

Simulating Rare Events in Dynamical Processes

Cristian Giardina · Jorge Kurchan · Vivien Lecomte ·
Julien Tailleur

Received: 24 June 2011 / Accepted: 6 September 2011 / Published online: 23 September 2011
© Springer Science+Business Media, LLC 2011

Abstract Atypical, rare trajectories of dynamical systems are important: they are often the paths for chemical reactions, the haven of (relative) stability of planetary systems, the rogue waves that are detected in oil platforms, the structures that are responsible for intermittency in a turbulent liquid, the active regions that allow a supercooled liquid to flow. . . . Simulating them in an efficient, accelerated way, is in fact quite simple.

In this paper we review a computational technique to study such rare events in both stochastic and Hamiltonian systems. The method is based on the evolution of a family of copies of the system which are replicated or killed in such a way as to favor the realization of the atypical trajectories. We illustrate this with various examples.

Keywords Large deviations · Rare events · Simulations

1 Introduction

When a dynamical system is complex enough, it becomes no longer feasible—or indeed, interesting—to describe every possible trajectory. A first step is then to study what a ‘typical trajectory’ does. For Hamiltonian dynamics, Statistical Mechanics provides us with powerful techniques to compute some properties of such typical trajectories, but for generic dynamics we must in most cases resort to simulations.

C. Giardina (✉)
Università di Modena e Reggio Emilia, viale A. Allegri, 9, 42121 Reggio Emilia, Italy
e-mail: c.giardina@tue.nl

J. Kurchan
CNRS UMR 7636 PMMH, ESPCI, 10 rue Vauquelin, 75005 Paris, France

V. Lecomte
Sorbonne Paris Cite, LPMA, UMR 7599 CNRS, Univ. Paris Diderot, 75205 Paris, France

J. Tailleur
Sorbonne Paris Cite, MSC, UMR 7057 CNRS, Univ. Paris Diderot, 75205 Paris, France

There are many situations in which the trajectories that matter are not the typical ones, but rather ‘rare’ ones reached from exceptional initial conditions, or particularly infrequently. Consider the following examples:

- Planetary systems are in general chaotic, and the different sets of present conditions, falling within the range of observational error, may lead to widely varying inferences about the past and future. Because we do not expect that an observed system has been created recently, or will be destroyed immediately, we must understand how this comes about, and we are naturally led to a statistical study of the trajectories.
- Molecular dynamics is in many cases characterized by long periods of vibrations around a local metastable configuration, punctuated by relatively rapid but infrequent ‘activation’ events, leading to a major rearrangement. Because they are the essential steps of chemical transformations, it is of the greatest importance to be able to simulate such events in an accelerated way, without having to wait for them to happen spontaneously. There is a vast literature on this subject.
- In a similar fashion, supercooled liquids and glasses are characterized by vibrational dynamics, with events localized in time *and space* where the transformations take place. These ‘dynamic heterogeneities’ are the analogues of reaction paths in chemical systems.
- It has long been known that, in a liquid undergoing fully developed turbulence, due to the presence of abnormally large fluctuations of velocities, the dynamics are intermittent. The natural question is which dynamic features are responsible for this.
- In the sea there have been reports of (‘rogue’) waves of exceptionally large amplitudes. They are rare, but much more common than one would expect from a Gaussian distribution. The subject is of obvious interest, and is still very much open.
- Transport of energy or particles across a sample is facilitated by exceptional ‘ballistic’ trajectories, or hindered by situations resembling traffic jams.
- When a system is subject to external forcing, the power injected (or the entropy production) during a given time is a quantity that depends on the particular trajectory it is following. The Second Law of thermodynamics sets limits on the expectation value of these quantities, but does not limit the extent of the (rare) fluctuations. Thus, one may extract work from a system while lowering the total entropy, but the probability of this goes down exponentially with its size, and with the interval of time.

All of these problems may be studied by simulating repeatedly, or for long times, the true dynamics. However, as one may imagine, this procedure soon becomes unfeasible. There are basically two types of methods to generate in a controlled way rare events. The *Path-sampling* method amounts to Monte Carlo dynamics in trajectory space, correctly designed to weigh each trajectory with the desired bias. A second strategy works directly in configuration space: one introduces a population of copies of the initial system and relies on a mixture including the original dynamics, supplemented with a ‘Darwinian pressure’—again, in a controlled way—to favor the exploration of atypical trajectories. In this review we concentrate on the second class.

The paper is organized as follows. The population dynamics with cloning is introduced in Sect. 2, where it is shown how it can be used to compute the large deviation function (or rather its Legendre transform) of extensive observables of the trajectories of a diffusive dynamics with drift and a multiplicative (cloning) term. The relative weight of the drift and cloning terms is analyzed in Sect. 3, where it is shown how a change of bases can help in adjusting their relative contribution. Then a series of examples from different contexts follows. Purely stochastic systems are studied in Sects. 4 and 5, where the large deviations of, respectively, the current in interacting particle systems and the dynamical activity in

kinetically constrained models are analyzed. Sections 6 and 7 consider examples of deterministic dynamics, such as the standard map and the Hamiltonian Fermi-Pasta-Ulam model, for which trajectory with large or small Lyapunov exponent are studied, or the Sinai billiard, for which the symmetry associated with the fluctuation theorem is easily verified. The last Sect. 8 suggests how the numerical method of cloning could be used also in the study of the stability of planetary systems.

2 Population Dynamics

To fix ideas, consider a noisy dynamics for a vector \mathbf{x} whose components evolve as:

$$\dot{x}_i = f_i(\mathbf{x}) + \eta_i(t) \tag{2.1}$$

with η_i a noise which for simplicity we shall suppose is Gaussian and white, with variance $2T_i$. The probability of a trajectory up to time t is found by writing $\eta_i = \dot{x}_i - f_i(\mathbf{x})$:

$$P[\mathbf{x}(t)] \propto e^{-\sum_i \int_0^t dt' \frac{(\dot{x}_i - f_i)^2}{4T_i}} \tag{2.2}$$

As an example, we wish to calculate the probability that a certain quantity $A[\mathbf{x}]$ takes a *time-averaged* value \mathcal{A}_0 :

$$p(\mathcal{A}_0) = \int D[\mathbf{x}] P[\mathbf{x}(t)] \delta \left\{ \int_0^t dt' A(\mathbf{x}) - t\mathcal{A}_0 \right\} \tag{2.3}$$

It is more practical to compute the Laplace transform:

$$\begin{aligned} Z_t(\alpha) &= \int d\mathcal{A}_0 p(\mathcal{A}_0) e^{\alpha t \mathcal{A}_0} = \int D[\mathbf{x}] P[\mathbf{x}(t)] e^{\{\alpha \int_0^t dt' A(\mathbf{x})\}} \\ &\propto \int D[\mathbf{x}] e^{-\sum_i \int_0^t dt' \frac{(\dot{x}_i - f_i)^2}{4T_i} + \alpha \int_0^t dt' A(\mathbf{x})} \end{aligned} \tag{2.4}$$

In particular, for large times $p(\mathcal{A}_0)$ becomes a peaked function $p(\mathcal{A}_0) \sim e^{-tI(\mathcal{A}_0)}$, with $I(\mathcal{A}_0)$ the large deviation function given by the Legendre transform [1]:

$$I(\mathcal{A}_0) = \sup_{\alpha} \left[\mathcal{A}_0 \alpha - \lim_{t \rightarrow \infty} \frac{1}{t} \log Z_t(\alpha) \right] \tag{2.5}$$

The last line of (2.4) may be interpreted as a sum over paths with a modified weight, and may be simulated with path sampling methods. The strategy we describe in this paper is instead to notice that (2.4) may be interpreted as describing the following dynamics:

- Consider a population of infinitely many non-interacting ‘clones’ of the system $\mathbf{x}^a(t)$ satisfying the original dynamics $\dot{\mathbf{x}}^a(t) = \mathbf{f}(\mathbf{x}^a) + \boldsymbol{\eta}^a$. The noise of each clone is independent from the others.
- At each time interval δt , each clone is either killed or replicated, so that it is replaced on average by $\exp(\alpha A(\mathbf{x}^a) \delta t)$ clones.

This population dynamics is such that the average cloning or pruning rate of clones yields at large times $Z_t(\alpha)$. In practice, we do not simulate infinitely many clones of the initial system and we explain in the following how to adapt the dynamics to work with a large, but

finite, *fixed* number of clones (typically in the hundreds). We shall see how this simple idea, originally applied in the context of Diffusion Monte Carlo [2], may be adapted to a number of different problems. The actual specific form of the population dynamics involved depends on the nature of the problem (continuous or discrete state space, continuous or discrete time, etc.): we shall specify this in each example below. Similar strategies to simulate rare events have been advocated in other context with great success, see for example [3–5].

We have mentioned so far large deviations of a quantity of the form:

$$F[\mathbf{x}(t)] = \int_0^t dt' A(\mathbf{x}(t')) \tag{2.6}$$

In many cases, the functionals F depend also on the time-derivatives $\frac{d\mathbf{x}}{dt}$, and even are functions that are non-local in time. In these cases, the cloning rate at time t depends as well on the configurations at time $t' < t$.

The algorithms presented in this review give not only access to large deviations of the observable F but also allow one to compute the average of any observable among the corresponding, atypical, histories weighted by $e^{\alpha F}$, allowing to answer questions such as “*what happens with the vorticity of a fluid at a time and place where energy dissipation is unusually large?*”

The average of an observable \mathcal{O} at the final time t

$$\overline{\mathcal{O}}(\alpha, t) = \frac{\langle e^{\alpha F} \mathcal{O}(x(t)) \rangle}{\langle e^{\alpha F} \rangle} \tag{2.7}$$

is recovered from the corresponding average among the clones at that time. The averages at intermediate times (for $0 \ll t' \ll t$) $\overline{\mathcal{O}}(\alpha, t') = \frac{\langle e^{\alpha F} \mathcal{O}(x(t')) \rangle}{\langle e^{\alpha F} \rangle}$ may also be recovered by attaching to each clone at time t' the observed value of \mathcal{O} , and then constructing the average $\overline{\mathcal{O}}(\alpha, t')$ among the clones which have survived until the final time t . In the large time limit $t \rightarrow \infty$, this average is not sensitive to the precise value of t' and a better sampling is achieved by attaching to each clone the average value of \mathcal{O} around time t' [6–8].

3 Biasing the Stationary Distribution: Drift Versus Cloning

Equation (2.4) is nothing but the path-integral representation of the equation:

$$\frac{dP}{dt} = -H_\alpha P \tag{3.1}$$

with $P(\mathbf{x})$ the probability distribution, and:

$$H_\alpha = - \sum_i T_i \frac{\partial^2}{\partial x_i^2} + \sum_i \frac{\partial}{\partial x_i} f_i - \alpha A \tag{3.2}$$

The three terms in H_α correspond to *diffusion*, *drift*, and *cloning*, respectively.

The technique of *dynamic importance sampling* can always be used to reshuffle the importance of drift and cloning. It is implemented by making a change of basis:

$$\tilde{H}_\alpha = e^{\phi(\mathbf{x})} H_\alpha e^{-\phi(\mathbf{x})} = \sum_i -T_i \frac{\partial^2}{\partial x_i^2} + \sum_i \frac{\partial}{\partial x_i} \tilde{f}_i - \tilde{A} \tag{3.3}$$

with:

$$\begin{aligned} \tilde{f}_i &= f_i + 2T_i \frac{\partial \phi}{\partial x_i} \\ \tilde{A} &= \sum_i \left[T_i \left(\frac{\partial \phi}{\partial x_i} \right)^2 + \frac{\partial \phi}{\partial x_i} f_i + T_i \frac{\partial^2 \phi}{\partial x_i^2} \right] + \alpha A \end{aligned} \tag{3.4}$$

In general, there is not an optimal choice for the field ϕ . We will see examples later in different contexts. Another way to understand (3.3) is to consider the dynamics (2.4) with a modified large deviation function:

$$A \rightarrow A + \frac{d\phi}{dt}; \quad F = \int_0^t A(t') dt' + \phi(t) - \phi(0) \tag{3.5}$$

Writing $\frac{d\phi}{dt} = \sum_i \frac{\partial \phi}{\partial x_i} \dot{x}_i$ and expressing \dot{x}_i in terms of the equation of motion, we recover the result (3.3), (3.4). Alternatively, we may of course always consider the modified dynamics as the original one with a cloning rate $A + \frac{d\phi}{dt}$.

Trajectories are thus reweighted according to initial and final configurations. The many-time expectation with respect to the original dynamics $\langle O(t_1)O(t_2) \dots O(t_n) \rangle$ for $t_1 < t_2 < \dots < t_n$, starting from a distribution $P_0(\mathbf{x})$, corresponds to averages with the modified dynamics of $\langle O(t_1)O(t_2) \dots O(t_n)e^{\phi(t_n)} \rangle$, starting from a distribution $e^{\phi} P_0(\mathbf{x})$.

It is important to realize that this is *not* the usual Monte-Carlo importance sampling technique used in equilibrium simulations, which consists simply of modifying the energy in the sampling protocol $E \rightarrow E + B$ (for some suitably chosen B), and compensating by calculating averages as follows:

$$\langle O \rangle_E \rightarrow \langle O e^{\beta B} \rangle_{E+B} \tag{3.6}$$

where $\langle \bullet \rangle_E$ stands for average using a Monte Carlo scheme with energy E . With such a technique, one cannot calculate many-time correlation functions, or trajectory probabilities, since the dynamics are unrelated to the original ones; as one can see easily for the case $B = -E$ where the modified dynamics are simple diffusion, unlike the original ones. In out of equilibrium situations, we do not have an explicit expression for the stationary distribution, and there is no simple way to modify the dynamics in order that they remain probability conserving and have a biased measure, i.e. there is no analog of (3.6).

3.1 Computing Large Moments of Instantaneous Quantities: The Example of Turbulence

It sometimes happens that we are interested in calculating the moments of an *instantaneous* quantity. Consider for example the case of Navier-Stokes equations for driven turbulence. A set of quantities that characterize intermittency are the so-called longitudinal-structure functions [9]

$$S_p(\mathbf{R}) = \langle |\mathbf{v}(\mathbf{x} + \mathbf{R}) - \mathbf{v}(\mathbf{x})|^p \rangle = \langle e^{p \ln |\mathbf{v}(\mathbf{x}+\mathbf{R})-\mathbf{v}(\mathbf{x})|} \rangle \tag{3.7}$$

In order to compute these moments efficiently, we put, in the notation of the previous paragraphs:

$$\phi = \frac{p}{2} \ln |\mathbf{v}(\mathbf{x} + \mathbf{R}) - \mathbf{v}(\mathbf{x})|^2 \tag{3.8}$$

We may run several parallel simulations of fully developed turbulence in the stationary state, each with its own realization of stochastic stirring, and supplement this with a cloning/pruning rate equal to the time-derivative of (3.8), which may be expressed in terms of the instantaneous velocities using the Navier-Stokes equations. The total average cloning rate yields, for large times, $S_p(R)$. Perhaps more interestingly, the configurations that dominate the modified dynamics are the ones that contribute to $S_p(R)$, and are continuously being sampled. To the best of our knowledge, this strategy has not been implemented yet.

4 Transport

We now describe large deviations in non-equilibrium stochastic models of transport. In such models the main observables (e.g. the current, the density, etc.) are functions of the sample path of a Markov chain in a high-dimensional state space.

4.1 Discrete-Time Markov Chains

Imagine a discretization in space of the noisy dynamics (2.1), so that the phase space is given by a finite set of configurations. If we assume that also time is discretized then the dynamics can be described by a Markov chain $\{\mathbf{x}_n\}$ with $(n = 1, 2, \dots, t)$. The evolution is specified by a transition probability matrix whose elements are $p(x, y) = P(\mathbf{x}_{n+1} = y | \mathbf{x}_n = x)$ and by an initial distribution $P(y) = P(\mathbf{x}_0 = y)$. We consider a functional $F[\mathbf{x}_n]$ which is the sum of the local contributions to the current, an additive function of the transitions along the trajectory up to time t :

$$F = F(\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_t) = \sum_{n=1}^t f(\mathbf{x}_{n-1}, \mathbf{x}_n) \quad (4.1)$$

Note that f is, unlike the example in the introduction, a function of the position at *two* successive times. For instance if one considers particles diffusing on a one dimensional lattice and chooses $f(\mathbf{x}_{n-1}, \mathbf{x}_n)$ to be ± 1 depending on whether particles jump to the right or the left, F is the time-integrated current flowing through the system from left to right. The ‘partition function’ (2.4) is given by

$$\begin{aligned} Z_t(\alpha) &= \langle e^{\alpha F(\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_t)} \rangle \\ &= \sum_{x_0, x_1, \dots, x_t} P(x_0) p(x_0, x_1) \cdots p(x_{t-1}, x_t) e^{\alpha f(x_0, x_1)} \cdots e^{\alpha f(x_{t-1}, x_t)} \end{aligned} \quad (4.2)$$

Just as in the previous section, we replace the initial evolution, given by a transition matrix $p(x, y)$, by a new evolution, given by a matrix $p(x, y)e^{\alpha f(x, y)}$. We may decompose this as a probability conserving transition matrix [6]:

$$p_\alpha(x, y) = p(x, y) e^{\alpha f(x, y)} \frac{1}{k(x)} \quad (4.3)$$

and a cloning factor

$$k(x) = \sum_y p(x, y) e^{\alpha f(x, y)} \quad (4.4)$$

We then have

$$Z_t(\alpha) = \sum_{x_0, x_1, \dots, x_{t-1}} P(x_0) p_\alpha(x_0, x_1) \cdots p_\alpha(x_{t-2}, x_{t-1}) k(x_0) \cdots k(x_{t-1}) \tag{4.5}$$

The convenient way to simulate (4.5) is to consider a cloning step of average factor $k(x)$ followed by an evolution step with the transition matrix $p_\alpha(x, y)$. The former may be implemented by substituting a given configuration by a number $(0, 1, 2, \dots)$ of equal clones, with expectation value of the number equal to $k(x)$, while the latter is a transition with probability $p_\alpha(x, y)$.¹ All in all, $\mathcal{N}(n, x)$ —the number of clones of in a configuration x at time n —evolves as

$$\mathcal{N}(n + 1, y) = \sum_x p_\alpha(x, y) k(x) \mathcal{N}(n, x) \tag{4.6}$$

This yields immediately that $Z_t(\alpha)$ is given by the ratio between the average total population at time t and the population at time 0 (at initial time every individual or clone has type distribution $P(x_0)$)

$$Z_t(\alpha) = \frac{\mathcal{N}(t)}{\mathcal{N}(0)} \tag{4.7}$$

To cope with possible extinction or explosion of the initial population one works with increments [6]

$$Z_t(\alpha) = \frac{\mathcal{N}(t)}{\mathcal{N}(t-1)} \frac{\mathcal{N}(t-1)}{\mathcal{N}(t-2)} \cdots \frac{\mathcal{N}(1)}{\mathcal{N}(0)} \tag{4.8}$$

This allows to keep the population size constant during a simulation (with a uniform sampling after the cloning with average factor $k(\cdot)$) and the $Z_t(\alpha)$ will be given by the products of all renormalization factors.

There are many ways of implementing the Diffusion Monte Carlo dynamics described by (4.3) and (4.4), which have been extensively discussed in the literature [10, 11]. For instance, one may choose to run the clones *sequentially*, rather than *simultaneously*, and use any cloning events as the starting point of new simulations [4]. This makes the algorithm easier to parallelize by reducing the overhead but the total number of clones is then harder to control.

4.2 An Example: The Totally Asymmetric Exclusion Process

The Exclusion Process on a lattice consists of particles which jump to their neighboring sites at a given rate, conditioned to the fact that the arrival site is empty. The large deviations of the total particle currents of a periodic chain of N sites with total asymmetry (TASEP) was considered in [6]: in this case only jumps to the right are allowed.

The technique described above amounts to running various independent copies of the chain, but cloning a copy in configuration x with an average rate proportional to

$$k(x) = 1 + \frac{(e^\alpha - 1)}{N} \times [\text{number of particles in } x \text{ with a free site to their right}] \tag{4.9}$$

¹The evolution step can be easily parallelized by splitting the total population of clones over several nodes. The cloning step however creates an overhead since one may have to copy clones from one node to another.

Fig. 1 A shock in the TASEP. Space-time diagram for a ring of $N = 100$ sites, $\alpha = -50/N$ and density 0.5. Time evolution of a single clone. The shock is dense and does not advance. Note the logarithmic scale on the y-axis





Fig. 2 A moving shock in the TASEP. Space-time diagram for a ring of $N = 100$ sites, $\alpha = -30/N$ and density 0.3. The shock drifts to the right



The numerical results obtained for $Z_r(\alpha)$ were compared to the analytic ones of Ref. [12] finding an excellent agreement with a very modest numerical effort. Moreover the algorithm allowed to probe the configurations of the system which are responsible for anomalous small value of the current, the *shocks*, and, in the case of a moving shock, to follow the evolution of the second class particle which set the front of the shock. In Fig. 1 we show a space-time diagram of the system with $N = 100$ particles, density 0.5 and $\alpha = -50/N$. The simulation was done with $L = 1000$ clones, each of them initialized with random (uniform) occupancy numbers, such that the configuration had density 0.5. As predicted by the theory [12] for this value of the density, the shock does not drift, although different initial conditions lead to different shock positions. Figure 2 shows the case $\alpha = -30/N$, and density 0.3: we see that the shock has a net drift to the right, again as predicted by the theory. Let us note here that the configuration corresponds to the end of the time-interval; but one could have sampled one at an intermediate time as explained just below (2.7).

The cloning algorithm has been applied for transport models such as the asymmetric exclusion process and the Kipnis-Marchioro-Presutti model [13–15] and to study symmetries in fluctuations far from equilibrium [16]. Such studies are useful as a test for the predictions of Fluctuating Hydrodynamics [13, 17], but also to probe the limits of the cloning method itself, when insufficient clone number may yield misleading results (a test criterion has been devised in [14]).

4.3 Continuous-Time Markov Chains

Many systems have dynamics that are naturally defined in continuous time. For instance, spin flips in the Ising model, that takes the system from a configuration x to another one y , can occur at any time with a given rate $W(x \rightarrow y)$. To simulate such systems, one can discretize time and choose a small time step dt , (transition probability writing $p(x, y) = dtW(x \rightarrow y)$). One then distinguishes between time steps during which a configuration change occurs (with probability, say, $dtW(x \rightarrow y)$) and those where nothing happens (with probability $1 - dt \sum_y W(x \rightarrow y)$). Doing this in the algorithm described in the previous sections, one arrives in the limit $dt \rightarrow 0$ at a continuous time version of the cloning algorithm.

One can however also work directly with continuous time simulations. Each configuration x has a total escape rate $r(x) = \sum_y W(x \rightarrow y)$, which is the rate at which the system jumps from configuration x to any other configuration. One can choose a time interval δt from an exponential clock, with probability $p(\delta t) = r(x) \exp[-r(x)\delta t]$, update the time $t \rightarrow t + \delta t$, and then decide which configuration changes to make. Going from x to y then occurs with probability $W(x \rightarrow y) / \sum_z W(x \rightarrow z)$. For traditional Monte Carlo algorithms, this method has two advantages. First, one does not have to decide which dt to use and the algorithm makes no discretization error. Second, there are no rejection events which can slow down severely discrete time simulations. However all this comes at the cost of having to generate two random numbers per configuration change (one for the time at which the change occurs, one for the target configuration) while discrete time Monte Carlo only needs one.

When simulating rare events, the continuous time method is more cumbersome to implement but overcomes the problem of diversity of time scales typically met in these simulations. For instance, depending on the value of the bias α , the TASEP presented above explores trajectories where the average time between two events ranges from order 1 (in a traffic jam, only the leading particle can jump forward) to order $1/N$ (when all particles can jump forward). When working with continuous time, the adjustment of the time-step is automatic. In other systems, such as the kinetically constrained models presented in Sect. 5, the situation is even worse. A typical trajectory can explore successive configurations where the waiting times may change by a factor of the order of the system size. In such case, a discrete time algorithm with a time step small enough to resolve the rapid configuration changes will have a prohibitively large number of rejection events when visiting the slow configurations.

To work directly in continuous time, as exposed in [18], the idea is to write the dynamical partition function as a sum over allowed values of F (cf. (2.6)):

$$Z_t(\alpha) = \langle e^{\alpha F} \rangle = \sum_x \underbrace{\sum_F e^{\alpha F} P(x, F, t)}_{\equiv \hat{P}(x, \alpha, t)} \tag{4.10}$$

where $P(x, F, t)$ is the probability density of being in configuration x at time t , and having observed a value F of the dynamical observable. The quantity $\hat{P}(x, \alpha, t)$ is its Laplace transform. As in (4.1), we can choose F to be the sum of contributions $f(x \rightarrow y)$ occurring at each configuration change. For instance, taking $f(x \rightarrow y) = +1$ (resp. -1) each time a particle jumps to the right (resp. left) in a 1d particle system corresponds to F being the total particle flux flowing through the system from right to left. We can also consider the case where F depends on the time average of some observable $A(x)$, as in the introduction

(see [8, 18]):

$$F = \sum_{k=1}^K f(x_{k-1} \rightarrow x_k) + \int_0^t dt' A(x(t')) \tag{4.11}$$

where $(x_0 \dots x_K)$ is the sequence of visited configurations of a given history presenting K changes of configurations. $A(x)$ can for instance be the magnetization of the configuration x of a spin system and one is then looking for trajectories that have atypical *time average* of the magnetization.

From the equation of evolution obeyed by $P(x, F, t)$, one obtains the evolution of $\hat{P}(x, \alpha, t)$:

$$\begin{aligned} \partial_t \hat{P}(x, \alpha, t) &= \sum_y e^{\alpha f(y \rightarrow x)} W(y \rightarrow x) \hat{P}(y, \alpha, t) \\ &\quad - \sum_y W(x \rightarrow y) \hat{P}(x, \alpha, t) + \alpha A(x) \hat{P}(x, \alpha, t) \end{aligned} \tag{4.12}$$

which is of the form $\partial_t |\hat{P}_\alpha\rangle = -H_\alpha |\hat{P}_\alpha\rangle$ where $|\hat{P}_\alpha\rangle$ is the vector of components $\hat{P}(x, \alpha, t)$. Just as in (3.2), the modified operator of evolution H_α does not conserve probability if $\alpha \neq 0$. We have to proceed as in the steps leading to (4.3) and split the evolution in two contributions, one conserving probability and the other a purely cloning term. To do so we introduce the modified transition rates $W_\alpha(y \rightarrow x) = e^{\alpha f(y \rightarrow x)} W(y \rightarrow x)$ and the corresponding escape rate $r_\alpha(x) = \sum_y W_\alpha(x \rightarrow y)$. We can then rewrite (4.12) as

$$\begin{aligned} \partial_t \hat{P}(x, \alpha, t) &= \overbrace{\sum_y W_\alpha(y \rightarrow x) \hat{P}(y, \alpha, t) - r_\alpha(x) \hat{P}(x, \alpha, t)}^{\text{probability conserving}} \\ &\quad + \underbrace{[r_\alpha(x) - r(x) + \alpha A(x)] \hat{P}(x, \alpha, t)}_{\text{cloning}} \end{aligned} \tag{4.13}$$

The first part is a modified dynamics of rates $W_\alpha(y \rightarrow x)$ while the second part corresponds to cloning at rate $r_\alpha(x) - r(x) + \alpha A(x)$. The method is then the same as for discrete time dynamics (Sect. 4.1): one takes a large number of copies of the system, each of them evolving in continuous time (i) through the modified rates $W_\alpha(y \rightarrow x)$ and (ii) subjected to a cloning probability $e^{[r_\alpha(x) - r(x) + \alpha A(x)] \Delta t}$ on each time interval Δt where the configuration does not change from x [8]. One can rescale the total clone population to keep its size constant, storing as previously the overall cloning factor. The dynamical partition function is then recovered from those factors as in (4.8) and the corresponding dynamical free energy $\mu(\alpha)$ is:


$$\mu(\alpha) = \lim_{t \rightarrow \infty} \frac{1}{t} \log Z_t(\alpha) \tag{4.14}$$

We provide in Appendix an example pseudo-code for the practical implementation of the algorithm.

4.4 An Example: Density Profiles in the ASEP

Exclusion processes (such as the TASEP studied above) are interesting transport models in which the cloning algorithms can be used and in particular compared to analytical results

Fig. 3 Density profile in the ASEP. $\alpha = -0.3$, α conjugated to the total current flowing through the system. System size is 400, with 200 particles, in periodic boundary conditions, with an asymmetry $E = \frac{1}{2} \log \frac{p}{q} = -0.2$



for the cumulant generating function $\mu(\alpha) = \lim_{t \rightarrow \infty} \ln Z_t(\alpha)/t$ [13, 18], including finite size effects [8]. In Fig. 3, we present an example of a mean profile at non-zero α for the asymmetric exclusion process (compared to the TASEP, particles can jump to the left and to the right with respective rates p and q). The parameter α is conjugated to the particle flux through the system. We observe on Fig. 3 that, to minimize the overall current, the system develops an asymmetric profile, where only the front particles can jump easily.

5 Fluctuations of Dynamical Activity

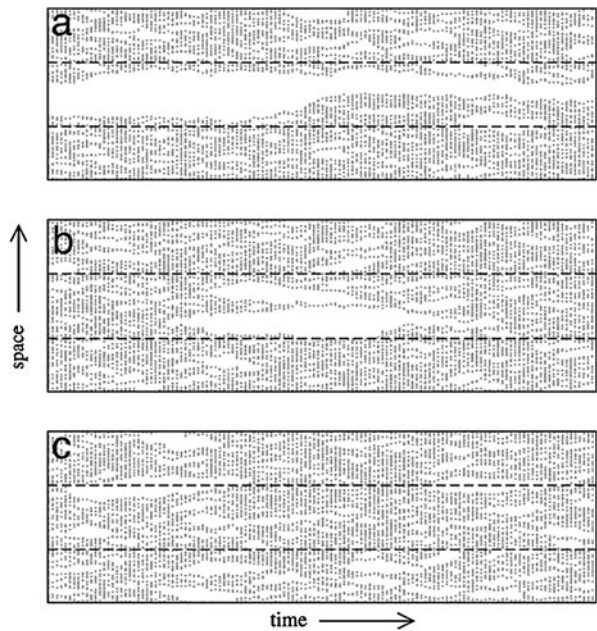
Driven systems may reach a non-equilibrium *steady state*, characterized by a non-zero current the probability distribution of which can be studied as described in the previous section. Another class of non-equilibrium systems is given by glassy systems. In the most simple cases, these systems are out of equilibrium not because they are driven but because their dynamics is so slow that a macroscopic system never reaches Boltzmann equilibrium (or any other steady state), despite the fact that the microscopic dynamics satisfy detailed balance. In this context, it can be interesting to study trajectories of atypical mobility, for instance to detect trajectories that are ‘faster’ or ‘slower’ than average, i.e. the dynamic heterogeneity. To quantify this, one introduces the *dynamical activity* [19–21] (also termed traffic [22, 23]), which provides a good description of dynamical heterogeneity in glass models, as we now discuss.

On a time window $[0, t]$ the dynamical activity K of a stochastic process is the number of configuration changes undergone by the system, and is thus a random variable that depends on the system’s trajectory.

Kinetically constrained models (KCMs), such as the Fredrickson-Andersen [24] or the Kob-Andersen [25] models are such that static (one-time) properties are trivial in the most simple cases, while their dynamical properties (e.g. two- or more times correlations) share common features with generic glassy phenomena (see [26, 27] for reviews on KCMs). They lend themselves rather easily for the study of their activity K , and for the analysis of the results.

Let us focus for simplicity on the one-dimensional Fredrickson-Andersen (FA) model. It consists in a 1d lattice of L sites. Each site is either excited (low density, active) or unexcited (high density, inactive). The sites may flip from inactive to active (at rate c), and from active to inactive (at rate $1 - c$). These transitions are allowed on a given site provided at least

Fig. 4 FA model (from Merolle et al. [19]). Space-time diagram of the FA model for atypical (a and b) and typical (c) histories. In the space direction, active sites are represented in *black dots* while inactive ones are *white*. The picture is reminiscent of the phase coexistence of a static medium at a solid-liquid coexistence point



one of the neighboring sites is active. This is the *kinetic constraint*, introduced as a way to mimic the facilitated dynamics of molecular glasses, whereby active regions enhance activity in their neighborhood. Clearly, for small values of c , the dynamics becomes very slow.

It was observed in [19] that the FA model presents “dynamical coexistence” of active and inactive regions in space-time (see Fig. 4), very similar to the phase coexistence of liquid and solid at the coexistence point in a first order static phase transition—if one forgets that one direction in Fig. 4 is the time.

The activity K of a configuration is defined as the number of active sites. In practice, one may weight the trajectories followed by the system by a factor e^{-sK} , to favor active ($s < 0$) or inactive ($s > 0$) histories (in this section we take the convention $s = -\alpha$ to follow the notation in the literature on KCMs). If the observed coexistence disappears for $s \neq 0$ (that is, if there is a dynamical phase transition), it means that the system indeed sits on a first-order dynamical coexistence point at $s = 0$.

The continuous time cloning algorithm [18] exposed in Sect. 4.3 enables us to compute numerically the dynamical partition function

$$Z_t(s) = \langle e^{-sK} \rangle \sim e^{t\mu_L(s)} \quad (5.1)$$

for this system and other KCMs [7, 28]. The average is taken on histories of duration t , in the large t limit, at fixed system size L . The non-analyticities of the dynamical free energy $\mu_L(s)$ in the large-size limit, signal the existence of a dynamical phase transition.

5.1 Dynamical Phase Coexistence

As shown in [7, 28], several KCMs display a phase transition, in the large system size limit, between an active phase ($s \leq 0$) where the dynamical free energy $\frac{1}{L}\mu_L(s)$ is finite, and an




Fig. 5 FA model (from Garrahan *et al.* [28]). Numerical evaluation of the “dynamical free energy” $\frac{1}{L}\mu_L(s)$ (left) and the density of occupied sites $\rho_L(s)$ (right) for histories weighted by e^{-sK} . As the system size increases, $\frac{1}{L}\mu_L(s)$ approaches its singular limit where $\frac{1}{L}\mu_L(s)$ is zero for $s > 0$. In the same limit, the density $\rho_L(s)$ displays a discontinuity at $s = 0$, indicating a first order dynamical transition

inactive phase ($s > 0$) where is identically zero (see Fig. 5, left, for the 1d FA model). The mean density of active sites (see Appendix for details on the practical computation of such a weighted mean)

$$\rho_L(s) = \frac{\langle e^{-sK} \frac{1}{L} \int_0^t \sum_{i=1}^L n_i \rangle}{\langle e^{-sK} \rangle} \tag{5.2}$$

(here $n_i \in \{0, 1\}$ is the activity at site i) also characterizes this transition (Fig. 5, right): it remains finite in the active phase $s \leq 0$ (a finite fraction of sites is active) while it goes to zero in the inactive phase (only a finite number of sites remains inactive). Several other glass formers display the same phenomenology (see [29] for a review), representative of dynamical heterogeneities, that is, of the coexistence in the system of regions with high and low dynamical activity.

An interesting question is to determine whether molecular models of glasses, such as Lennard Jones mixtures, also present such a dynamical phase transition. A conceptual difficulty that arises is to find a physically relevant measure of the mobility, that generalizes the concept of dynamical activity to this context. In [30], the activity was defined as the number of events where particles move sufficiently far in a given time-interval, thus averaging out short-scale vibrations, whereas in [31], the activity was taken to be a time-average of the modulus of the forces, in a continuous version of the model. In both approaches, numerical results support the existence of a phase transition at some critical value s_c . An open issue is to characterize the inactive phase and to determine whether the effective finite-size critical transition parameter $s_c(L)$ goes to 0 as L goes to infinity or not (that is to say: does the standard dynamics at $s = 0$ lie exactly at the critical point?).

More generically, the phase transitions are also present in p -spin models [32] and in trap models [33], where numerical approach support analytical results. These results are in favor of a generic link between glassiness and dynamical phase coexistence, whose precise nature remains to be understood.

6 Fluctuation of Chaoticity in Dynamical Systems

As explained in the previous section, large deviation theory plays nowadays an important role in non-equilibrium statistical physics to study and quantify dynamical phase transitions.

The first studies of large deviations of dynamical observables were however inspired by another field, that of dynamical systems. It was argued in the 70s, following the seminal works of Sinai, Ruelle, Bowen and others [34–37] that quantitative studies of dynamical systems should rely on a construction analogous to statistical mechanics of trajectory space, where the quantities playing the role of energy functionals for the trajectories are functions of the Lyapunov exponents. This line of thought was very successful in terms of formalism and theory, but progress was severely hampered by the difficulty of computing anything in all but the most schematic systems. Indeed, many of the examples studied very low dimensional systems—mostly maps of the interval, with notable exception of the Lorenz gas [38]. As we show in the two following sections, the development of recent methods to compute the fluctuations of Lyapunov exponents can fill this gap and hopefully lead to new insights in the field of dynamical systems of *many* bodies.

For sake of concreteness, we will focus on Hamiltonian dynamics but one should keep in mind that the method is much more general and can be applied, for instance, to dissipative systems. We consider a system with $2N$ degrees of freedom whose dynamics is given by

$$\dot{\mathbf{x}}_i = f_i[\mathbf{x}(t)]; \quad \text{with} \quad \begin{cases} \mathbf{x} = (q_1, \dots, q_N, p_1, \dots, p_N) \\ \mathbf{f} = (\frac{\partial H}{\partial p_1}, \dots, \frac{\partial H}{\partial p_n}, -\frac{\partial H}{\partial q_1}, \dots, -\frac{\partial H}{\partial q_N}) \end{cases} \quad (6.1)$$

As usual to quantify the chaoticity of a trajectory we introduce the Lyapunov exponents. We consider an infinitesimal perturbation $\delta\mathbf{x}(t)$ whose dynamics reads

$$\delta\dot{\mathbf{x}} = -A \cdot \delta\mathbf{x}; \quad \text{with} \quad A_{ij} = -\frac{\partial f_i[\mathbf{x}(t)]}{\partial x_j} \quad (6.2)$$

The evolution of the norm of such a perturbation is given by

$$\frac{d}{dt} |\delta\mathbf{x}|^2 = -\sum_{ij} 2\delta x_i A_{ij} \delta x_j \quad (6.3)$$

Introducing the normalized tangent vectors $v_i = \frac{\delta x_i}{|\delta\mathbf{x}|}$ whose evolutions are given by

$$\dot{v}_i = -\sum_j A_{ij} v_j + v_i \sum_{kl} v_k A_{kl} v_l \quad (6.4)$$

(6.3) can be recast as

$$\frac{d}{dt} |\delta\mathbf{x}(t)|^2 = -\sum_{ij} 2v_i A_{ij} v_j |\delta\mathbf{x}(t)|^2 \quad (6.5)$$

and finally solved to yield

$$|\delta\mathbf{x}(t)| = |\delta\mathbf{x}(0)| e^{-\sum_{ij} \int_0^t v_i(t') A_{ij}[\mathbf{x}(t')] v_j(t') dt'} \quad (6.6)$$

The largest Lyapunov exponent is then given by $\lambda = \lim_{t \rightarrow \infty} \lambda(t)$, where the finite time Lyapunov exponent $\lambda(t)$ is

$$\lambda(t) = \frac{1}{t} \log \frac{|\delta\mathbf{x}(t)|}{|\delta\mathbf{x}(0)|} = -\frac{1}{t} \int_0^t dt' \left\{ \sum_{ij} v_i(t') A_{ij}[\mathbf{x}(t')] v_j(t') \right\} \quad (6.7)$$

More generally, the exponential expansion of k -dimensional volume elements, rather than vectors $\delta \mathbf{x}$, yields in a similar way the sum of the first k Lyapunov exponents.

To characterize the fluctuations of chaoticity amounts to sampling the distribution of $\lambda(t)$

$$P(\lambda, t) = e^{S(\lambda, t)} \underset{t \rightarrow \infty}{\sim} e^{t s(\lambda)} \tag{6.8}$$

One can understand that the exponent is generically extensive in time, as in usual thermodynamic systems: one cuts a long trajectory of duration t in many segments of duration δt much larger than the typical correlation time τ . Each segment can thus be considered independent of the others and the probability that the total trajectory has an exponent λ is

$$P(\lambda, t) = \sum_{(\lambda_1 + \dots + \lambda_{t/\delta t}) \delta t = \lambda t} P_1(\lambda_1, \delta t) \dots P_{t/\delta t}(\lambda_{t/\delta t}, \delta t) \tag{6.9}$$

$$= \sum_{(\lambda_1 + \dots + \lambda_{t/\delta t}) \delta t = \lambda t} e^{S_1(\lambda_1, \delta t) + \dots + S_{t/\delta t}(\lambda_{t/\delta t}, \delta t)} \tag{6.10}$$

The exponent of each term of the r.h.s. is the sum of $t/\delta t$ terms of order one and is thus of order t . At large times, $t/\delta t \gg 1$, the distribution $P(\lambda, t)$ concentrates around its typical value, and the scaling law (6.8) is thus verified. This scaling breaks down in the presence of diverging correlation times, a signature of dynamical phase transitions.

As in statistical mechanics, the derivation of the entropy $s(\lambda)$ is difficult and one rather works in a ‘‘canonical’’ ensemble by introducing a dynamical partition function

$$Z_t(\alpha) = \langle e^{\alpha t \lambda(t)} \rangle \underset{t \rightarrow \infty}{\sim} e^{t \mu(\alpha)} \tag{6.11}$$

where the average $\langle \cdot \rangle$ is made with respect to $P(\lambda, t)$, i.e. over initial conditions, noise realizations, etc. $\mu(\alpha)$ plays the role of $-\beta F$ in statistical mechanics, where F is a free energy, and is called topological pressure.

From the definition of the finite time Lyapunov exponent (6.7), one sees that the computation of $Z_t(\alpha)$ amounts to the large deviation computation presented in the introduction, with the observable A now given by

$$A(\mathbf{x}) = - \sum_{i,j} v_i A_{ij}(\mathbf{x}) v_j; \quad F = \int dt A(\mathbf{x}) \tag{6.12}$$

Let us now make a point that will be valid for all deterministic systems. In such cases, the only source of fluctuations are the initial conditions. If the system is chaotic enough, this should not be very important but, for example, in the case of mixed system, starting from a regular island or a chaotic region yields a very different result, because trajectories do not take from one to the other. In this review we consider a shortcut to this problem which consists of adding a small amount of stochastic noise, so that the dynamics effectively samples the whole trajectory space (for a discussion of the low noise limit see [40]). We thus consider a slightly different set of equations

$$\dot{q}_i = p_i; \quad \dot{p}_i = - \frac{\partial H}{\partial q_i} + \sqrt{2\epsilon} \eta_i \tag{6.13}$$

The algorithm presented in the introduction of this paper can now be applied to our noisy Hamiltonian dynamics. We consider a population of \mathcal{N} clones in phase space of positions

and momenta \mathbf{q} and \mathbf{p} . To each clone we associate a normalized tangent vector \mathbf{v} . We then choose a time step dt and a noise intensity ϵ and run the simulation over a large time $t = Mdt$. At $t = 0$, the \mathcal{N} copies of the system start from an arbitrary initial configuration (the noise ensures the ergodicity of the algorithm). At each time step $t' = ndt$, we do the following [39]:

- (1) For each clone
 - (\mathbf{q}, \mathbf{p}) evolve with the noisy Hamiltonian dynamics (6.13),
 - \mathbf{v} evolves according to the linearized dynamics

$$\dot{v}_i = -A_{ij}v_j \tag{6.14}$$

- \mathbf{v} is then renormalized to unity and we store the renormalization factor $N(n) = \frac{|v(t+dt)|}{|v(t)|} \simeq e^{-v^T \cdot A \cdot v dt}$.
 - (2) Each clone of the system is then pruned or replicated, with its rate $N(n)^\alpha$. To do so, we pull a random number ϵ uniformly between 0 and 1 and we compute² $\tau = \lfloor \epsilon + N(n)^\alpha \rfloor$,
 - if $\tau = 0$, the clone is deleted,
 - if $\tau > 1$, we create $\tau - 1$ copies of the clone.
 - (3) The total population is now composed of $\mathcal{N}(n + 1)$ clones, instead of the initial $\mathcal{N}(n)$ ones. We then store $R(n) = \frac{\mathcal{N}(n+1)}{\mathcal{N}(n)}$,
 - if $\mathcal{N}(n + 1) < \mathcal{N}(n)$, we copy $\mathcal{N}(n + 1) - \mathcal{N}(n)$ clones, chosen at random,
 - if $\mathcal{N}(n + 1) > \mathcal{N}(n)$, we delete $\mathcal{N}(n + 1) - \mathcal{N}(n)$ clones, chosen at random.
- Finally, we end up again with $\mathcal{N}(n + 1) = \mathcal{N}(n) = \mathcal{N}(0)$ clones.

The dynamical partition function is then obtained from $R(n)$ through

$$Z_t(\alpha) = \prod_{n=1}^M R(n) \tag{6.15}$$

while the topological pressure is given by

$$\mu_t(\alpha) = \frac{1}{t} \sum_{n=1}^M \log R(n) \tag{6.16}$$

Let us now illustrate this algorithm, called “Lyapunov Weighted Dynamics”, with a low dimensional system (the standard map) and a large dimensional one (a FPU chain of 1024 particles).

6.1 The Standard Map

The standard map is defined by the dynamics

$$p_{n+1} = p_n + \frac{k\delta}{2\pi} \sin(2\pi q_n); \quad q_{n+1} = q_n + \delta p_{n+1} \tag{6.17}$$

with $(q_n, p_n) \in [0, 1] \times [-1, 1]$. It is one of the traditional models used to study transition to chaos. It goes from an integrable system when $k = 0$ to a more and more chaotic one when k increases. In Fig. 6 we show the typical trajectories that are localized by the Lyapunov

² $\lfloor x \rfloor$ is the largest integer smaller than x .





Fig. 6 Typical Configurations for $\alpha = \pm 0.04$. Phase space trajectories of the standard map are shown in light gray (light red in the web version) whereas the trajectories localized by the Lyapunov Weighted Dynamics appear in dark gray (dark blue in the web version)

Fig. 7 Standard map.

Dynamical free energy $\mu(\alpha)$ (red crosses) and average Lyapunov exponent $\langle \lambda \rangle_\alpha = Z_t^{-1}(\alpha) \langle \lambda e^{\alpha \lambda t} \rangle$ (blue squares) as a function of the bias α . The discontinuity at $\alpha = 0$ of $\langle \lambda \rangle_\alpha = \mu'(\alpha)$ is the signature of a phase coexistence between chaotic and integrable trajectories in space time



Weighted Dynamics for very small bias ($\alpha = \pm 0.04$). One sees that as soon as the system is biased in favor of integrable trajectories ($\alpha < 0$), the dynamics localizes on integrable islands, whereas a tiny bias favoring chaotic trajectories ($\alpha > 0$) detects the chaotic layers surrounding these islands.

Computing the topological pressure (Fig. 7) shows that the system lies at a critical point where chaotic and integrable trajectories coexist in phase space, in the manner of a first order phase transition.

6.2 FPU Chains

Beyond the computation of dynamical free energies (or topological pressure), the algorithm can be used to sample trajectories of atypical chaoticity. Let us show here on a high-dimensional system, with 2048 degrees of freedom, which are the trajectories that realize large deviations of the chaoticity in anharmonic chains of oscillators. We consider the following Hamiltonian

$$\mathcal{H} = \sum_{i=1}^N \frac{p_i^2}{2} + \sum_{i=1}^N \left[\frac{(x_{i+1} - x_i)^2}{2} + \beta \frac{(x_{i+1} - x_i)^4}{4} \right] \tag{6.18}$$

where $x_{N+1} = x_1$. This system, studied in the 50s by Fermi, Pasta, Tsingou and Ulam, corresponds to N particles connected by anharmonic springs. The limit $\beta = 0$ corresponds to

Fig. 8 Equilibrium simulations of the FPU chain ($N = 128, \alpha = 0$). Time-line of each of the 128 particles around their [arbitrary] equilibrium positions. We see a superposition of localized breathers, ballistic solitons and small fluctuations





Fig. 9 Looking for regular trajectories ($N = 128, \alpha = 5 N$). Simulation at fixed energy ($E = 1$) with fixed boundary conditions, starting from microcanonical equilibrium. The figure shows the time-line of each particles around its [arbitrary] equilibrium position. Several solitons are ballistically propagating from one end of the system to the other, where it elastically bounce of fixed boundary condition. The Lyapunov exponent of this trajectory is equal to half the average one

an integrable case: the springs are harmonic and the Fourier modes correspond to N independent harmonic oscillators or frequencies

$$\omega_k = 2 \sin\left(\frac{\pi k}{N}\right) \tag{6.19}$$

There has been continuous interest in this model (for a review see [41]) because of its rich phenomenology, and in particular, there has been some recent studies of the (Gaussian) fluctuations of its Lyapunov exponent [42]. As soon as β is non-zero, the dynamics are chaotic. However, starting from well chosen initial conditions, the model admits long-lived solitonic modes, related to the Korteweg-de Vries modified equation [43]. Similarly, a modulational instability leads to short-lived chaotic breathers [44, 45], when energy is injected in high-frequency modes. If one runs an equilibrium simulation of the anharmonic chain, one typically observes a mixture of short-lived localized structures (solitons, breathers) and a phonon bath (Fig. 8).

When applying the Lyapunov Weighted Dynamics, we add a small stochastic noise to the system, taking care that the noise conserves the total energy and momentum and thus preventing a slow, unphysical drift in these quantities.

If one biases the system in favor of regular trajectories, the phonons and breathers completely disappear and we observe a long-lived gas of solitons, propagating ballistically (see Fig. 9). In this case, it is important to set the center of mass velocity to zero, because other-





Fig. 10 Looking for chaotic trajectories ($N = 128, \alpha = 5N$). This simulation is realized at fixed energy and total momentum (energy density $E = 1$) with periodic boundary conditions. The *gray level* represent the total energy of each particles. Starting from an equilibrium configuration, the dynamics reveals a chaotic breathers whose Lyapunov exponent is three time larger than the average one

Fig. 11 Looking for chaotic trajectories ($N = 1024, \alpha = 5N$). This simulation is realized at fixed energy and total momentum (energy density $E = 1$) with periodic boundary conditions. The *gray level* represent the total energy of each particles



wise the system can eliminate completely chaoticity by concentrating all its energy on the center of mass motion.

On the other extreme, a bias in favor of chaotic trajectories localizes long-lived chaotic breathers (see Fig. 10). We used periodic boundary conditions for this simulation to reduce the interactions between the wandering breather and the boundaries of the system. Note that running the same simulation in a much larger system ($N = 1024$) shows that the breathers are much more localized than the solitons (Fig. 11).

Interestingly, the values of the bias α we have to use here are not of order one. Indeed, as N increases, the distribution of the largest Lyapunov exponent becomes more and more peaked. Let us assume for instance that $s(\lambda)$ is extensive with some power of the system size, so that one can write

$$P(\lambda_1, t) = \exp[N^\xi t \tilde{s}(\lambda_1)] \tag{6.20}$$

with $\tilde{s}(\lambda)$ of order 1 in both t and N . From the expression

$$Z_t(\alpha) = \langle e^{\alpha \lambda t} \rangle = \int d\lambda \exp[N^\xi t \tilde{s}(\lambda) + \alpha \lambda t] \tag{6.21}$$

one sees that the integral is dominated by a value λ^* such that:

$$\tilde{s}'(\lambda^*) = -\frac{\alpha}{N^\xi} \tag{6.22}$$

When $N \rightarrow \infty$, λ^* satisfies $s'(\lambda^*) = 0$ and is thus the typical value of the Lyapunov exponent. One should thus use a bias that scales as $\alpha = N^\xi \tilde{\alpha}$ to observe large deviations of the Lyapunov exponents. Similarly, to access the dynamical free energy, one has to compute the

exponent ξ and define

$$\tilde{\mu}(\tilde{\alpha}) = \frac{1}{tN^\xi} \log Z_t(\alpha) \quad (6.23)$$

Such a calculation, which, as far as we know, has not been done so far, would tell if the FPU chain lies at a critical point where breathers, solitons and phonons coexist in a first order phase transition manner. The computation of the dynamical free energy for large dimensional systems is now achievable numerically and is one of the exciting goal that are facing us.

7 Work and Entropy Production

When a system is subjected to an external drive, the total energy absorbed (and the resulting entropy production), are quantities that fluctuate depending on the initial microscopic configuration of the system and on the thermal bath, if there is one. Work and entropy production are important quantities, because they concern the state of the system and are the subject of the Second Law of thermodynamics. The Second Law as such concerns only average quantities, and not the fluctuations. It was only relatively recently realized that a wider framework—based on considering the effect of time-reversal on the dynamics—allows to derive a set of relations that are obeyed by the fluctuations—well beyond the linear regime—and yields the Second Law constraints as particular cases.


- (i) The *transient* Fluctuation Theorem relates, in the same context, the probability of a given work W , and that of its opposite: $P(W)/P(-W) = e^{W/T}$ [46, 47].
- (ii) The Jarzynski relation states that the average of $e^{-W/T}$ over all processes starting from an equilibrium distribution at temperature T is *one* [48].

Both are very general, model-independent results, and were later shown to be particular cases of the more general relation, Crooks' relation.

- (iii) The *stationary* fluctuation theorem involves the same relation for the work as the transient version, in a stationary (non-equilibrium) situation, and is valid only in the limit of large times. The particular case in which the dynamics is deterministic (the Gallavotti-Cohen theorem [47]) deserves special attention: the theorem is non trivial because the nature of the stationary distribution is then dependent upon the ergodicity properties of the system. These conditions involve not only chaoticity properties of the attractor, as one would expect from any problem in ergodic theory, but also the fact that attractor and repeller sets are sufficiently intertwined: large deviation trajectories that commute between them generate the reversals in entropy production [49].

Systems with macroscopic, *hydrodynamic* degrees of freedom may have extremely large fluctuations when subjected to strong forcing, due to excitation of macroscopic structures [50]. The typical example is the (Rayleigh-Bénard) convection of a fluid between a hot T_h lower plate and a colder T_c top plate [51]. The heat is transported by fluid currents that have macroscopic fluctuations, enormous compared with $k_b T_h$. The fluctuation theorem as such involves the temperatures T_h, T_c that are irrelevant for these fluctuations. The only way in which the appearance of a Fluctuation Relation for the hydrodynamic modes may be justified, is to invoke the existence of a large *effective temperature*, related to the macroscopic fluctuations. Bonetto and Gallavotti [52] have conjectured that this could be justified by considering the restricted space in which the macroscopic takes place. These questions are very

Fig. 12 The Sinai billiard. The radii are $R_1 = 0.39$, $R_2 = 0.79$. We also show an example of trajectory for the external field $\vec{E} = (1, 0)$



much open, and in order to make progress it would be useful to simulate the limits beyond which the fluctuation theorem ceases to hold rigorously, because that is where new concepts may arise. These are the limits in which large deviations are particularly hard to observe, if one has to wait for them to happen spontaneously.

7.1 Sinai Billiard

The method of cloning has been shown to work efficiently in the verification of the Gallavotti-Cohen theorem on a simple chaotic system given by the Sinai billiard. This system consists of a particle moving inside a billiard as in Fig. 12, with periodic boundary conditions. It is under the action of a force field \vec{E} , and is subject to a deterministic thermostat that keeps the velocity modulus constant $|\vec{v}| = 1$. Between bounces, the equations of motion are:

$$\begin{aligned} \ddot{x}_i &= -E_i + \gamma(t)\dot{x}_i, \quad i = 1, 2; \\ \gamma(t) &= \sum_i E_i \dot{x}_i. \end{aligned} \tag{7.1}$$

We wish to calculate the fluctuations of the dissipated power γ and thus the dynamical partition function

$$Z_t(\alpha) = \langle e^{\alpha \int_0^t \gamma(t') dt'} \rangle \tag{7.2}$$

The fluctuation theorem arises from the symmetry


$$\mu(\alpha) = \mu(-1 - \alpha) \tag{7.3}$$

with $\mu(\alpha) = \lim_{t \rightarrow \infty} \frac{1}{t} \ln Z_t(\alpha)$. Therefore with reference to the notation of the first section we have now

$$A(x) = \gamma(x) \tag{7.4}$$

As in the previous section, the dynamics is deterministic, and hence to allow different clones to diversify, we introduce a small stochastic noise, (cf. paragraph leading to (6.13)) and

Fig. 13 The Gallavotti-Cohen theorem. Plot of $\mu(\lambda)$ vs. λ for the driven Lorentz gas. Data for $\vec{E} = (E, 0)$, $E = 1, 2$ and noise intensity $\Delta = 10^{-3}, 10^{-4}$. The Gallavotti-Cohen theorem implies the symmetry around $\lambda = -1/2$. The continuous lines represents a polynomial fit, quadratic for $E = 1$ (Gaussian behavior), 4-th order for $E = 2$



check the stability of results in the limit of small noise. We evolve the system for macroscopic intervals \mathcal{T} , and clone at time $t' = n\mathcal{T}$ with a factor

$$k_{t'} = e^{\alpha \int_{t'}^{t'+\mathcal{T}} \gamma(t'') dt''}$$

Before each deterministic step of time \mathcal{T} , clones are given random kicks of variance Δ in position and/or velocity direction. The time-interval \mathcal{T} and the noise intensity Δ are chosen so that twin clones have a chance to separate during time \mathcal{T} , and this depends on the chaotic properties of the system. In the present case, $0.1 \leq \mathcal{T} \leq 1$ allows for a few collisions, which guarantees clone diversity for $10^{-3} \leq \Delta \leq 10^{-4}$.

In Fig. 13 we show the results of $\mu(\alpha)$ for $-2 \leq \alpha \leq 1$, and for $\vec{E} = (E, 0)$ with $E = 1$ and $E = 2$, both corresponding to very large current deviations (in the figure α is called λ).

8 Planetary Systems

Planetary systems are the epitome of deterministic systems. With their relatively small number of interacting bodies, they could easily be considered the systems that are further from statistical treatment. And yet, statistical analysis of orbits becomes necessary: when we discover a planetary system we find that many amongst the observationally allowed configurations are only stable in the immediate past or future [53]. Since we do not expect that just by chance we came across a system that has just ejected (or will soon eject) a planet, we tend to favor amongst configurations compatible within error with the data, those that have an unusually high level of stability.

On a related line, it has been shown [54] that just considering a shift in the Earth’s present position of the order of one hundred meters, the fate of Mercury may change dramatically, in some cases leading its orbit to intersect the one of Venus. Consider for example the study by Laskar [54]. In a first calculation, he integrated the orbit of Mercury starting from different configurations, obtained by displacing the position of the earth by about 150 meters. The orbits obtained this way were qualitatively similar, and yet different. Next, he repeated the calculation but making a few clones of the trajectories, and choosing the one with largest eccentricity. After a few such steps, he reached orbits with great eccentricities, that could cross the orbit of Venus. We recognize here a strategy that is very close to the one we are describing here, for the particular cases $\alpha = 0$ and $\alpha = \infty$, respectively. The small displacements are in fact playing the role of our noise.

Indeed, if at each cloning step we had cloned or killed configurations in a fraction proportional to α times the eccentricity change during the corresponding time interval (cf. Sect. 3: the eccentricity plays here the role of ϕ described there) we would have obtained the full probability distribution of, say, the eccentricity at each time. Denote $N(\alpha)$ the total number of clones at time t obtained without normalizing the clone population, or keeping track of the normalizations if they were done. $N(\alpha)$ is the Laplace transform of the probability $P(e)$:

$$N(\alpha) = \int de e^{-\alpha e} P(e) \tag{8.1}$$

Just as in the example of Sinai’s billiard, because the system is chaotic, the displacements (or the noise level), may be essentially negligible—for example, compatible with all other external sources of displacements which we have neglected—and yet yield all the variety of trajectories.

It would be very interesting to see these methods applied to studying in detail the possible future and past evolution of planetary systems, with a large deviation statistical analysis. Many interesting questions concerning the self-organization of the stability of our solar system could be investigated this way.

Appendix: Cloning in Continuous Time: An Example Pseudo-code

In this appendix we provide an example pseudo-code for the cloning of a system described by a configuration `conf`, evolving with Markov dynamics in continuous time (see Sect. 4.3). The dynamics of each clone consists in a succession of (i) Poissonian waiting times (sampled with the function `random.poisson`) between jumps, (ii) change of configuration, or “jumps” (performed by `evolve()`) and (iii) cloning, keeping the total number of clone constant. The way in which the weighted average of a time-extensive observable `obs` is computed is also explicited: a value of `obs` is attached to each clone and copied/pruned with it.

```

alpha=0.1          # parameter conjugated to the observable F
N=500             # number of clones
time=0            # initial time
tmax=1000         # maximum simulation time
cloning=0         # logarithm of the global cloning factor
                  # at the end the ldf is given by cloning/time
conf.init()       # initialization of the clones:
                  # conf[1] to conf[N] are set to given configurations
escaperate.init() # initialization of the alpha-dependent escape rates
obs.init()        # initialize an observable obs that we want to average
                  # over weighted histories

# initialisation of first jump times
for c from 1 to N do:
    jumptime[c]=random.poisson(escaperate[c]) # Poisson law of rate escaperate[c]

# main loop

while t<tmax do
    (c,t)=next(jumptime) # returns the first clone c to jump, and its jumptime t
    conf[c].evolve()     # evolves the configurations clone c
                        # note that the observable obs is evolved accordingly
    deltaT=random.poisson(escaperate[c])
                        # determines the time interval until the next jump
    jumptime[c]+=deltaT  # updates the jumptime
    
```



```

K=conf[c].clfact(deltaT)# yields the cloning factor
                        #  $K=e^{(\text{deltaT}*(\text{deltaescaperate}[c]+\text{alpha}*A[c]))}$ 
cloning+=log((N+K-1)/N) # updates the log of the global cloning factor

k=floor(K+random.real())# integer number k representing the number of clones
                        # replacing the current clone c

cases
  k=0: # clone c is suppressed, i.e. replaced by another one chosen at random
        do newc=random.integer(N) while newc==c
        conf[c]=conf[newc]
        obs[c]=obs[newc]
        jumptime[c]=jumptime[newc]
  k=1: # nothing is done
  k>1: # k-1 copies of c have to be done; then, among the total N+k-1
        # resulting clones, k-1 of them are pruned so as to keep N constant
        indices=randomarray(N,k)
        # puts in indices k-1 *different* random integers between 1 and
        # Nclones+k-1
        # (both included); only those less or equal than N will be replaced by c
        for newc in indices do:
          if newc<=N do:
            conf[newc]=conf[c]
            obs[newc]=obs[c]
            jumptime[newc]=jumptime[c]

# output of results
ldf=cloning/time
print('large deviation function = ',ldf)
meanobs=sum(obs[c] for c in range(N))/N/time
print('weighted mean of observable = ',meanobs)

```

References

1. Touchette, H.: *Phys. Rep.* **478**, 1 (2009)
2. Anderson, J.B.: *J. Chem. Phys.* **63**, 1499 (1975)
3. Aldous, D., Vazirani, U.: In: *Proc. 35th IEEE Sympos. on Foundations of Computer Science* (1994)
4. Grassberger, P.: *Comput. Phys. Commun.* **147**, 64–70 (2002)
5. Del Moral, P., Doucet, A., Jasra, A.: *J. R. Stat. Soc., Ser. B, Stat. Methodol.* **68**, 411–436 (2006)
6. Giardina, C., Kurchan, J., Peliti, L.: *Phys. Rev. Lett.* **96**, 120603 (2006)
7. Garrahan, J.-P., Jack, R.L., Lecomte, V., Pitard, E., van Duijvendijk, K., van Wijland, F.: *J. Phys. A, Math. Theor.* **42**, 075007 (2009)
8. Tailleur, J., Lecomte, V.: In: *Modeling and Simulation of New Materials. AIP Conf. Proc.*, vol. 1091, p. 212 (2009)
9. Frisch, U.: *Turbulence: The Legacy of A.N. Kolmogorov*. Cambridge University Press, Cambridge (1995)
10. El Makrini, M., Jourdain, B., Lelièvre, T.: *Diffusion Monte Carlo method: numerical analysis in a simple case. ESAIM, Math. Model. Numer. Anal.* **41**, 189 (2007)
11. Cochran, W.: *Sampling Techniques*. Wiley Eastern, New Delhi (2007)
12. Bodineau, T., Derrida, B.: *Phys. Rev. E* **72**, 066110 (2005)
13. Hurtado, P.I., Garrido, P.L.: *Phys. Rev. Lett.* **102**, 250601 (2009)
14. Hurtado, P.I., Garrido, P.L.: *J. Stat. Mech.*, P02032 (2009)
15. Hurtado, P.I., Garrido, P.L.: *Phys. Rev. E* **81**, 041102 (2010)
16. Hurtado, P.I., Pérez-Espigares, C., del Pozo, J.J., Garrido, P.L.: *Proc. Natl. Acad. Sci. USA* **108**, 7704 (2011)
17. Hurtado, P.I., Garrido, P.L.: [arXiv:1106.0690](https://arxiv.org/abs/1106.0690)
18. Lecomte, V., Tailleur, J.: *J. Stat. Mech.*, P03004 (2007)
19. Merolle, M., Garrahan, J.-P., Chandler, D.: *Proc. Natl. Acad. Sci. USA* **102**, 10837 (2005)
20. Lecomte, V., Appert-Rolland, C., van Wijland, F.: *C. R. Phys.* **8**, 609 (2007)
21. Bodineau, T., Lefevere, R.: *J. Stat. Phys.* **133**, 1 (2008)
22. Maes, C., Netocný, K., Wynants, B.: *Markov Process. Relat. Fields* **14**, 445 (2008)
23. Maes, C., Netocný, K.: *Europhys. Lett.* **82**, 30003 (2008)
24. Fredrickson, G.H., Andersen, H.C.: *Phys. Rev. Lett.* **53**, 1244 (1984)

25. Kob, W., Andersen, H.C.: *Phys. Rev. E* **48**, 4364 (1993)
26. Ritort, F., Sollich, P.: *Adv. Phys.* **52**, 219 (2003)
27. Garrahan, J.-P., Sollich, P., Toninelli, C.: In: *Dynamical Heterogeneities in Glasses, Colloids and Granular Materials*. Oxford University Press, London (2011)
28. Garrahan, J.-P., Jack, R.L., Lecomte, V., Pitard, E., van Duijvendijk, K., van Wijland, F.: *Phys. Rev. Lett.* **98**, 195702 (2007)
29. Chandler, D., Garrahan, J.-P.: *Annu. Rev. Phys. Chem.* **61**, 191 (2010)
30. Hedges, L.O., Jack, R.L., Garrahan, J.-P., Chandler, D.: *Science* **323**, 1309 (2009)
31. Pitard, E., Lecomte, V., Van Wijland, F.: [arXiv:1105.2460](https://arxiv.org/abs/1105.2460) (2011)
32. van Duijvendijk, K., Jack, R.L., van Wijland, F.: *Phys. Rev. E* **81**, 011110 (2010)
33. van Duijvendijk, K., Schehr, G., van Wijland, F.: *Phys. Rev. E* **78**, 011120 (2008)
34. Sinai, Y.G.: *Russ. Math. Surv.* **27**, 21 (1972)
35. Bowen, R.: *Equilibrium States and the Ergodic Theory of Anosov Diffeomorphism*. Lecture Notes in Math., vol. 470. Springer, Berlin (1975)
36. Ruelle, D.: *Invent. Math.* **34**, 231 (1976)
37. Ruelle, D.: *Thermodynamic Formalism*. Addison-Wesley, Reading (1978)
38. van Beijeren, H., Mülken, O.: *Phys. Rev. E* **71**, 036213 (2005)
39. Tailleur, J., Kurchan, J.: *Nat. Phys.* **3**, 203–207 (2007)
40. Kurchan, J.: *J. Stat. Phys.* **128**, 1307 (2007)
41. Berman, G.P., Izrailev, F.M.: *Chaos* **15**, 015104 (2005)
42. Kuptsov, P.V., Politi, A.: [arXiv:1102.3141](https://arxiv.org/abs/1102.3141) (2011)
43. Kruskal, M.D., Zabusky, N.J.: *J. Math. Phys.* **5**, 231 (1964)
44. Cretegnny, T., Dauxois, T., Ruffo, S., Torcini, A.: *Physica D* **121**, 109–126 (1998)
45. Trombettoni, A., Smerzi, A.: *Phys. Rev. Lett.* **86**, 2353–2356 (2001)
46. Evans, D.J., Cohen, E.G.D., Morriss, G.P.: *Phys. Rev. Lett.* **71**, 2401 (1993)
47. Gallavotti, G., Cohen, E.G.D.: *Phys. Rev. Lett.* **74**, 2694 (1995)
48. Jarzynski, C.: *Phys. Rev. Lett.* **78**, 2690 (1997)
49. Kurchan, J.: *J. Stat. Phys.* **128**, 1307 (2007)
50. Portelli, B., Holdsworth, P.C.W., Pinton, J.F.: *Phys. Rev. Lett.* **90**, 104501 (2003)
51. Ciliberto, S., Laroche, C.: *J. Phys. (Paris)* **8**, 215–219 (1998)
52. Bonetto, F., Gallavotti, G.: *Commun. Math. Phys.* **189**, 263 (1997)
53. Barnes, R., Quinn, T.: *Astrophys. J.* **611**, 494 (2004)
54. Laskar, J.: *Astron. Astrophys.* **287**, L9 (1994)