

Simulation of large deviation functions using population dynamics

Julien Tailleur* and Vivien Lecomte[†]

*SUPA, School of Physics and Astronomy, University of Edinburgh, Mayfield Road, Edinburgh EH9 3JZ, Scotland

[†]DPMC, Université de Genève, 24, Quai Ernest Ansermet 1211 Genève, Switzerland

Abstract. In these notes we present a pedagogical account of the population dynamics methods recently introduced to simulate large deviation functions of dynamical observables in and out of equilibrium. After a brief introduction on large deviation functions and their simulations, we review the method of Giardinà *et al.* for discrete time processes and that of Lecomte *et al.* for the continuous time counterpart. Last we explain how these methods can be modified to handle static observables and extract information about intermediate times.

PACS: 05.40.-a, 05.70.Ln, 05.60.-k

The main achievement of equilibrium statistical mechanics is probably the simplification it offers in the study of static observables in steady state. Indeed, the resolution of dynamical equations is replaced by static averages and ensemble approaches. This can still be very difficult, but from a conceptual point of view, the problem is much simpler. On the other hand, when one is interested in dynamical observables, like currents of particles, or in out-of-equilibrium situations, like for glassy or driven systems, such static ensemble approaches are not available anymore and there is no general formalism on which one can rely.

For the last ten years, physicists have been interested in large deviation functions mainly because they are good candidates to extend the concept of thermodynamic potentials to out-of-equilibrium situations and to dynamical observables (for a review, see [1]). Of course the mere definition of out-of-equilibrium potentials is not useful in itself and the challenge is thus to go beyond their construction. To do so, two strategies can be followed. First, one can try to derive general properties of large deviation functions. An example of success in this direction is provided by the Fluctuation Theorem [2, 3, 4, 5], which can be read as a symmetry of large deviation functions and is one of the few general results valid out-of-equilibrium. Another strategy is to consider specific examples and to compute the large deviation function explicitly. For some simple yet non-trivial interacting particle systems, exact computations have been possible (for a review see [6]), but one has to rely on numerics for more generic systems. From the algorithmic point of view, two paths can be followed. For small system sizes, exact procedures can be used (see for instance [8, 9]) but as soon as mesoscopic systems are considered, one has to rely on importance sampling approaches. Generalising a procedure followed previously to study rare events in chemical reactions [10, 11, 12], Kurchan and co-workers developed methods to compute large deviation functions in dynamical systems [13], and discrete [14] or continuous [15] time Markov chains. We shall concentrate here on the

statistical mechanical aspects and review the methods available for Markov chains.

Let us consider a system and call $\{\mathcal{C}\}$ its set of configurations. As time goes from 0 to a final time t , the system typically jumps into successive configurations $\mathcal{C}_0, \mathcal{C}_1, \dots, \mathcal{C}_K$ at distinct times t_1, \dots, t_K . A *dynamical* observable Q is defined as a sum along the history of small contributions $q_{\mathcal{C}_{k+1}\mathcal{C}_k}$, one for each *transition* between two successive configurations. The simplest example of such observable is probably a current of particles Q in a one-dimensional lattice gas, which is either incremented or decremented every time a particle jumps to the right or to the left, respectively. This contrasts with static observables, like the number of particles at a given site, which depend solely on the configuration of the system at a given time. We will see below that the algorithms used to obtain the large deviation functions slightly differ in these two cases. To characterise the fluctuations of the observable Q , the first thing one can do is to extend the microcanonical approach to the space of trajectories and compute the corresponding macrostate entropy

$$s(q) = \lim_{t \rightarrow \infty} \frac{1}{t} \log P[Q(t) = qt]. \quad (1)$$

However one knows from usual statistical mechanics that working in the microcanonical ensemble is often harder than in the canonical one and we rather introduce a dynamical partition function and the corresponding dynamical free energy

$$Z(\beta, t) = \left\langle e^{-\beta Q(t)} \right\rangle; \quad \psi(\beta) = \lim_{t \rightarrow \infty} \frac{1}{t} \log Z(\beta, t). \quad (2)$$

These definitions differ slightly from those used in the mathematical literature, where one rather speaks about rate functions $-s(q)$ and cumulants generating functions $\psi(-\beta)$.

The main purpose of this short review is to explain how one can compute $\psi(\beta)$ using an approach relying on population dynamics. But let us first sketch why direct sampling would be inefficient. Consider for simplicity the case where $s(q)$ has a single maximum at q_0 , which satisfies $s(q_0) = 0$ (for normalisation purpose). Fluctuations around $Q = q_0 t$ which occur with probabilities of order one must typically be of order $1/\sqrt{t}$ so that

$$P(Q = qt) \simeq e^{ts(q)} \simeq e^{\frac{1}{2}t(q-q_0)^2 s''(q_0)} \sim 1, \quad (3)$$

whence a probability of larger fluctuations exponentially small in t . On the other hand, the dynamical partition function has a weight $e^{-\beta Q}$ (with $Q \sim qt$) exponential in t . As a result, there is a competition between the exponentially rare fluctuations and their exponential weight such that for β of order 1 the trajectories which dominate the average in (2) are exponentially rare. A more quantitative way to rephrase this can be read in the Legendre relation between $s(q)$ and $\psi(\beta)$: $\psi(\beta) = \max_q [s(q) - \beta q]$. The maximum is realised for a value q^* which dominates the average in (2). It thus differs from q_0 – which maximises only $s(q)$ – by a factor independent of t . From (3) we see that the corresponding trajectories indeed have a probability exponentially small in t . To have a good sampling over N unbiased simulations, N should thus be of order e^t , an impracticable requirement. Direct sampling is thus a hopeless strategy to observe large deviations.

DISCRETE TIME

We present in this section the method first introduced by Giardinà, Kurchan and Peliti [14] to simulate cumulant generating functions in discrete time Markov chains. In this case, the dynamics is defined by the transition *probabilities* $U(\mathcal{C} \rightarrow \mathcal{C}') = U_{\mathcal{C}'\mathcal{C}}$ between configurations and the corresponding the master equation reads

$$P(\mathcal{C}, t) = \sum_{\mathcal{C}'} U_{\mathcal{C}\mathcal{C}'} P(\mathcal{C}', t-1). \quad (4)$$

Conservation of probability enforces the matrix U to be stochastic, i.e. for all \mathcal{C}' , $\sum_{\mathcal{C}} U_{\mathcal{C}\mathcal{C}'} = 1$. Starting from a fixed configuration \mathcal{C}_0 the explicit solution of (4) is

$$P(\mathcal{C}, t) = \sum_{\mathcal{C}_1 \dots \mathcal{C}_{t-1}} U_{\mathcal{C}\mathcal{C}_{t-1}} U_{\mathcal{C}_{t-1}\mathcal{C}_{t-2}} \dots U_{\mathcal{C}_1\mathcal{C}_0} = [U^t]_{\mathcal{C}\mathcal{C}_0}. \quad (5)$$

The dynamical observable Q can be written as a sum over configuration changes $Q(t) = q_{\mathcal{C}_t\mathcal{C}_{t-1}} + \dots + q_{\mathcal{C}_1\mathcal{C}_0}$. To compute the dynamical free energy $\psi(\beta)$, we first rewrite the dynamical partition function as

$$Z(\beta, t) = \left\langle e^{-\beta Q(t)} \right\rangle = \sum_{\mathcal{C}_1 \dots \mathcal{C}_t} U_{\mathcal{C}_t\mathcal{C}_{t-1}} e^{-\beta q_{\mathcal{C}_t\mathcal{C}_{t-1}}} \dots U_{\mathcal{C}_1\mathcal{C}_0} e^{-\beta q_{\mathcal{C}_1\mathcal{C}_0}} = \sum_{\mathcal{C}} [U_\beta^t]_{\mathcal{C}\mathcal{C}_0}, \quad (6)$$

where we have introduced the matrix $[U_\beta]_{\mathcal{C}\mathcal{C}'} = U_{\mathcal{C}\mathcal{C}'} e^{-\beta q_{\mathcal{C}\mathcal{C}'}}$. Let us note that $\psi(\beta)$ is given by the log of the largest eigenvalue of U_β . A possible strategy, used for instance in [7, 8, 9], is thus to compute numerically this eigenvalue. The matrix U_β is however exponentially large in the system size, which limits this strategy to small systems. The main advantage of this method is to yield a numerical approximation of an exact expression - as pointed out in [9] - as opposed to our approach, efficient for large systems, but relying on importance sampling. Note that the ‘exact’ approach can be used to check the validity of the importance sampling approach for small system sizes, before going to larger ones, as was actually done for the simulations presented in [15].

Comparison of expressions (5) and (6) leads one to think that $\psi(\beta)$ could be obtained from a new dynamics, induced by U_β . However, the matrix U_β is not stochastic, as in general

$$Y_{\mathcal{C}'} \equiv \sum_{\mathcal{C}} [U_\beta]_{\mathcal{C}\mathcal{C}'} = \sum_{\mathcal{C}} U_{\mathcal{C}\mathcal{C}'} e^{-\beta q_{\mathcal{C}\mathcal{C}'}} \neq 1 \quad (7)$$

and we should not understand (6) as a stochastic process with conserved probability but as a population dynamics with branching and death, where the population size is not constant. To do so, let us define

$$U'_{\mathcal{C}\mathcal{C}'} = \frac{[U_\beta]_{\mathcal{C}\mathcal{C}'}}{Y_{\mathcal{C}'}} = \frac{U_{\mathcal{C}\mathcal{C}'}}{Y_{\mathcal{C}'}} e^{-\beta q_{\mathcal{C}\mathcal{C}'}}. \quad (8)$$

The matrix U' is stochastic and (6) now writes

$$\left\langle e^{-\beta Q(t)} \right\rangle = \sum_{\mathcal{C}_1 \dots \mathcal{C}_t} U'_{\mathcal{C}_t\mathcal{C}_{t-1}} Y_{\mathcal{C}_{t-1}} \dots U'_{\mathcal{C}_1\mathcal{C}_0} Y_{\mathcal{C}_0}. \quad (9)$$

This expression is now closer to (5) as U' is stochastic, and can be interpreted as follows: N agents evolve with the stochastic dynamics defined by U' and are replicated with a rate $Y_{\mathcal{C}}$ when they are in configuration \mathcal{C} . This interpretation is possible as $Y_{\mathcal{C}}$ depends solely on the initial configuration \mathcal{C} and can thus be interpreted as a configuration-dependent reproduction rate. This was less apparent in (6), where factors $e^{-\beta q_{\mathcal{C}\mathcal{C}'}}$ depend on both initial and final configurations.

This interpretation as a population dynamics can be implemented using a diffusion Monte Carlo algorithm. Let us consider an ensemble of N_0 agents ($N_0 \gg 1$) evolving in the configuration space $\{\mathcal{C}\}$. At each time step $\tau \rightarrow \tau + 1$,

- (1) Each agent evolves according to the β -modified dynamics $U'_{\mathcal{C}\mathcal{C}'}$,
- (2) Each agent in configuration \mathcal{C} is replicated/killed with probability $Y_{\mathcal{C}}$, *i.e.* is replaced by y copies, where

$$y = \begin{cases} \lfloor Y_{\mathcal{C}} \rfloor + 1 & \text{with probability } Y_{\mathcal{C}} - \lfloor Y_{\mathcal{C}} \rfloor \\ \lfloor Y_{\mathcal{C}} \rfloor & \text{with probability } 1 - (Y_{\mathcal{C}} - \lfloor Y_{\mathcal{C}} \rfloor) \end{cases} \quad (10)$$

Concretely, the agent is replaced by y copies of itself, so that the population size is increased by $y - 1$ (decreased by 1 if $y = 0$).

Let us show that the size of the population at time t yields the large deviation function. For a given history, the number $N(\mathcal{C}, \tau)$ of copies in configuration \mathcal{C} at intermediate time τ satisfies $N(\mathcal{C}_\tau, \tau) = U'_{\mathcal{C}_\tau \mathcal{C}_{\tau-1}} Y_{\mathcal{C}_{\tau-1}} N(\mathcal{C}_{\tau-1}, \tau - 1)$, so that for the whole history

$$N(\mathcal{C}_t, t) = U'_{\mathcal{C}_t \mathcal{C}_{t-1}} Y_{\mathcal{C}_{t-1}} \dots U'_{\mathcal{C}_1 \mathcal{C}_0} Y_{\mathcal{C}_0} N(\mathcal{C}_0, 0). \quad (11)$$

Consequently, we see from (9) that the population size $N(t) = \sum_{\mathcal{C}} N(\mathcal{C}, t)$ behaves as $N(t)/N_0 = \langle e^{-\beta Q(t)} \rangle \sim e^{t\psi(\beta)}$. As usual in importance sampling approaches, the ensemble average $\langle \cdot \rangle$ has been replaced by an average over a finite number of simulations. Whereas in principle correct, this approach is however impracticable since the population size varies exponentially in time and we thus add a third step to the previous algorithm:

- (3) After the cloning step, the population is rescaled by a factor X_τ to its initial size N_0 , by uniformly pruning/replicating the agents.

At each time step τ , the rescaling factor is given by $X_\tau = \frac{N(\tau-1)}{N(\tau)}$ so that $X_t \dots X_0 = \frac{N_0}{N(t)}$ and finally

$$\psi(\beta) = - \lim_{t \rightarrow \infty} \frac{1}{t} \log \langle X_t \dots X_0 \rangle. \quad (12)$$

CONTINUOUS TIME DYNAMICS

For a continuous time dynamics defined by rates $W(\mathcal{C} \rightarrow \mathcal{C}')$ the master equation reads

$$\partial_t P(\mathcal{C}, t) = \sum_{\mathcal{C}' \neq \mathcal{C}} W(\mathcal{C}' \rightarrow \mathcal{C}) P(\mathcal{C}', t) - r(\mathcal{C}) P(\mathcal{C}, t), \quad (13)$$

where $r(\mathcal{C}) = \sum_{\mathcal{C}'} W(\mathcal{C} \rightarrow \mathcal{C}')$ is the escape rate from configuration \mathcal{C} .

There are many different ways of deriving the algorithm presented in the previous section and we shall follow here a derivation of the continuous time algorithm slightly different from the one we introduced in [15]. The formal solution of (13) reads

$$P(\mathcal{C}, t) = \sum_{K \geq 0} \sum_{\mathcal{C}_1 \dots \mathcal{C}_{K-1}} \int_{t_0}^t dt_K \int_{t_0}^{t_K} dt_{K-1} \dots \int_{t_0}^{t_2} dt_1 \quad (14)$$

$$\rho(t_K | \mathcal{C}_{K-1}, t_{K-1}) \dots \rho(t_1 | \mathcal{C}_0, t_0) e^{-(t-t_K)r(\mathcal{C})} \frac{W(\mathcal{C}_0 \rightarrow \mathcal{C}_1)}{r(\mathcal{C}_0)} \dots \frac{W(\mathcal{C}_{K-1} \rightarrow \mathcal{C})}{r(\mathcal{C}_{K-1})},$$

where $\rho(t_k | \mathcal{C}_{k-1}, t_{k-1}) = r(\mathcal{C}_{k-1}) \exp[-(t_k - t_{k-1})r(\mathcal{C}_{k-1})]$ represents the probability distribution of the time intervals between jumps. The sum over K corresponds to all the possible numbers of jumps between 0 and t , the sum over the \mathcal{C}_k 's to the different configurations which can be visited. The integrals over t_k account for all the possible times at which jumps occur. Last, the ratio $\frac{W(\mathcal{C}_{k-1} \rightarrow \mathcal{C}_k)}{r(\mathcal{C}_{k-1})}$ gives the *probability* that the system goes to configuration \mathcal{C}_k , when it quits configuration \mathcal{C}_{k-1} . Multiplying the second line of (14) by $e^{-\beta Q}$ yields an explicit formula for $Z(\beta, t)$:

$$Z(\beta, t) = \sum_{K \geq 0} \sum_{\mathcal{C}_1 \dots \mathcal{C}_{K-1}, \mathcal{C}} \int_{t_0}^t dt_K \int_{t_0}^{t_K} dt_{K-1} \dots \int_{t_0}^{t_2} dt_1 \quad (15)$$

$$\rho(t_K | \mathcal{C}_{K-1}, t_{K-1}) \dots \rho(t_1 | \mathcal{C}_0, t_0) e^{-(t-t_K)r(\mathcal{C})}$$

$$\frac{W(\mathcal{C}_0 \rightarrow \mathcal{C}_1)}{r(\mathcal{C}_0)} e^{-\beta q_{\mathcal{C}_1 \mathcal{C}_0}} \dots \frac{W(\mathcal{C}_{K-1} \rightarrow \mathcal{C})}{r(\mathcal{C}_{K-1})} e^{-\beta q_{\mathcal{C} \mathcal{C}_{K-1}}}.$$

Further introducing the biased rates $W_\beta(\mathcal{C} \rightarrow \mathcal{C}') = W(\mathcal{C} \rightarrow \mathcal{C}') e^{-\beta q_{\mathcal{C}' \mathcal{C}}}$, the corresponding escape rates r_β and time distributions between two jumps ρ_β , (15) can be rewritten (after some algebra) as

$$Z(\beta, t) = \sum_{K \geq 0} \sum_{\mathcal{C}_1 \dots \mathcal{C}_{K-1}, \mathcal{C}} \int_{t_0}^t dt_K \int_{t_0}^{t_K} dt_{K-1} \dots \int_{t_0}^{t_2} dt_1 \quad (16)$$

$$\rho_\beta(t_K | \mathcal{C}_{K-1}, t_{K-1}) \dots \rho_\beta(t_1 | \mathcal{C}_0, t_0) e^{-(t-t_K)r_\beta(\mathcal{C})}$$

$$Y(\mathcal{C}_0)^{t_1-t_0} \frac{W_\beta(\mathcal{C}_0 \rightarrow \mathcal{C}_1)}{r_\beta(\mathcal{C}_0)} \dots Y(\mathcal{C}_{K-1})^{t_K-t_{K-1}} \frac{W_\beta(\mathcal{C}_{K-1} \rightarrow \mathcal{C})}{r_\beta(\mathcal{C}_{K-1})} Y(\mathcal{C})^{t-t_K},$$

where $Y(\mathcal{C}_k) = e^{r_\beta(\mathcal{C}_k) - r(\mathcal{C}_k)}$. $Z(\beta, t)$ is thus a weighted sum over all possible trajectories generated by the biased rates W_β , where the weights are given by the factors $Y(\mathcal{C}_k)^{t_{k+1}-t_k}$. A first idea which can come to mind is to simply evolve the population with the rates W_β , without cloning and to simply average $e^{\int_0^t d\tau (r_\beta(\mathcal{C}(\tau)) - r(\mathcal{C}(\tau)))}$ over these trajectories, as was proposed in [16]. This however fails as soon as t is large, for the same reason as the one described in the introduction: the weight is exponentially large in t and large fluctuations of the exponent are exponentially rare. One thus has to use a biased sampling to compute the average (16). Following the philosophy of the ‘‘Go with the Winner methods’’ [19], the general idea is to stochastically replace a trajectory with weight \mathcal{W} by \mathcal{W}' trajectories with weight 1, so that trajectories which high rates are favoured whereas those with small weights are not investigated.

If the re-weighting procedure is made systematic, every time an agent c_α changes of configuration at time t^α , one gets the following algorithm:

- (0) The time is set to t^α .

- (1) c_α jumps from its configuration \mathcal{C} to another configuration \mathcal{C}' with probability $W_\beta(\mathcal{C} \rightarrow \mathcal{C}')/r_\beta(\mathcal{C})$.
- (2) The time interval Δt until the next jump of c_α is chosen from the Poisson law ρ_β of parameter $r_\beta(\mathcal{C}')$.
- (3) The agent c_α is either cloned or pruned with a rate $\mathcal{Y}(\mathcal{C}') = e^{\Delta t(r_\beta(\mathcal{C}') - r(\mathcal{C}'))}$
 - a) One computes $y = \lfloor \mathcal{Y}(\mathcal{C}') + \varepsilon \rfloor$ where ε is uniformly distributed on $[0, 1]$.
 - b) If $y = 0$, the copy c_α is erased.
 - c) If $y > 1$, we make $y - 1$ new copies of c_α .
- (4) If $y = 0$, one agent $c_\beta \neq c_\alpha$ is chosen at random and copied, while if $y > 1$, $y - 1$ agents are chosen uniformly among the $N + y - 1$ agents and erased. We store the rescaling factor $X = \frac{N}{N+y-1}$.

To reconstruct the dynamical free energy, we keep track of all X factors

$$\frac{1}{t} \log \langle X_1 \dots X_t \rangle = \frac{1}{t} \log \langle e^{-\beta Q(t)} \rangle \sim -\psi(\beta) \quad \text{as } t \rightarrow \infty \quad (17)$$

Once again the step (4) ensures constant population.

THE CASE OF STATIC OBSERVABLES

The methods presented above only apply for dynamical observables, which can be decomposed as sums of individual contributions over each configuration change. One could also be interested in averages of static observables along the histories $O = \int_0^t d\tau o(\tau) = \sum_k (t_{k+1} - t_k) o(\mathcal{C}_k)$. In this case, the above procedure simplifies and the algorithm is identical apart from two points.

- First, there is no bias in the rate. The agents are evolved with the unmodified Markov rates $W(\mathcal{C} \rightarrow \mathcal{C}')$.
- Second, the cloning rate is simply given by $e^{-\beta(t_{k+1} - t_k) o(\mathcal{C}_k)}$.

This can best be seen in formula (15) by replacing the weight $e^{-\beta q_{\mathcal{C}_{k+1} \mathcal{C}_k}}$ which depends on the configurations before and after the jump by $e^{-\beta(t_{k+1} - t_k) o(\mathcal{C}_k)}$ which depends solely on the configuration \mathcal{C}_k and can thus be seen as a constant cloning rate for the whole time the system spends in the configuration \mathcal{C}_k .

INTERMEDIATE TIMES

As pointed out in [14], the configurations probed along the simulation are representative of the typical ones at *final time* t in the evolution, rather than at intermediate times ($0 \ll \tau \ll t$). In particular, the weighted average value $\langle O(t) \rangle_\beta$ of a static observable $O(\mathcal{C}(t))$ at final time t is obtained in the algorithm by computing the average of O among the agents at the end of the simulation.

In general, the value of $\langle O(\tau) \rangle_\beta$ at intermediate times $0 \ll \tau \ll t$ differs from the one at final time, yet $\langle O(\tau) \rangle_\beta$ is of particular interest since it is representative of

configurations visited during most of the weighted evolution (see [20] for examples). A way to compute $\langle O(\tau) \rangle_\beta$ numerically can be read in its formal expression:

$$\begin{aligned} \langle O(\tau) \rangle_\beta = & \sum_{K \geq 0} \sum_{\mathcal{C}_1 \dots \mathcal{C}_{K-1}, \mathcal{C}} \int_{t_0}^t dt_K \int_{t_0}^{t_K} dt_{K-1} \dots \int_{t_0}^{t_2} dt_1 O(\tau) \\ & \rho_\beta(t_K | \mathcal{C}_{K-1}, t_{K-1}) \dots \rho_\beta(t_1 | \mathcal{C}_0, t_0) e^{-(t-t_K)r_\beta(\mathcal{C})} \\ & Y(\mathcal{C}_0)^{t_1-t_0} \frac{W_\beta(\mathcal{C}_0 \rightarrow \mathcal{C}_1)}{r_\beta(\mathcal{C}_0)} \dots Y(\mathcal{C}_{K-1})^{t_K-t_{K-1}} \frac{W_\beta(\mathcal{C}_{K-1} \rightarrow \mathcal{C})}{r_\beta(\mathcal{C}_{K-1})} Y(\mathcal{C})^{t-t_K} \end{aligned} \quad (18)$$

One simply runs the same algorithm as before, which generates the bias on trajectories, except that whenever an agent arrives at a time t_k such that $t_{k-1} \leq \tau < t_k$, the corresponding value $O(\mathcal{C}_{k-1})$ is attached to the agent. Then, each time an agent is cloned, the corresponding value of O is copied accordingly. At the end of the simulation, $\langle O(\tau) \rangle_\beta$ is obtained from the average of the values of $O(\tau)$ attached to the surviving clones. Of course, thanks to the cloning process between τ and t , this average differs from the one we could have done at the intermediate time τ in the simulation.

The same kind of scheme also applies to compute the weighted average of any observable O depending on the whole history of the system and the crucial step is to copy the value of the observable when cloning events occur. Note in particular that the determination of $\langle O(\tau) \rangle_\beta$ can be quite noisy since only a few instances of $O(\tau)$ have survived at time t . In the long time limit one may gain similar information by studying $\frac{1}{t} \langle \int_0^t d\tau O(\mathcal{C}(\tau)) \rangle_\beta$ [20] which is a less noisy dynamical observable.

DISCUSSION

These numerical methods have been applied successfully in many different situations but it is important to keep in mind their limitations. First, as pointed out in [17] convergence problems are met when the evolution operator is gapless. This can for instance happen in systems where the configuration space is unbounded, as in the Zero Range Process but will however not be a problem as long as the configuration space remains finite.

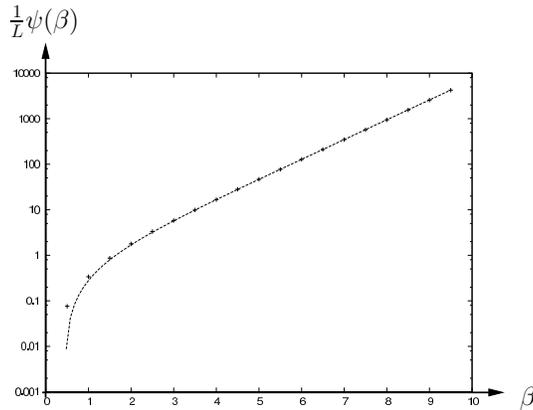


FIGURE 1. Large deviation function $\psi(\beta)/L$ (in log scale) for the total current of particles in the SSEP at density $\rho = 1/2$ ($L = 400$ sites). The points (+) are the result of the continuous-time numerical algorithm. The line is the analytic result (19) valid for very large deviations $|\beta| \gg 1$.

The other important limitation is the finiteness of the population of agents, which can be problematic for very large deviations yielding large cloning rates. For all applications considered in [15], we thus ensured that the cloning factor would never grow larger than few percent of the total population size and agreement with theory - when at hand - was very good. As an example, one can compare (figure 1) for the simple symmetric exclusion process (SSEP) the numerical result in the regime of very large deviations with the analytical result [18]

$$\frac{1}{L}\psi(\beta) = 2 \cosh \beta \frac{\sin \pi \rho}{\pi} - 2\rho(1 - \rho) - 2 \frac{\sin^2(\pi \rho)}{\pi^2} + \mathcal{O}(e^{-|\beta|}) \quad (19)$$

valid for $|\beta| \gg 1$, for the total current of particles in the system. Agreement is very good although the values of β correspond to very large deviations. For large cloning rates, it may also be necessary to modify step (4) of the continuous time algorithm so that agents are not pruned uniformly but according to the weight they carry since their last change of configuration $\propto e^{(t-t_k)[r_\beta(\mathcal{C}_k) - r(\mathcal{C}_k)]}$ but in our simulations we never ran into this problem. In this worst case scenario, the efficiency of the continuous time implementation would fall back to that of the discrete time where at every step the whole population is re-sampled.

ACKNOWLEDGMENTS

We thank R.J. Harris, A. Rákos and H. Touchette for many useful discussions and P. Hurtado for suggesting this review. JT acknowledges funding from EPSRC grants EP/030173. VL was supported by the Swiss FNS, under MaNEP and division II.

REFERENCES

1. H. Touchette, arXiv:0804.0327 (to appear in Phys. Rep.)
2. D.J. Evans, E.G.D. Cohen and G.P. Morriss, *Phys. Rev. Lett.* **71** 2401 (1993)
3. G. Gallavotti and E.G.D. Cohen, *Phys. Rev. Lett.* **74** 2694 (1995)
4. J. Kurchan, *J. Phys. A* **31** 3719 (1998)
5. J.L. Lebowitz and H. Spohn, *J. Stat. Phys.* **95** 333 (1999)
6. B. Derrida, *J. Stat. Mech.* P07023 (2007)
7. T. Bodineau and B. Derrida, *Phys. Rev. E* **72** 066110 (2005)
8. J. Mehl, T. Speck, and U. Seifert, *Phys. Rev. E* **78** 011123 (2008)
9. M. Baiesi, C. Maes, K. Netočný, Accepted in J. Stat. Phys., arXiv:0807.0145
10. S. Tănase-Nicola and J. Kurchan, *Phys. Rev. Lett.* **91** 188302 (2003)
11. S. Tănase-Nicola and J. Kurchan, *J. Stat. Phys.* **116** 1201 (2004)
12. J. Tailleur, S. Tănase-Nicola and J. Kurchan, *J. Stat. Phys.* **122** 557 (2006)
13. J. Tailleur and J. Kurchan, *Nat. Phys.* **3** 203 (2007)
14. C. Giardinà, J. Kurchan and L. Peliti, *Phys. Rev. Lett.* **96** 120603 (2006)
15. V. Lecomte and J. Tailleur, *J. Stat. Mech.* P03004 (2007).
16. A. Imparato, L. Peliti, *J. Stat. Mech.* L02001 (2007)
17. A. Rákos, R.J. Harris *J Stat. Mech.* P05005 (2008)
18. V. Lecomte, *PhD thesis*, Paris VII University (2007)
19. P. Grassberger, *Comp. Phys. Comm.* **147** 64-70 (2002)
20. J.P. Garrahan, R.L. Jack, V. Lecomte, E. Pitard, K. van Duijvendijk and F. van Wijland, arXiv:0810.5298