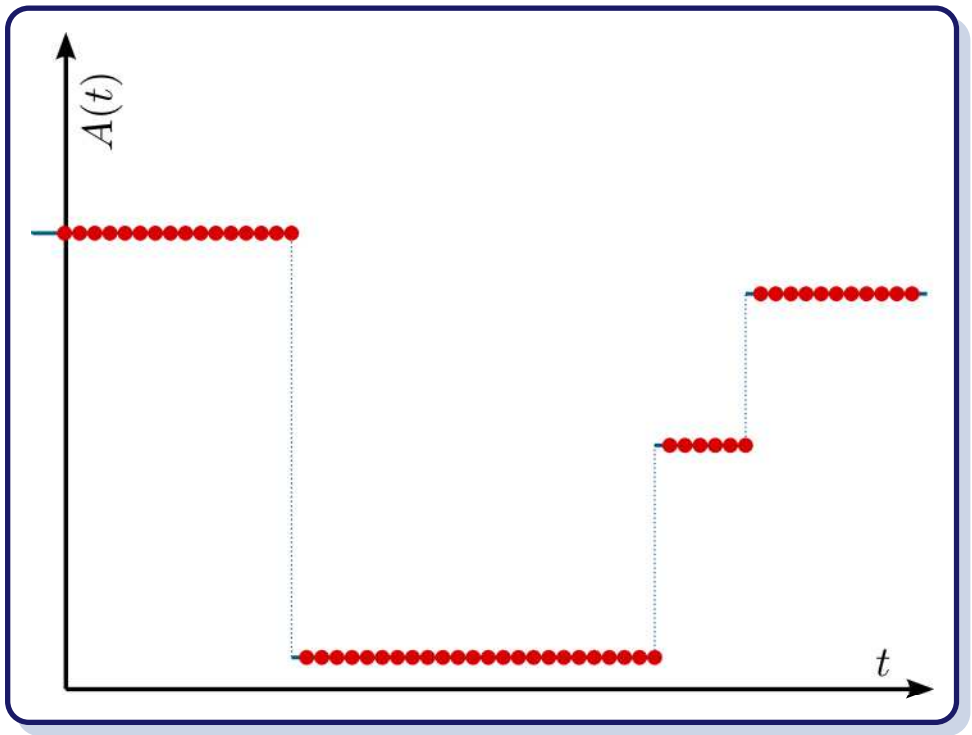

Measuring ergodicity

- The error on averages decreases with the square root of the number of **uncorrelated** samples

- Autocorrelation function:

$$\langle A(t) A(0) \rangle = \frac{\frac{1}{T} \int_0^T (A(s+t) - \langle A \rangle)(A(s) - \langle A \rangle) \, \mathrm{d}s}{\langle A^2 \rangle - \langle A \rangle^2}$$

Measuring ergodicity

- The error on averages decreases with the square root of the number of **uncorrelated** samples

- Autocorrelation function:

$$\langle A(t) A(0) \rangle = \frac{\frac{1}{T} \int_0^T (A(s+t) - \langle A \rangle)(A(s) - \langle A \rangle) \, \mathrm{d}s}{\langle A^2 \rangle - \langle A \rangle^2}$$

  - Will it rain as much as today?

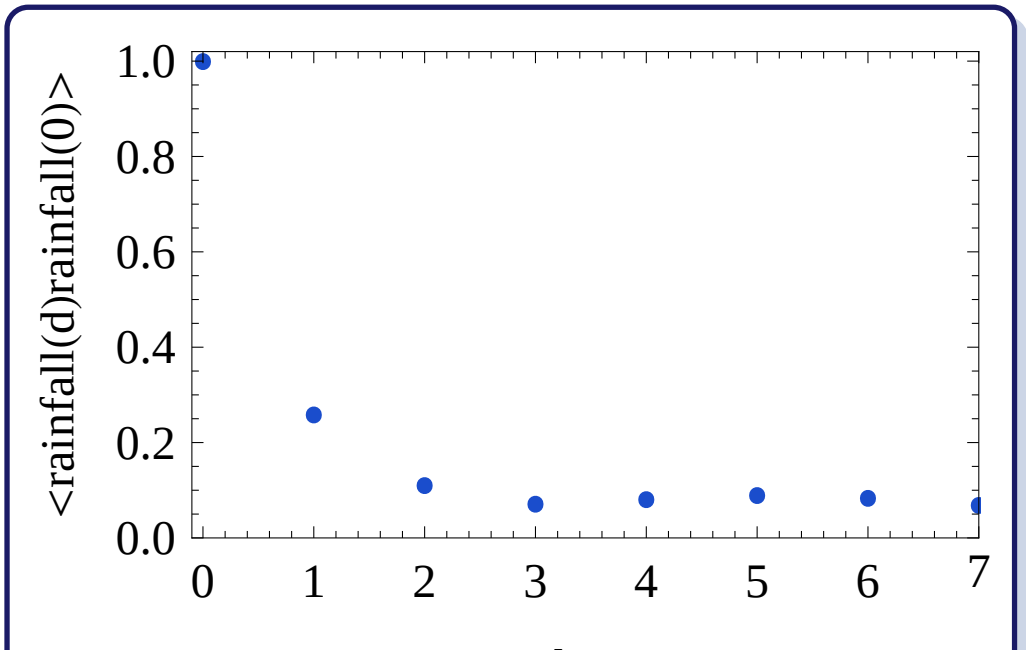

Measuring ergodicity

- The error on averages decreases with the square root of the number of **uncorrelated** samples

- Autocorrelation function:

$$\langle A(t) A(0) \rangle = \frac{\frac{1}{T} \int_0^T (A(s+t) - \langle A \rangle)(A(s) - \langle A \rangle) \, \mathrm{d}s}{\langle A^2 \rangle - \langle A \rangle^2}$$

- Averages over a time interval T will be affected by an error which decreases as $1/\sqrt{T/2\tau}$

$$\tau = \int_0^\infty \langle A(t) A(0) \rangle \, \mathrm{d}t$$

- Computing autocorrelation functions is hard: must sample for hundreds of times τ!

Time scales

- The evolution of a non-trivial system is a combination of fast and slow components
- Correspondingly, the autocorrelation function shows different time scales

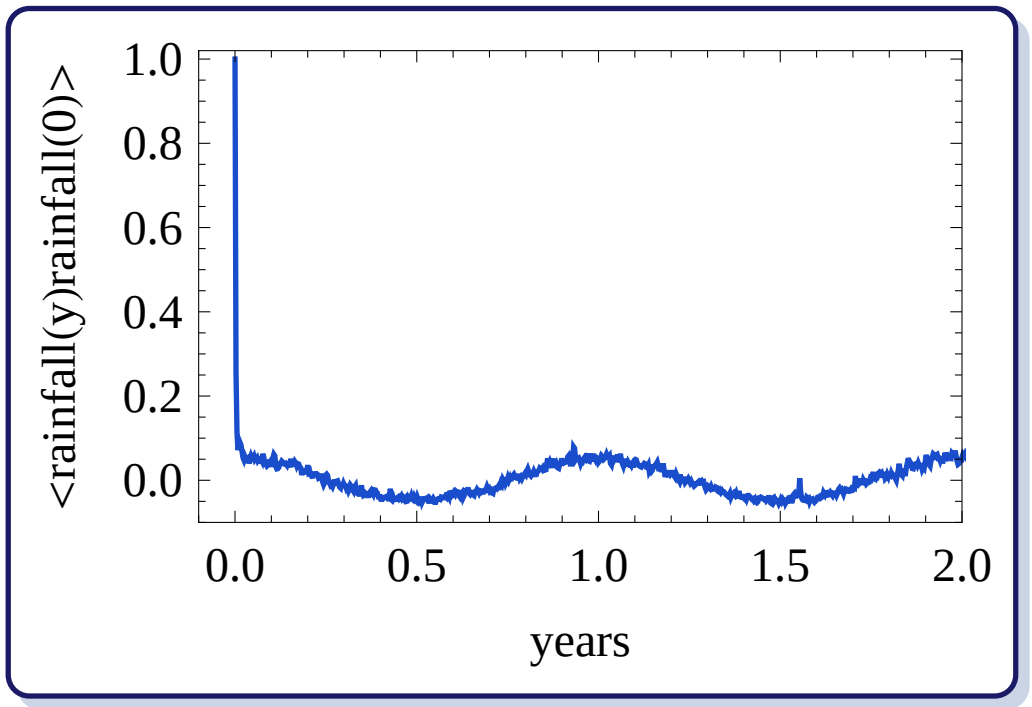

seasonal dependence of rainfall

Time scales

- The evolution of a non-trivial system is a combination of fast and slow components
- Correspondingly, the autocorrelation function shows different time scales
- Different observables may have different time scales

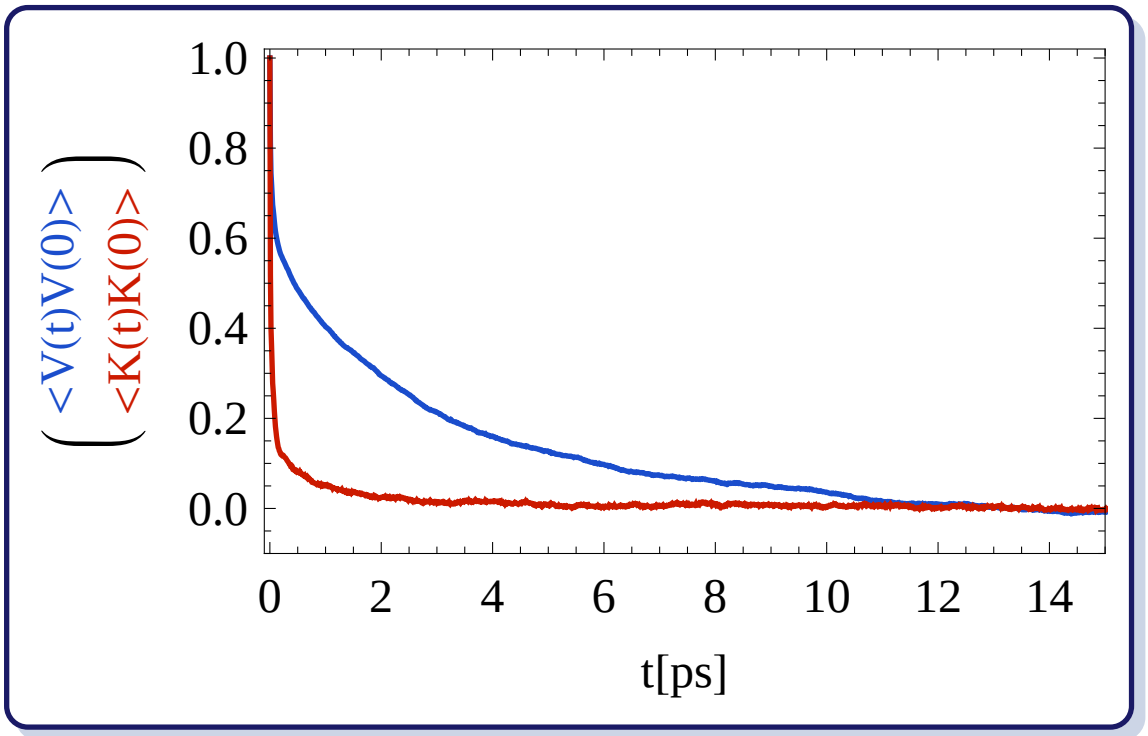

potential and kinetic energy in liquid water

Constant-Temperature MD

- Relax isolated-system hypothesis: closer to "real life" than microcanonical MD for small systems
- One must modify Hamilton's equations (Andersen, Langevin, Nosé-Hoover, stochastic rescaling...)
  - Mimick the effect of a heat bath (open system, total energy fluctuates)
  - Can we define a conserved quantity (useful to check timestep)?
- The canonical ensemble ($P(p,q) \propto e^{-\beta\left[\frac{p^2}{2m}+V(q)\right]}$) is sampled:
  - Initial equilibration (bring the system quickly to temperature)
  - **Dynamical properties are altered**
  - Efficient **sampling of static properties** (how to improve ergodicity?)

Local and global thermostats

- A **global** thermostat enforces the distribution of the **total** kinetic energy

$$P(K)\,\mathrm{d}K \propto K^{\left(N_f/2-1\right)} e^{-K/k_B T}\,\mathrm{d}K$$

  - Little disturbance on the dynamics, relies on internal couplings

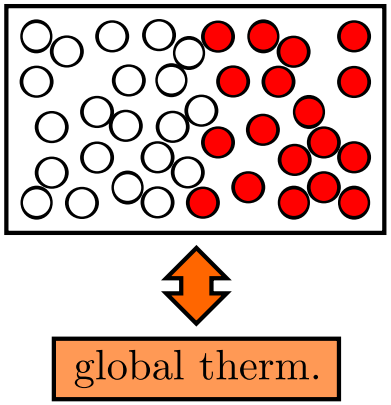

- A **local** enforces canonical distribution of individual degrees of freedom

$$P(p_i)\,\mathrm{d}p_i \propto e^{-p_i^2/2mk_B T}\,\mathrm{d}p_i$$

  - Greater disturbance, actively counteracts local imbalance

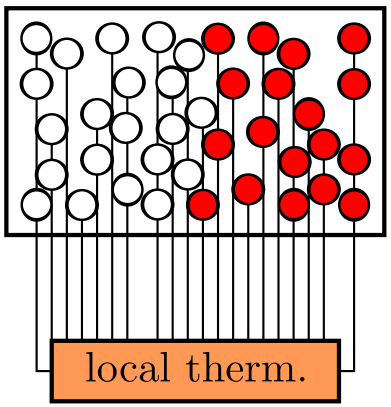

Langevin Dynamics

- A linear, Markovian stochastic equation for the momenta

$$\dot{p}(t) = -\gamma p(t) + \sqrt{2m\gamma T}\xi(t)$$

- Constant temperature is achieved by the balance of friction and gaussian white noise $\Rightarrow$ fluctuation-dissipation theorem,

$$\left\langle \xi(t)\,\xi\left(t'\right)\right\rangle = \delta\left(t - t'\right)$$

- Test Langevin thermostat on a 1-d harmonic oscillator

A Langevin Oscillator

- Langevin dynamics on a 1-D oscillator with $\omega = 1$. Trajectory of kinetic and potential energy and position, $\gamma = 0$, $\gamma = 1$ and $\gamma = 10^3$.

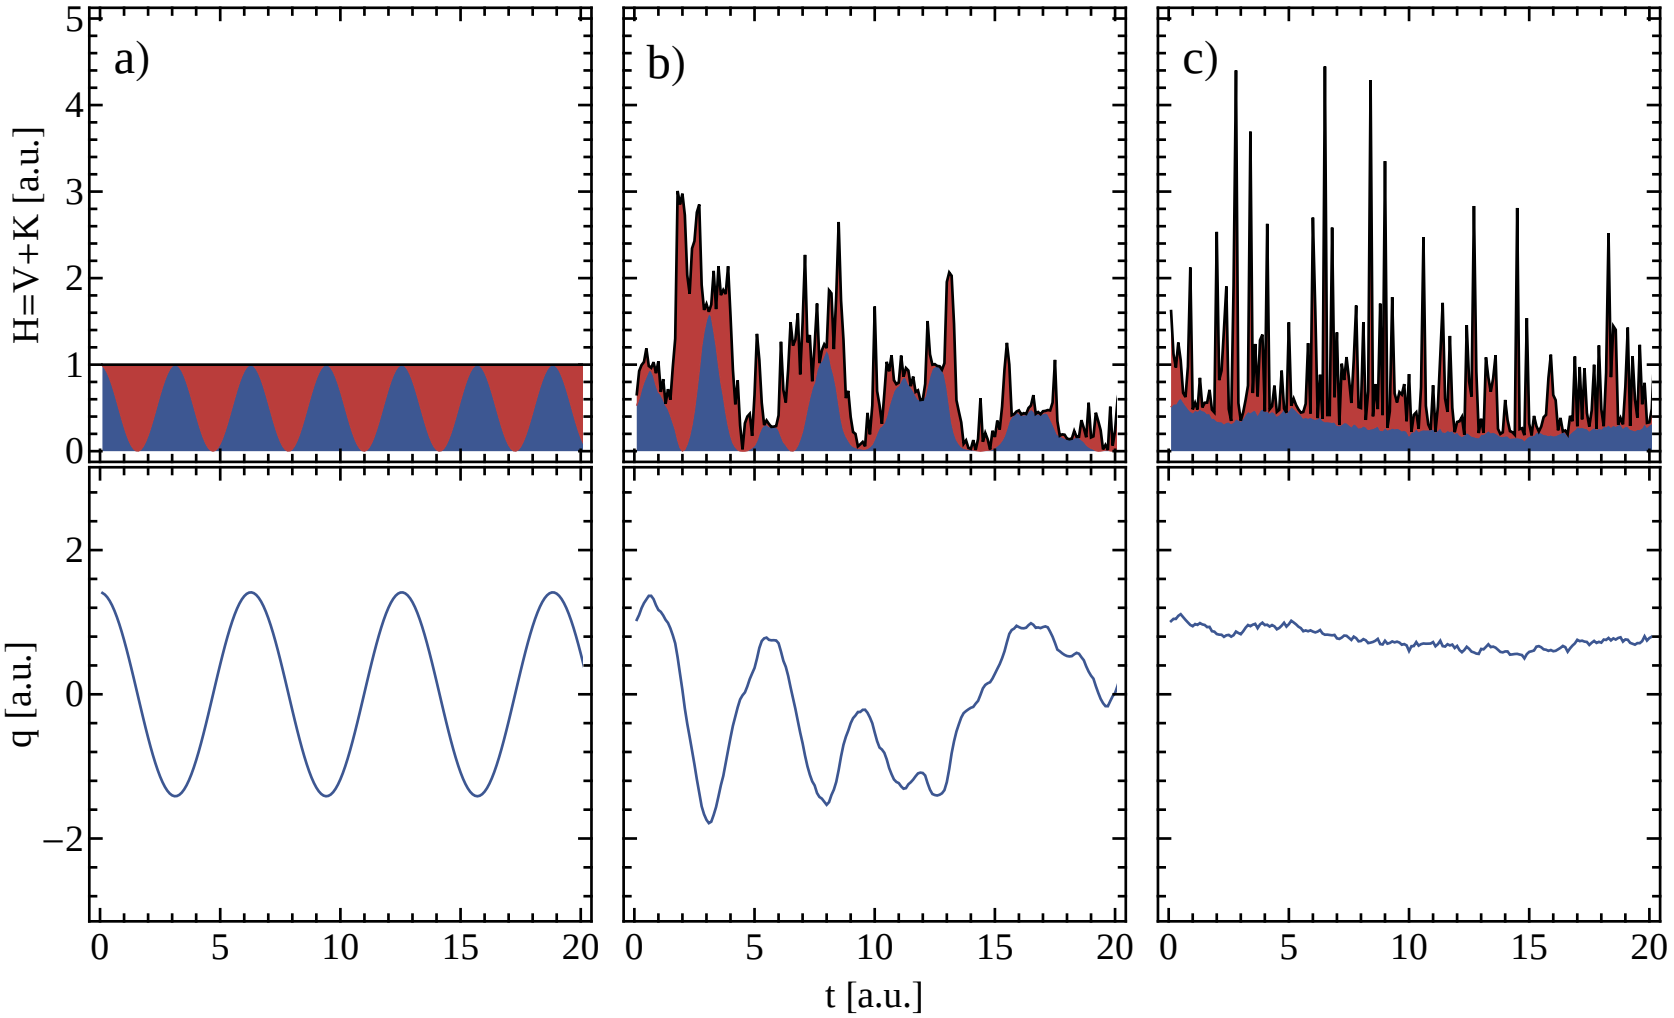

A Langevin Oscillator

- We can compute analytically correlation times, and distinguish different regimes

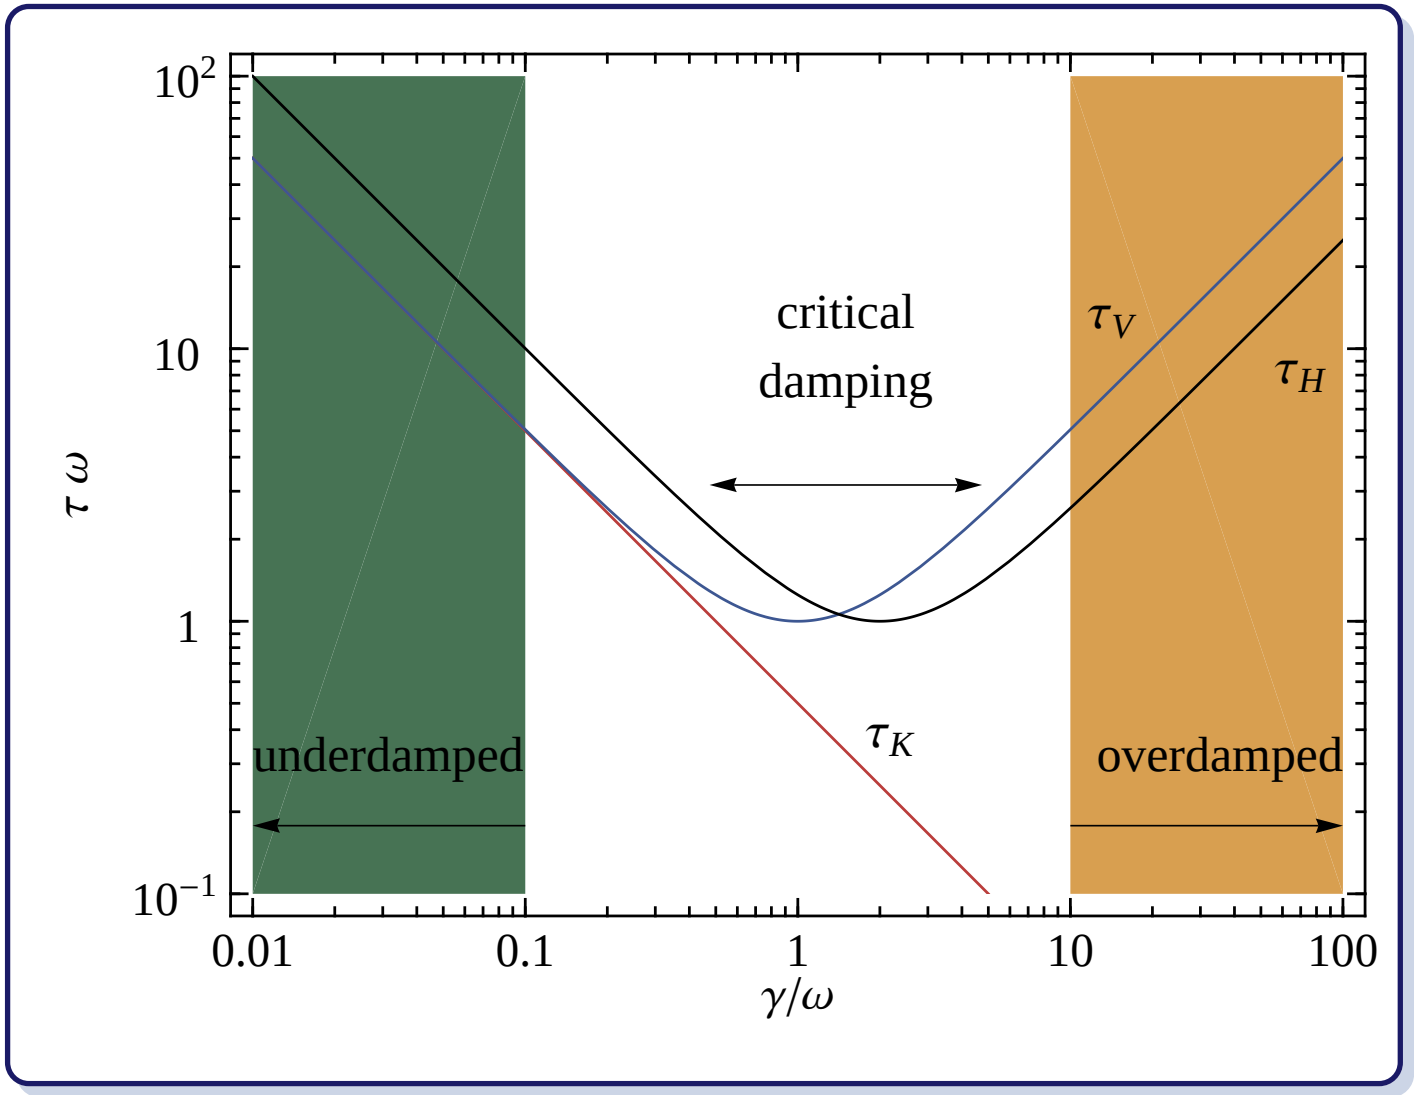

Optimal-sampling GLE

- What can we do if there are multiple frequencies? Only one would respond optimally to a Langevin thermostat!

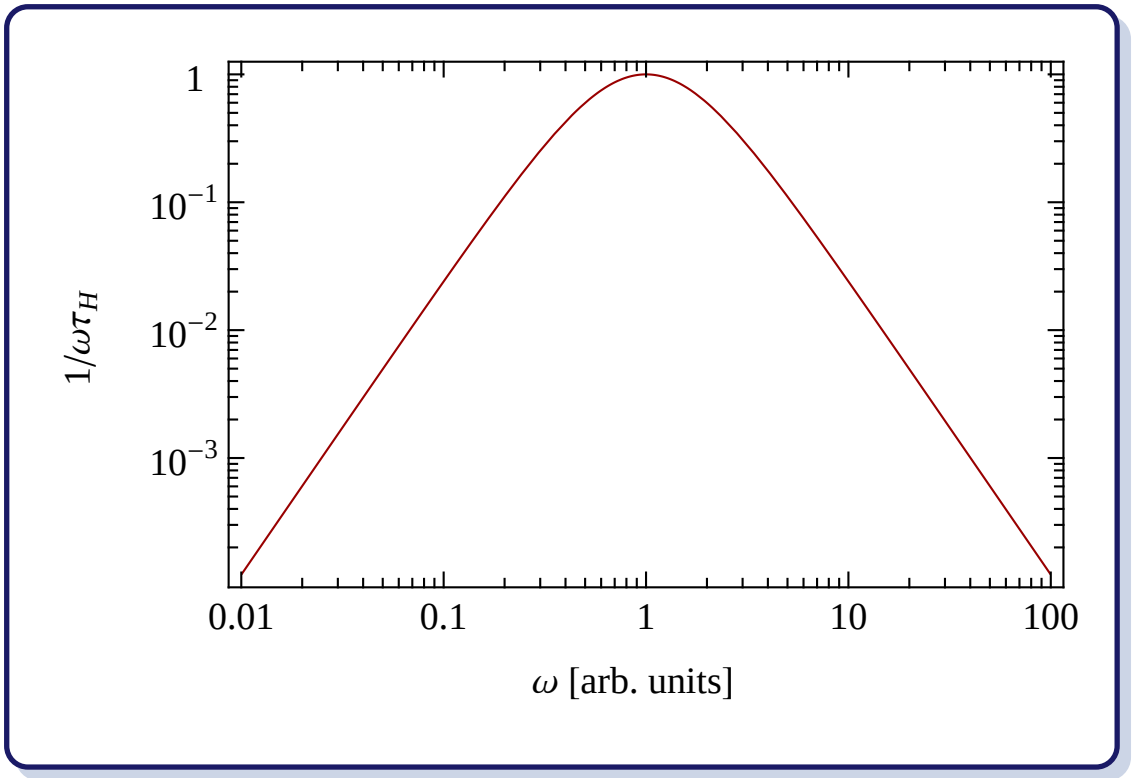

$$\dot{p}(t) = -\gamma p(t) + \sqrt{2m\gamma T}\xi(t)$$
$$\langle \xi(t)\xi(0)\rangle = \delta(t)$$

Optimal-sampling GLE

- What can we do if there are multiple frequencies? Only one would respond optimally to a Langevin thermostat! **Use non-Markovian noise to obtain constant efficiency!**

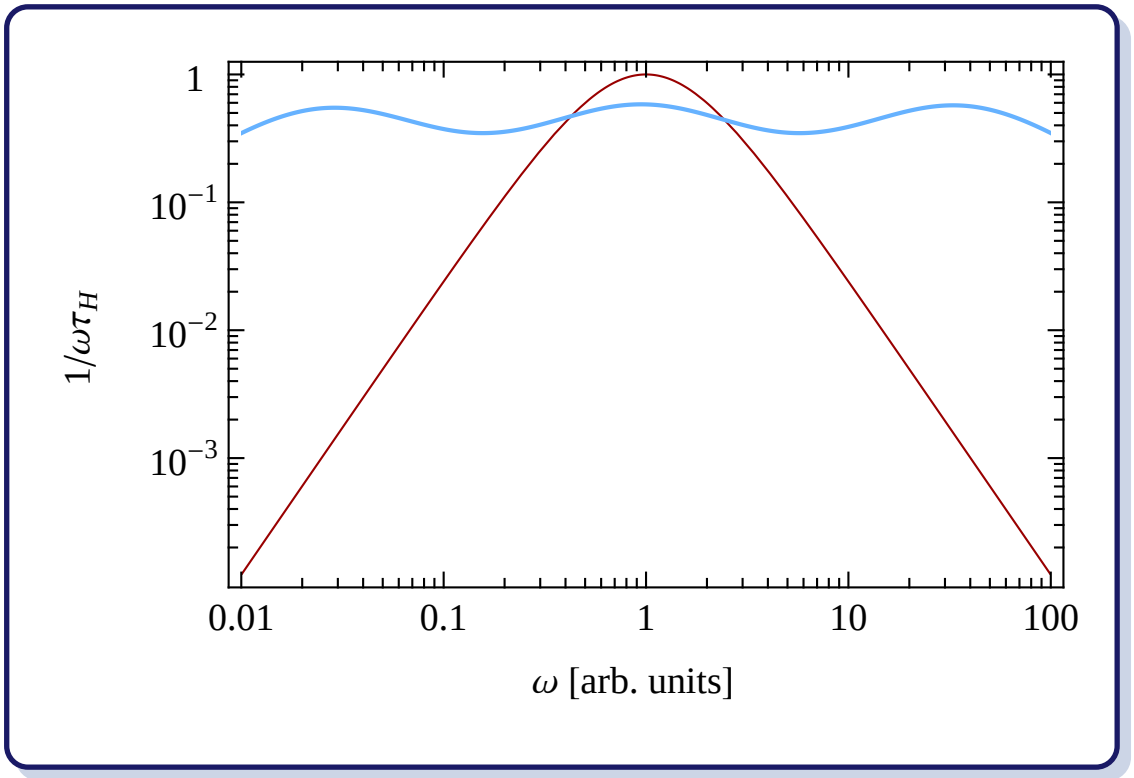

$$\dot{p}(t) = -\int_0^\infty K(t')\, p(t-t')\, \mathrm{d}t' + \sqrt{mT}\zeta(t)$$
$$\langle \zeta(t)\, \zeta(0) \rangle = K(t)$$

Let's get into water

- A critical discussion of different schemes applied to liquid water
  - Water is difficult! Normal modes span several order of magnitude in frequency
  - Diffusive motion requires complex rearrangements in H-bonds network
- Classical dynamics of liquid water using a flexible, TIP4P-like model
- We monitor total potential energy, cell's dipole moment (necessay to evaluate ϵ, difficult to converge) and kinetic temperature projected on internal modes, librations and center of mass motion.

Overdamped LE

- White-noise **local** Langevin thermostat, $\gamma^{-1} = 1$ fs
  - Lightning-fast decorrelation of velocities
  - Overdamped dynamics, configurational sampling is greatly slowed down

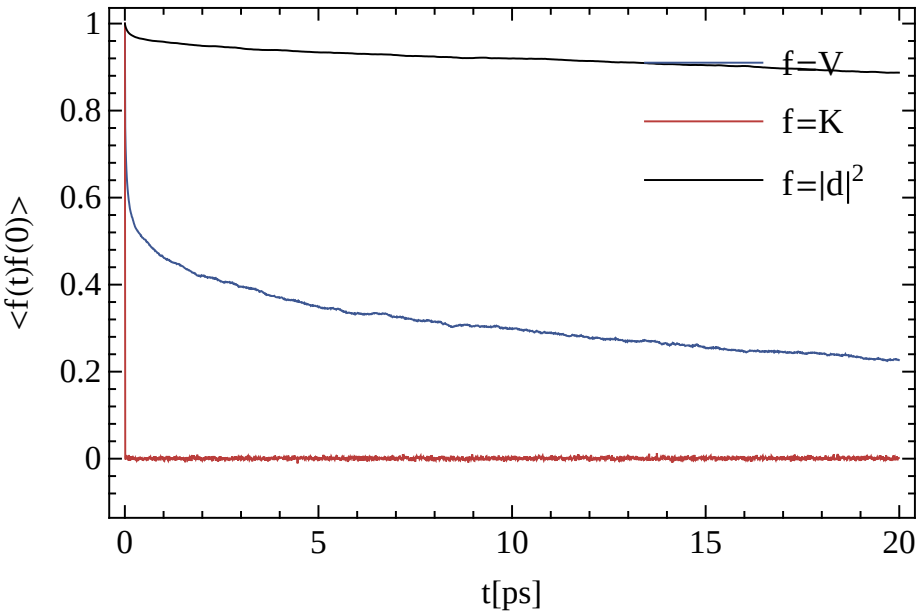

Potential & kinetic energy

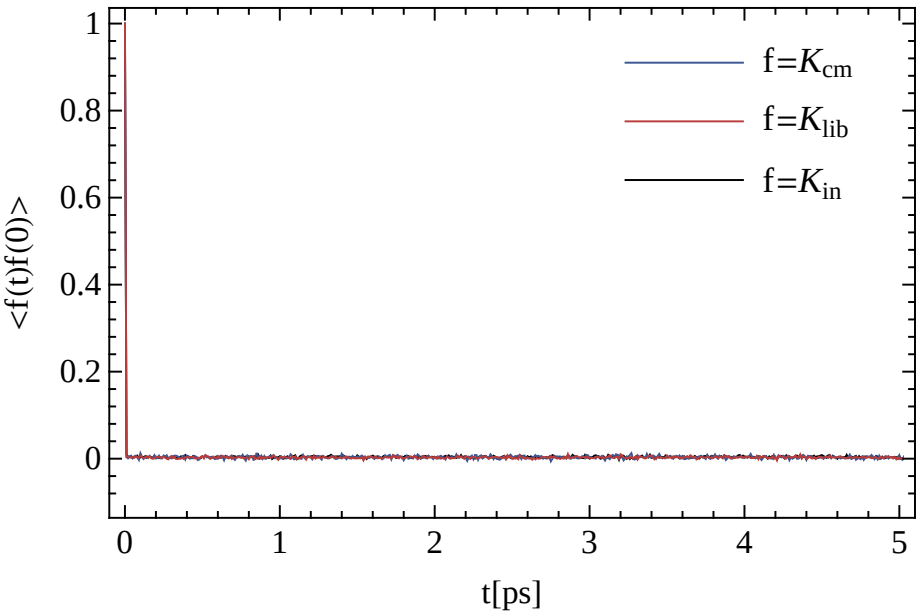

Projected kinetic energy

Mild LE

- White-noise **local** Langevin thermostat, mild friction $\gamma^{-1} = 1$ ps
  - Slower relaxation of momenta
  - No overdamping $\Rightarrow$ reasonable sampling of positions

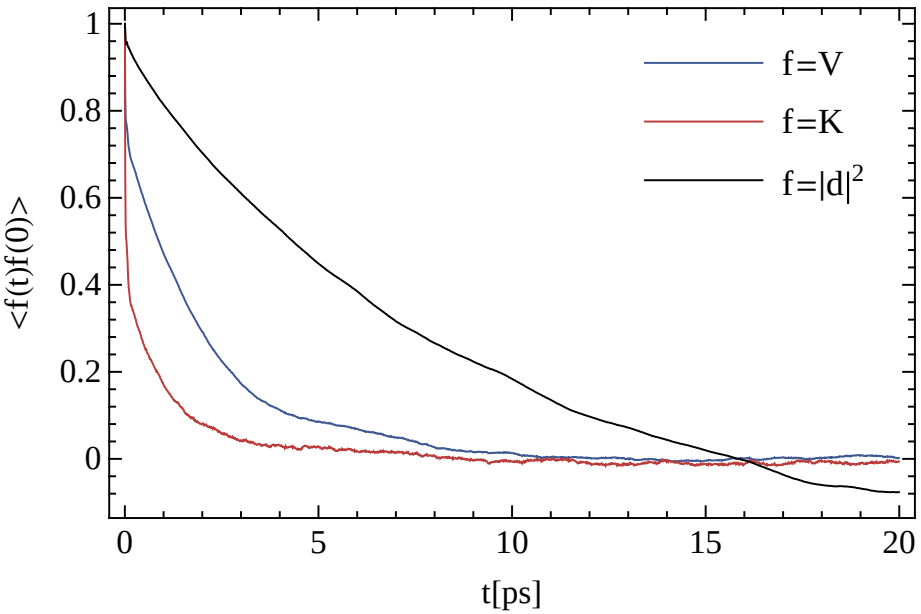

Potential & kinetic energy

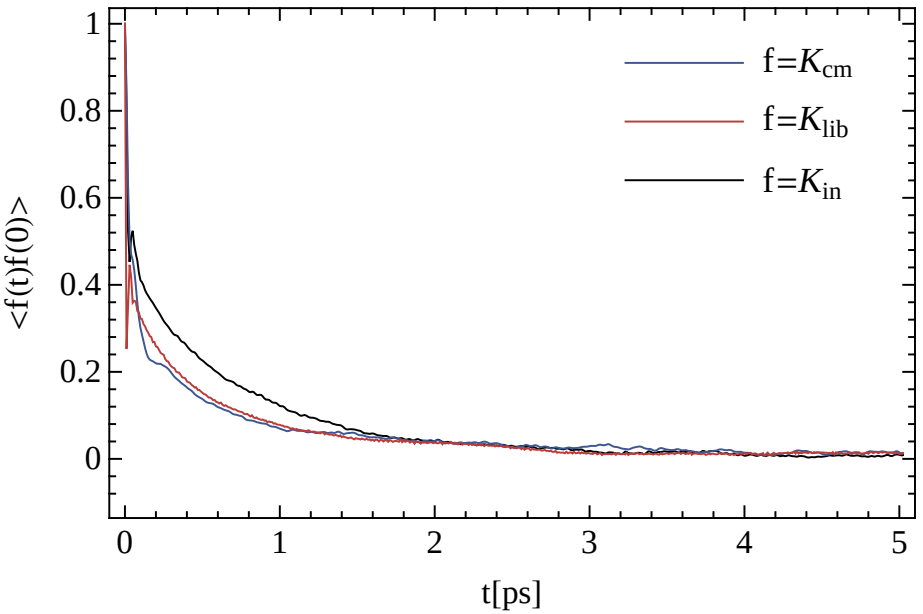

Projected kinetic energy

Optimal Sampling LE

- Optimal-sampling LE, fitted to encompass the whole range of vibrations
  - Efficient sampling of all normal modes
  - Reduced overdamping, avoided slowing-down of configuration sampling

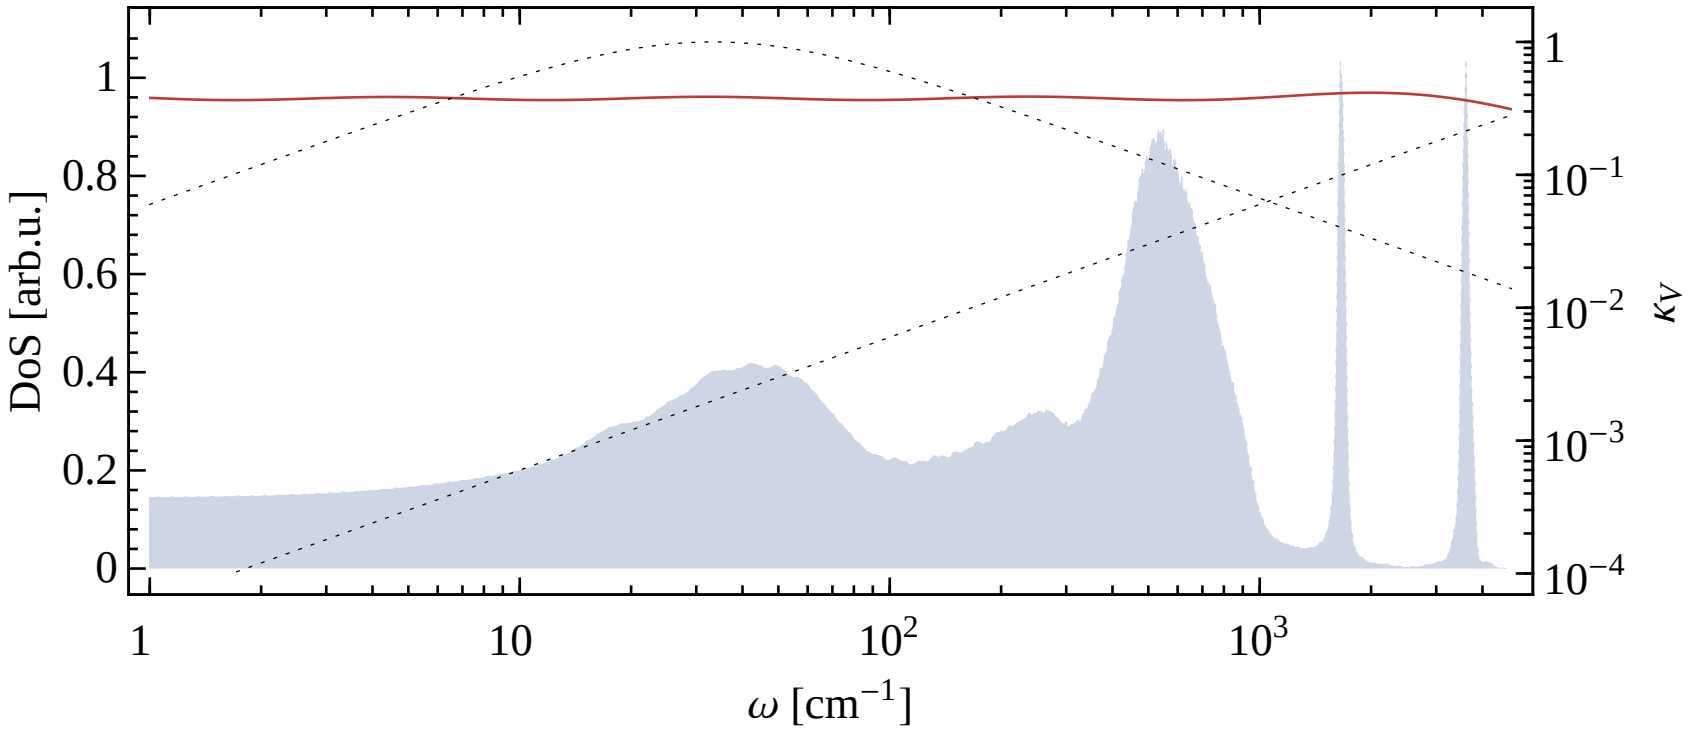

sampling efficiency of white noise and colored-noise,

Optimal Sampling LE

- Optimal-sampling LE, fitted to encompass the whole range of vibrations
  - Efficient sampling of all normal modes
  - Reduced overdamping, avoided slowing-down of configuration sampling

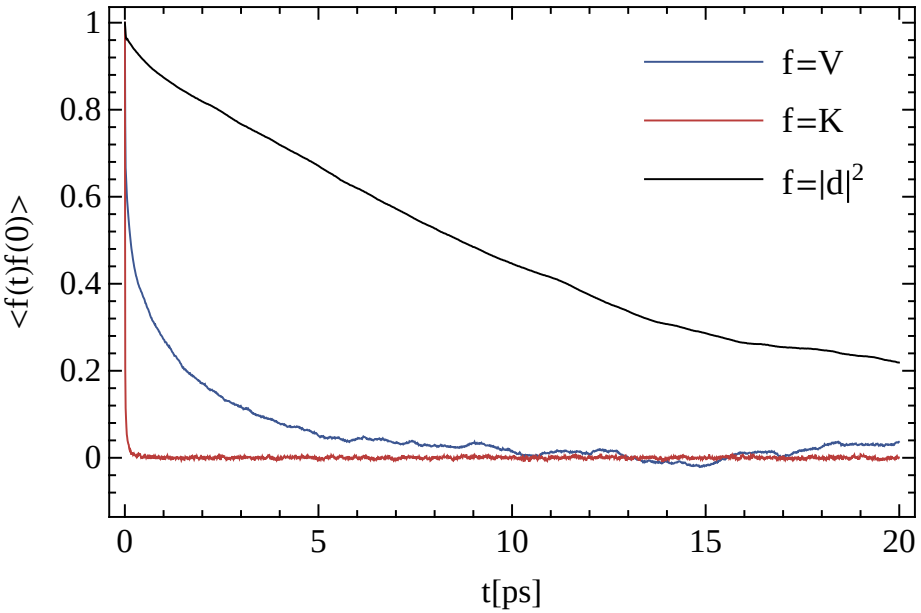

Potential & kinetic energy

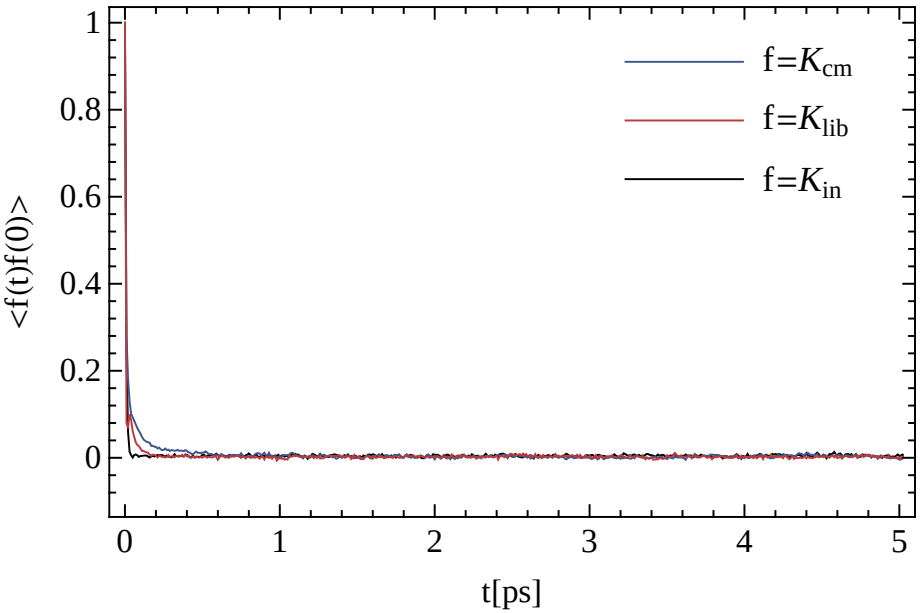

Projected kinetic energy

What if we go global?

- Stochastic velocity rescaling, $\gamma^{-1} = 1$ fs. The dynamics is not disturbed and **total** kinetic energy is sampled very efficiently!
  - Very efficient sampling of the difficult property of total dipole moment
  - Projected temperatures relax slowly... do we really care?

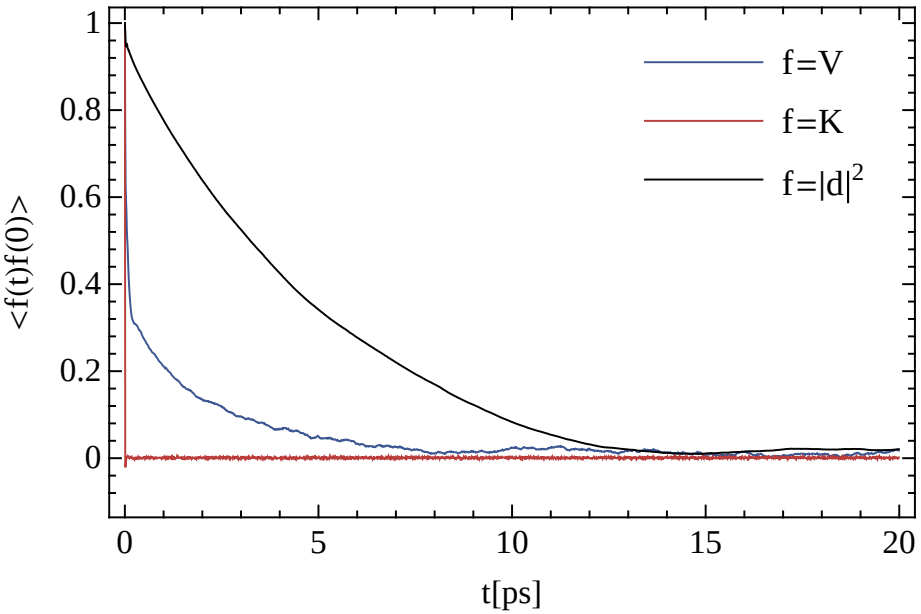

Potential & kinetic energy

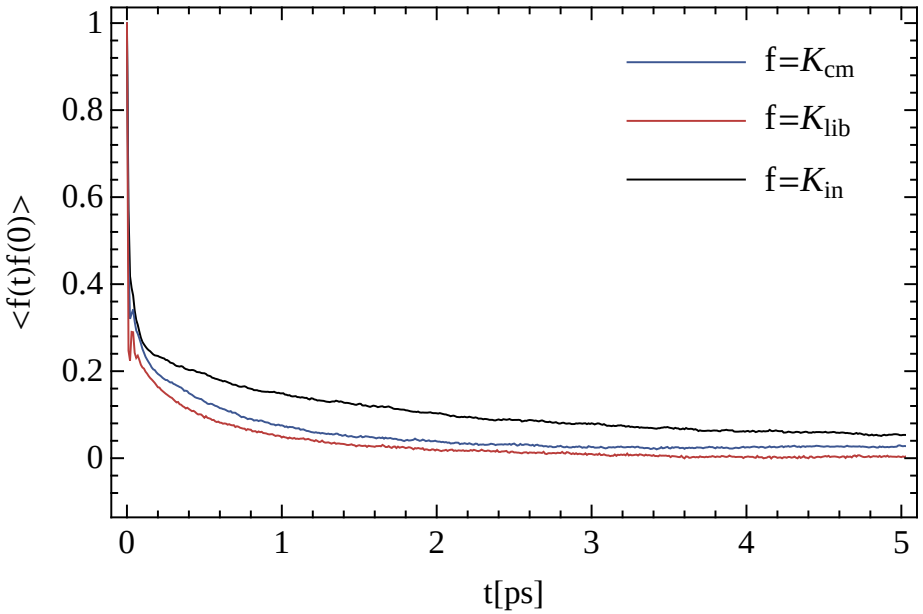

Projected kinetic energy

The magic of MD

- Global thermostats work nicely for slow configurational properties because they do not disturb slow, diffusive modes

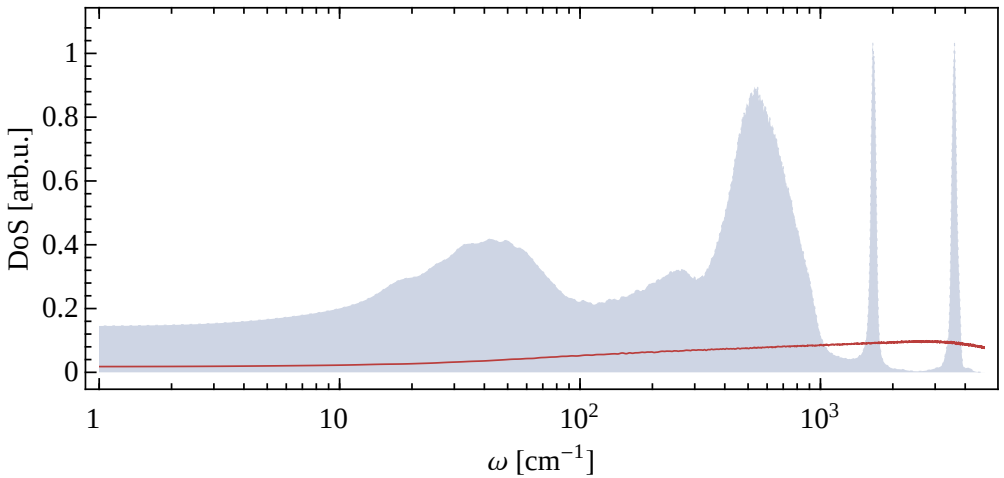

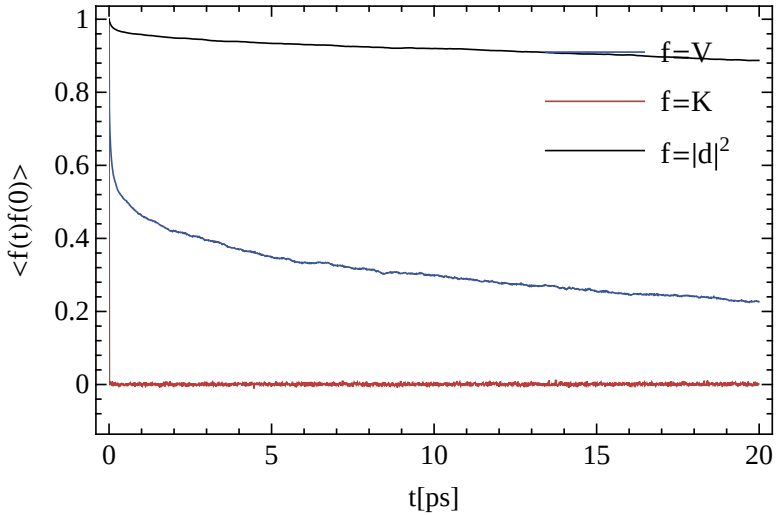

Local Langevin thermostat, $\gamma^{-1} = 1$ fs

The magic of MD

- Global thermostats work nicely for slow configurational properties because they do not disturb slow, diffusive modes

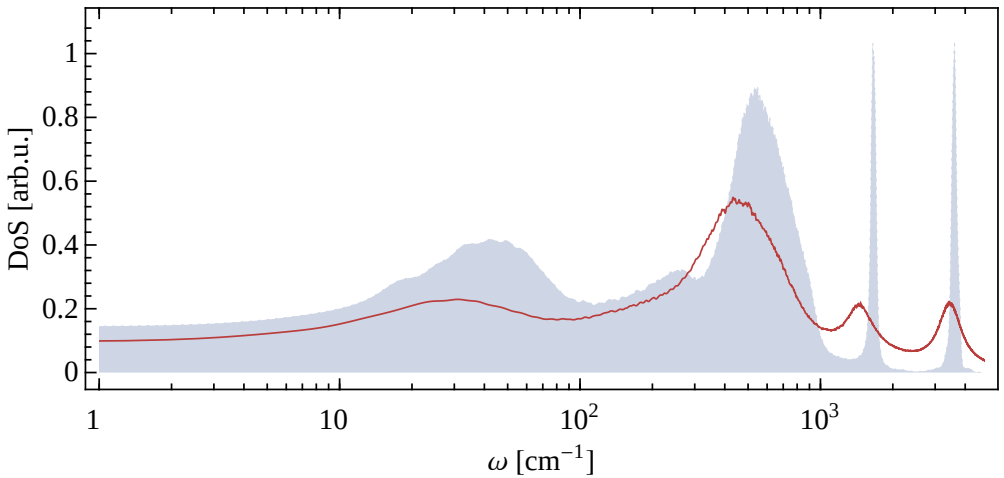

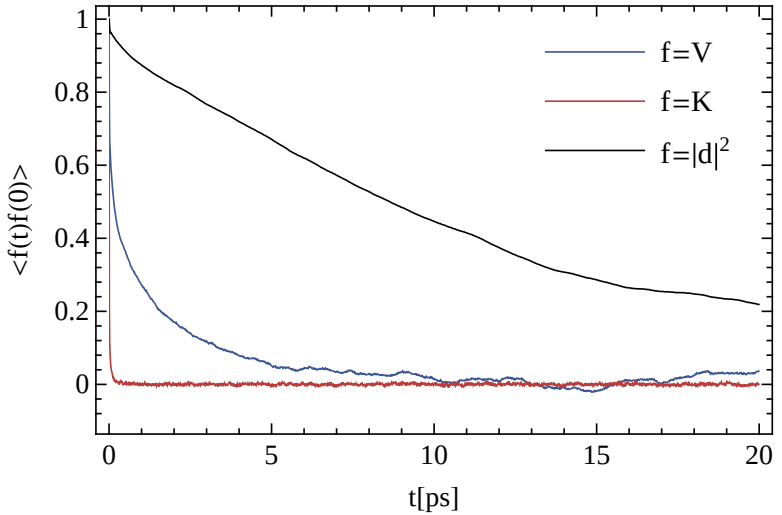

Optimal sampling GLE

The magic of MD

- Global thermostats work nicely for slow configurational properties because they do not disturb slow, diffusive modes

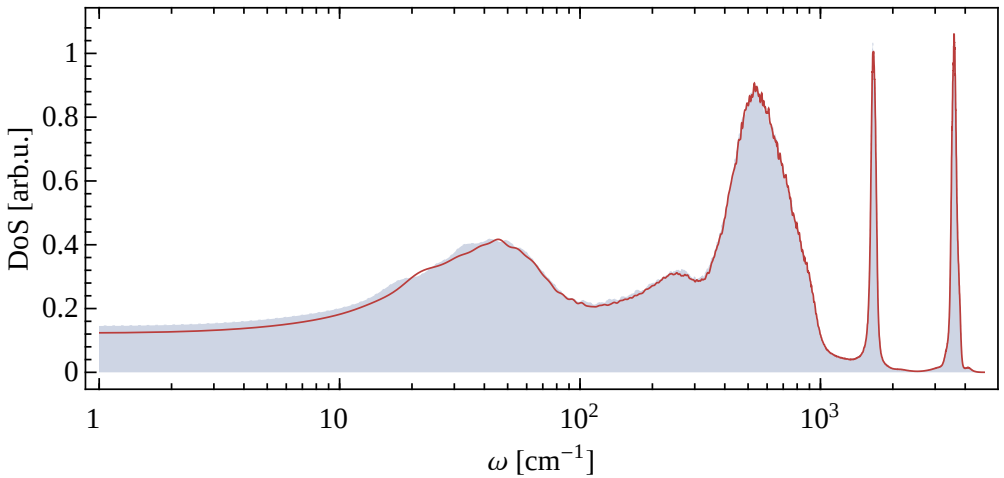

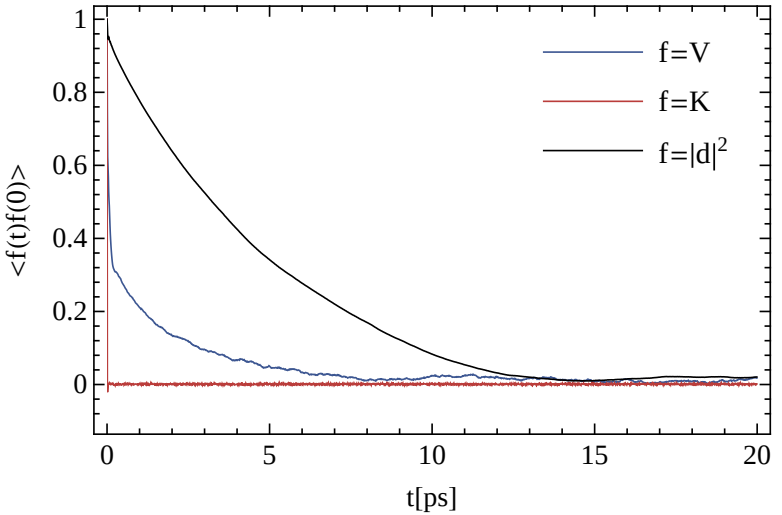

Stochastic velocity rescaling, $\gamma^{-1} = 1$ fs

- Vibrational density of states is almost equal to NVE!

No free lunch!

- One must pay attention when using global thermostats: local equilibration relies on intrinsic ergodicity of the system
- This is particularly dangerous when performing metadynamics, or quasi-equilibrium free-energy methods in general
- Energy is injected in localized modes, but only the **total** kinetic energy is monitored
  - Total temperature is rescaled $\Rightarrow$ one feels safe but...

No free lunch!

- One must pay attention when using global thermostats: local equilibration relies on intrinsic ergodicity of the system
- This is particularly dangerous when performing metadynamics, or quasi-equilibrium free-energy methods in general

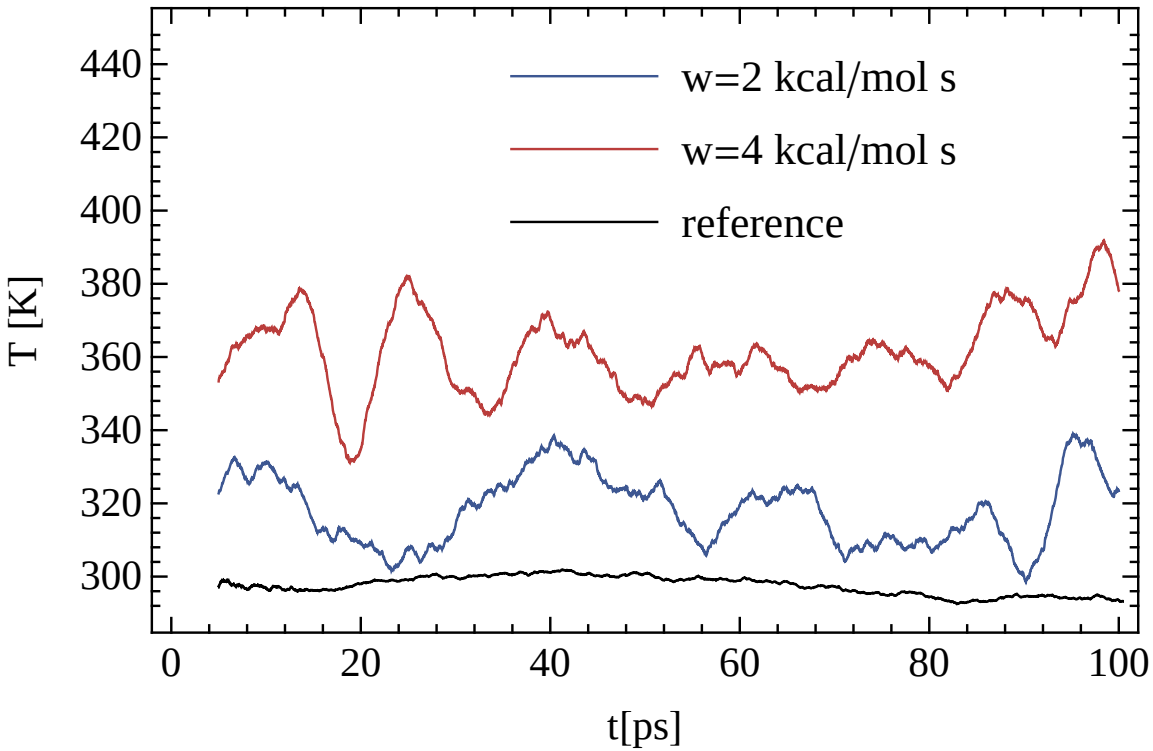

computation of PMF of a $CaCO_3$ molecule in water;
kin. temperature of the ions during the meta run

A GLOCAL thermostat

- Let's put all the ideas together. We want to use a local thermostat, but leave diffusive, collective motions alone. We must think **global**, and act **local**!
- Within GLE framework one can estimate and minimize the disturbance on selected frequencies ($\eta(\omega)$ parameter). Also, we require effective coupling by maximizing $\kappa_V = 1/\omega\tau_V$.

A GLOCAL thermostat

- Let's put all the ideas together. We want to use a local thermostat, but leave diffusive, collective motions alone. We must think **global**, and act **local**!
- Within GLE framework one can estimate and minimize the disturbance on selected frequencies ($\eta(\omega)$ parameter). Also, we require effective coupling by maximizing $\kappa_V = 1/\omega\tau_V$.

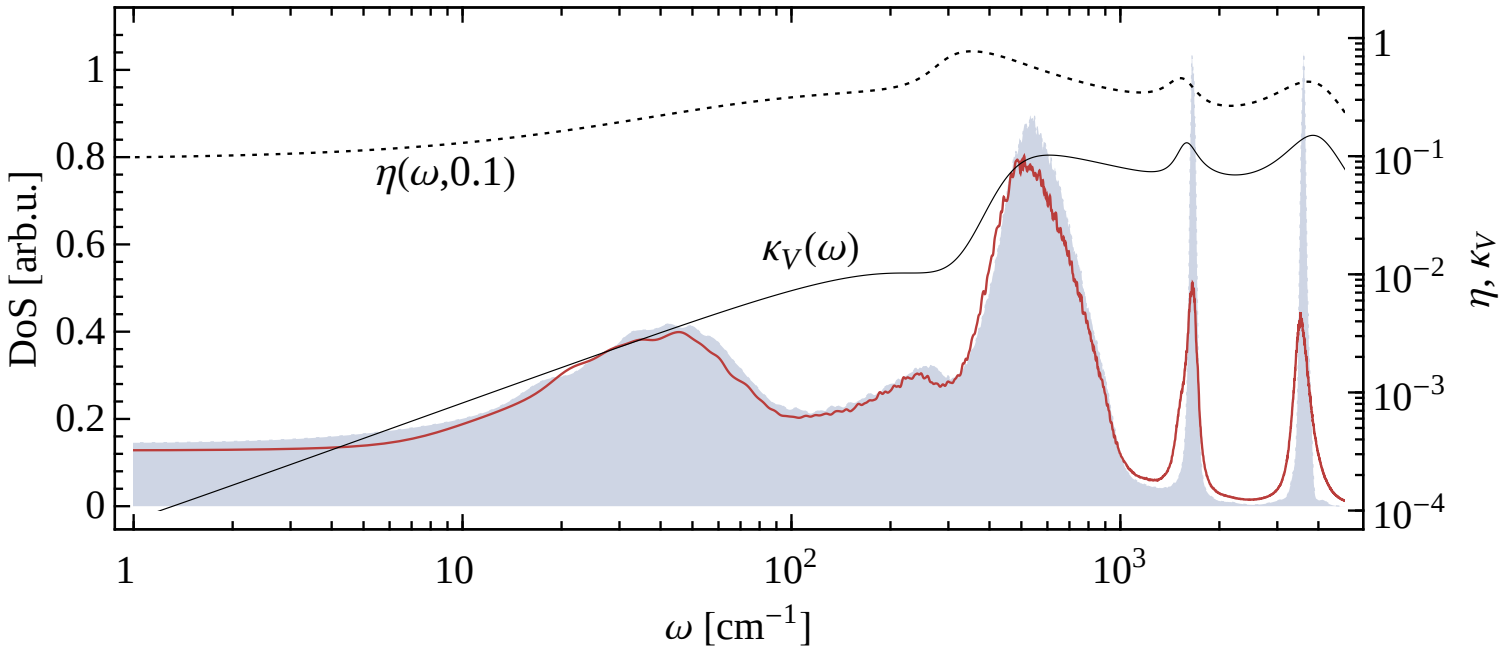

sampling efficiency, local disturbance and actual density of states for the "water thermostat"

Bechmark: global *vs* GLOCAL

- Comparison of sampling properties of stochastic rescale vs GLE

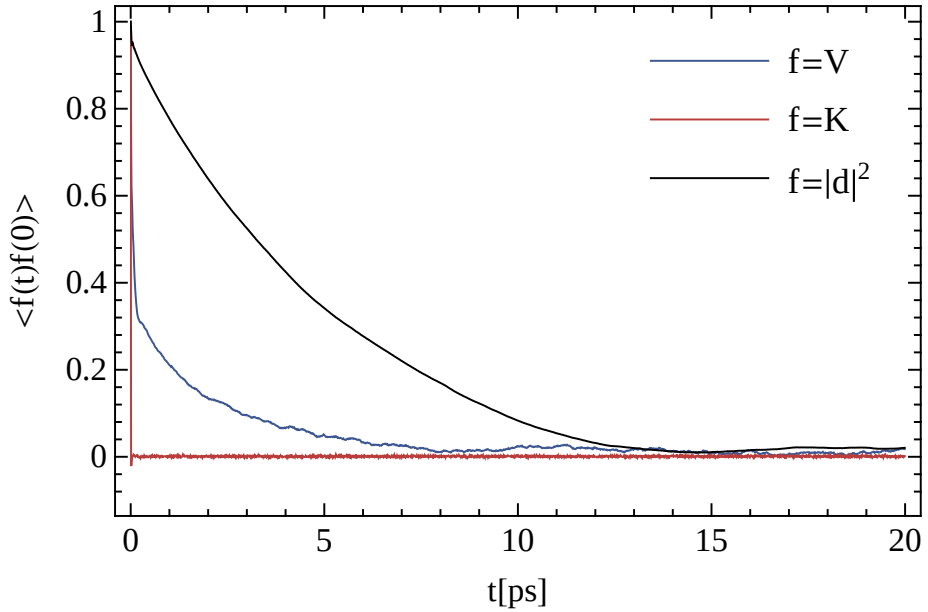

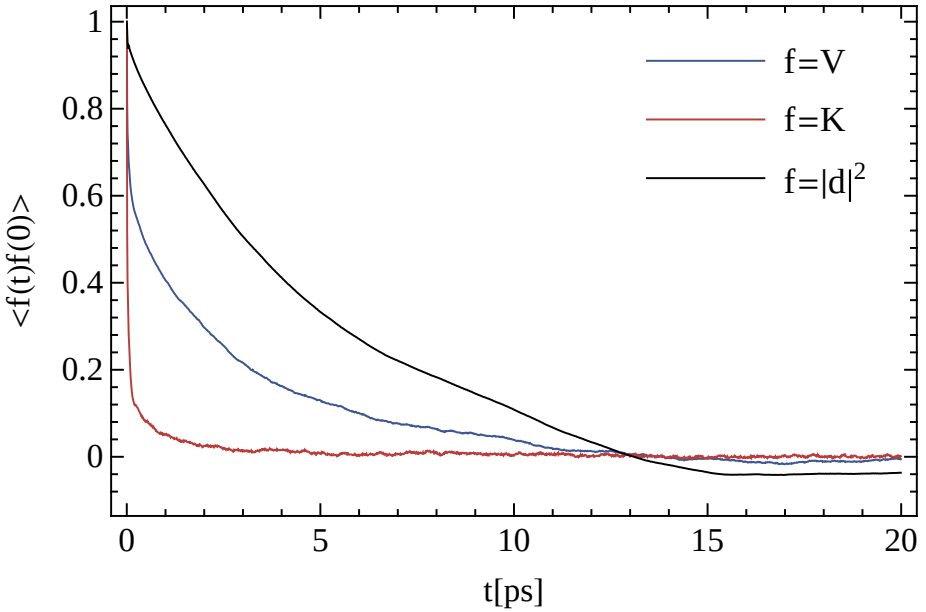

global, $\gamma^{-1} = 1$ fs GLOCAL

Potential & kinetic energy, total dipole

Bechmark: global *vs* GLOCAL

- Comparison of sampling properties of stochastic rescale vs GLE

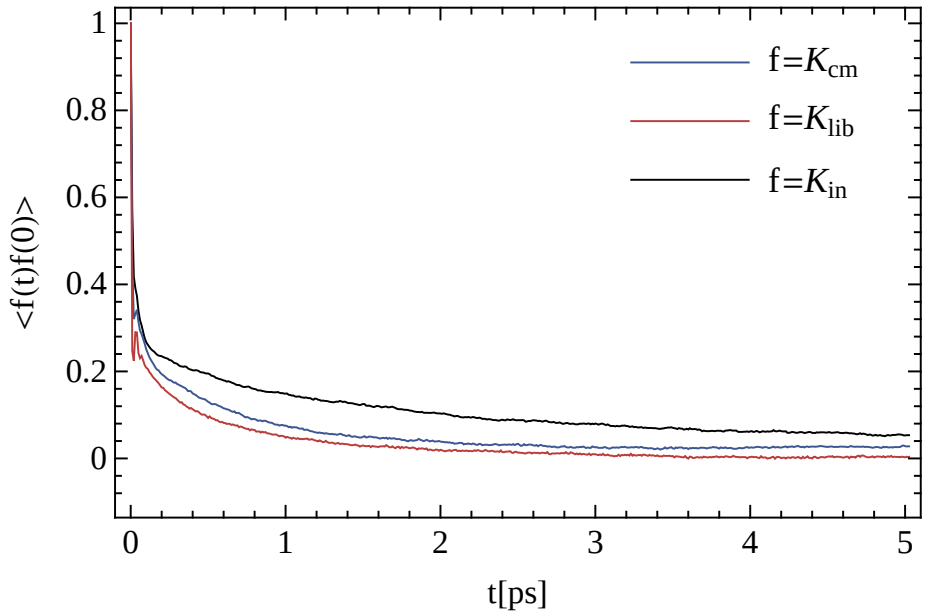

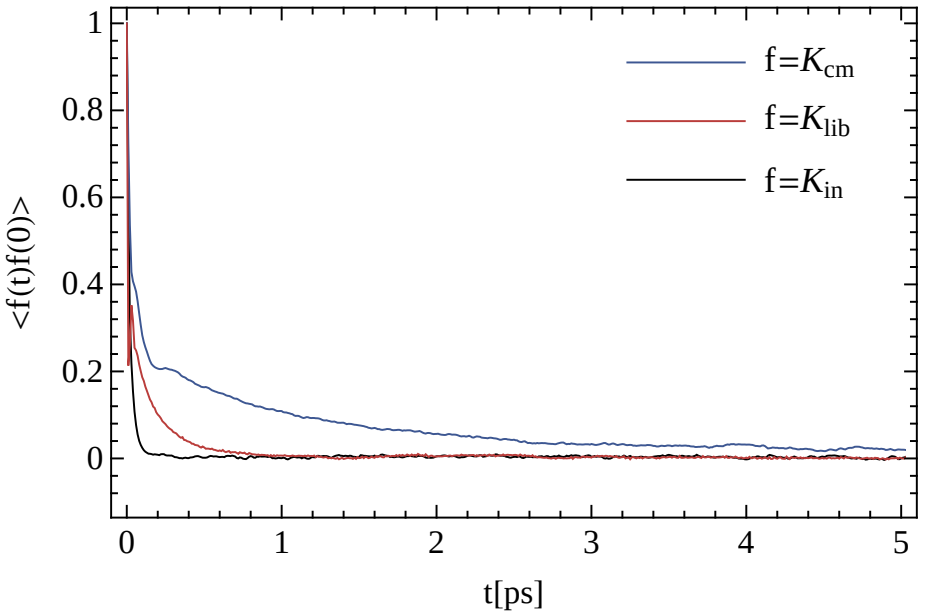

global, $\gamma^{-1} = 1$ fs GLOCAL

Kin. energy projected on different modes

Bechmark: global *vs* GLOCAL

- Comparison of sampling properties of stochastic rescale vs GLE

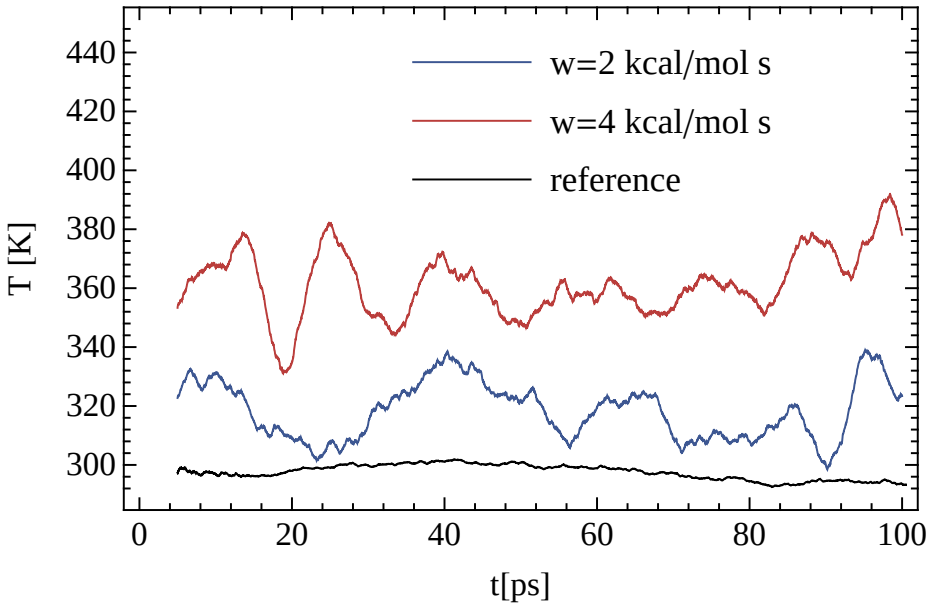

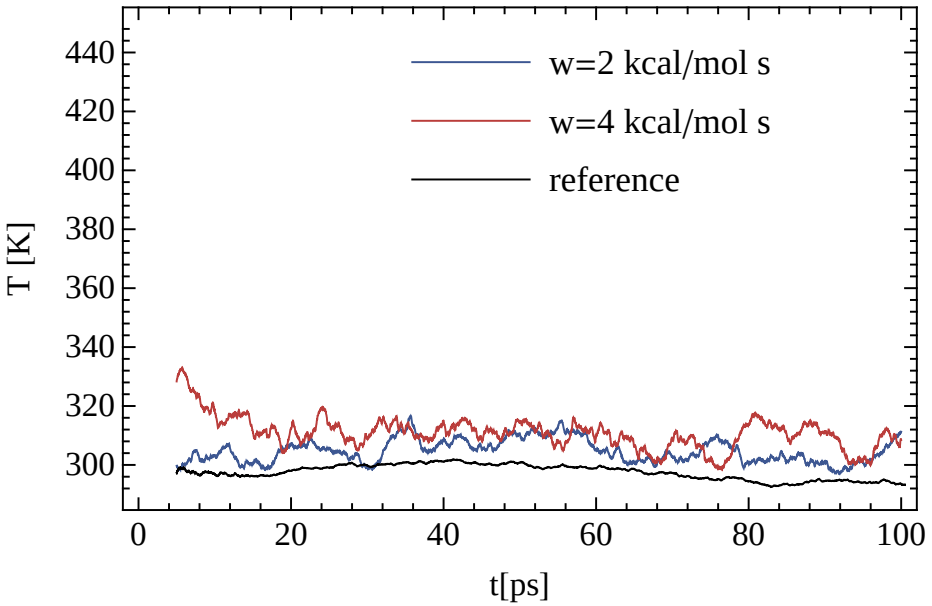

global, $\gamma^{-1} = 1$ fs GLOCAL

Temperature of ions during meta run of $CaCO_3$ PMF

Conclusions: a list of caveats

Thermostatting can impact your simulation in many different ways!

- Different observables might have very different relaxation times, and an observable might have correlations on multiple time scales !
- Molecular dynamics is very good at sampling diffusive motion. Aggressive thermostatting might degrade sampling efficiency !
- Do not look at total kinetic temperature alone: that can be made to uncorrelate very quickly by just resampling momenta at every time-step !
- Everything becomes more tricky when doing biased dynamics: impacts not only efficiency but also the actual result! !
- Global thermostats do very well on strongly coupled systems, but one must be careful, as they might hide non-equilibrium conditions.

Conclusions: why colors?

- Testing the thermostat is boring and expensive. Still, a bad choice can cause larger statistical and even systematic errors!
- GLE framework allows to predict the properties of the dynamics from many points of view:
  - sampling efficiency in the harmonic limit
  - disturbance of the dynamical properties
- Optimal-sampling GLE provides a no-brainer local thermostat which will be strong on local modes and won't overdamp diffusion
- With a little effort, even better performance can be obtained: it's truly **à la carte** thermostating!
- You can do much more: quantum thermostat, δ-termostat, more to come... **all within the same framework.**
- ... will you help me implement it in established MD codes?

THANKS!

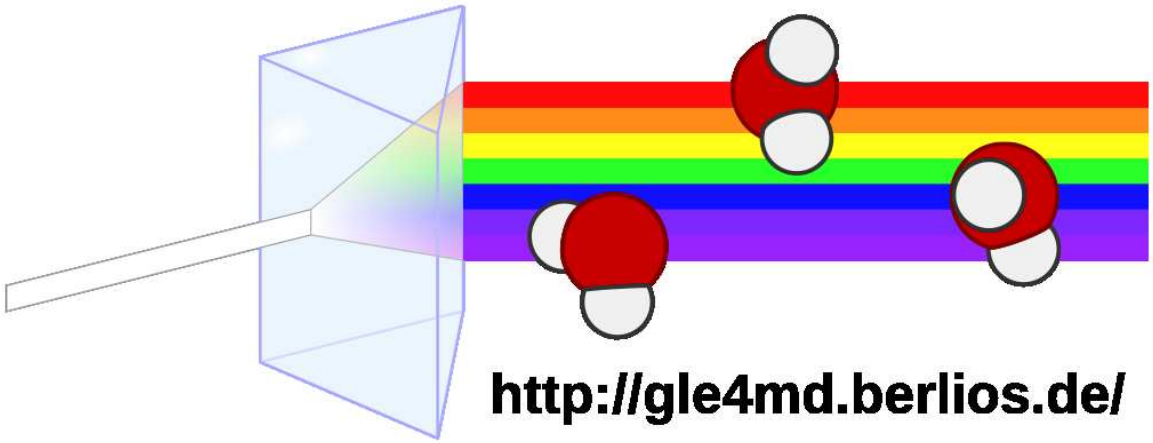

http://gle4md.berlios.de/