



*Citation for published version:*

Brewer, T, Clark, SR, Bradford, R & Jack, RL 2018, 'Efficient characterisation of large deviations using population dynamics', *Journal of Statistical Mechanics-Theory and Experiment*, pp. 1-29.  
<https://doi.org/10.1088/1742-5468/aab3ef>

*DOI:*

[10.1088/1742-5468/aab3ef](https://doi.org/10.1088/1742-5468/aab3ef)

*Publication date:*

2018

*Document Version*

Early version, also known as pre-print

[Link to publication](#)

This is an author-created, un-copyedited version of an article published in *Journal of Statistical Mechanics: Theory and Experiment*. IOP Publishing Ltd is not responsible for any errors or omissions in this version of the manuscript or any version derived from it. The Version of Record is available online at <https://doi.org/10.1088/1742-5468/aab3ef>

**University of Bath**

**Alternative formats**

If you require this document in an alternative format, please contact:  
[openaccess@bath.ac.uk](mailto:openaccess@bath.ac.uk)

**General rights**

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

**Take down policy**

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

# Efficient characterisation of large deviations using population dynamics

Tobias Brewer<sup>1</sup>, Stephen R. Clark<sup>1,2</sup> Russell Bradford<sup>3</sup>, and Robert L. Jack<sup>1,4,5</sup>,

<sup>1</sup> Department of Physics, University of Bath, Bath BA2 7AY, United Kingdom

<sup>2</sup> Max Planck Institute for the Structure and Dynamics of Matter, University of Hamburg CFEL, Hamburg, Germany

<sup>3</sup> Department of Computer Science, University of Bath, Bath BA2 7AY, United Kingdom

<sup>4</sup> Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Wilberforce Road, Cambridge CB3 0WA, United Kingdom

<sup>5</sup> Department of Chemistry, University of Cambridge, Lensfield Road, Cambridge CB2 1EW, United Kingdom

**Abstract.** We consider population dynamics as implemented by the cloning algorithm for analysis of large deviations of time-averaged quantities. Using the simple symmetric exclusion process as a prototypical example, we investigate the convergence of the results with respect to the algorithmic parameters, focussing on the dynamical phase transition between homogeneous and inhomogeneous states, where convergence is relatively difficult to achieve. We discuss how the performance of the algorithm can be optimised, and how it can be efficiently exploited on parallel computing platforms.

## 1. Introduction

Large deviation theory [1] is concerned with the probabilities of rare events in random processes. Despite their scarcity, such events can have dramatic consequences, especially if one considers the behaviour of systems on long time scales, as in geology [2] or climate science [3]. In theoretical physics, a number of recent studies (for example [4, 5, 6, 7, 8]) have concentrated on rare events in which a system does not behave ergodically – that is, time-averaged observable quantities have non-typical values, even when measured over long time periods. While, it is challenging to capture rare events in experiments (see however [9]), there are a range of computational approaches for the characterisation of rare events [10, 11, 12, 13, 14, 15, 16, 17] – these are vital since analytical calculations are usually possible only in very simple systems [4]. Computational analyses of large deviations using these methods have lead to new insights in glassy materials [18, 19], protein-folding [8, 20, 21] and integrable systems [22]. In particular, analysis of the associated rare-event mechanisms can reveal properties of metastable or unusual dynamical states that can aid understanding of the typical behaviour.

When considering the large deviations of time-averaged quantities, the two dominant computational methods are transition path sampling [10] and a cloning (or population dynamics) algorithm [12]. This article investigates properties of the cloning method, applied to the symmetric simple exclusion process (SSEP). This is a prototypical interacting-particle system, for which analytical results are available [4, 23, 24], allowing a direct comparison between theory and simulation. While the model is very simple, it exhibits a range of surprising rare-event phenomena, including events where many particles assemble into a large (macroscopic) cluster, as well as hyperuniform states [25], where density fluctuations are strongly suppressed. These regimes of behaviour are separated by dynamical phase transitions [26, 27, 28]: at these points, numerical calculations become challenging and require analysis of large system sizes and long time scales.

The cloning algorithm was introduced more than ten years ago [12, 29], based on earlier ideas in statistical physics [30, 31] and quantum mechanics [32]. It has been applied to a range of systems [22, 33, 34, 35]. The method is powerful, but it requires simulations of many copies (“clones”) of the system, and its results are accurate only in the limit where the number of clones tends to

infinity. For finite numbers of clones, there are both systematic and random errors, which have been analysed recently in Refs. [36, 37, 38, 39]. The scaling of these errors has been determined using an analytical description of the algorithm and has been verified using a numerical approach to measure how quickly an estimator converges towards its true value [36]. Simple guiding models (or control forces) have been used to improve the efficiency of convergence [40, 35, 39]. In previous analyses of the convergence of the algorithm, it was often assumed that the clone population is larger than the total number of states visited by the model [36, 37]: this is not the case in typical applications so we analyse here the case where the clone population is much smaller than the total number of states.

We have implemented the cloning method on high-performance computing platform, which allows us to investigate large numbers of clones. We present a detailed analysis of the dynamical phase transition that occurs in the SSEP. The finite-size scaling properties of this system differ from conventional phase transitions [23, 24], and we discuss the physical reasons for this. In addition, we discuss the systematic errors inherent in the cloning method; we provide practical heuristics as to how the significance of these errors can be assessed; we provide some simple optimisations of the method in order to reduce these errors. Finally, we discuss our computational implementation, and how this can be optimised to improve performance when the number of clones is large.

The form of this paper is as follows: Section 2 provides background on large deviation theory and the cloning method; Section 3 describes the model and its dynamical phase transition, including some numerical results that allow characterisation of finite-size effects. Section 4 explains in more detail how we apply the method to the SSEP, and how this method is optimised to reduce errors. After that, Section 5 analyses the convergence of the algorithm with respect to the number of clones, and the (long) time scale associated with the rare events of interest. Finally, Section 6 discusses the computational implementation, and we summarise our conclusions in Section 7.

## 2. Large deviation theory and the cloning algorithm

### 2.1. Large deviation theory for time averaged-quantities

Consider a physical system described by a Markov process on a discrete set of states (for example, the SSEP). The state of the system at time  $t$  is  $\mathcal{C}_t$ , the transition rate from state  $\mathcal{C}$  to  $\mathcal{C}'$  is  $W(\mathcal{C} \rightarrow \mathcal{C}')$ , and we define the *escape rate* as  $r(\mathcal{C}) = \sum_{\mathcal{C}'} W(\mathcal{C} \rightarrow \mathcal{C}')$ . Let  $\Theta$  denote a trajectory of the system, during the time interval  $[0, t_{\text{obs}}]$ , for some *observation time*  $t_{\text{obs}}$ . In trajectory  $\Theta$ , suppose that the configuration changes happen at times  $t_1, t_2, \dots, t_K$ , and let the state of the system just after the  $k$ th change be  $\mathcal{C}_k$  (also let the initial configuration be  $\mathcal{C}_0$  and define  $t_{K+1} = t_{\text{obs}}$  and  $t_0 = 0$ ). Then, denoting the probability of trajectory  $\Theta$  by  $P(\Theta)$ , one has (see for example [19]):

$$P(\Theta) = \left[ \prod_{k=0}^{K-1} W(\mathcal{C}_k \rightarrow \mathcal{C}_{k+1}) \right] \cdot \exp \left[ - \sum_{k=0}^K r(\mathcal{C}_k)(t_{k+1} - t_k) \right]. \quad (1)$$

Now let  $A_t$  be a (random) observable quantity that depends on the behaviour of the system during the time-interval  $[0, t]$ . For example, in the SSEP,  $A_t$  will be the total number of particle hopping events in  $[0, t]$ , as discussed in [4, 25]. Alternatively,  $A_t$  might be a time integral of the form  $\int_0^t b(\mathcal{C}_{t'}) dt'$ , where  $b$  is some function that depends on the configuration. With either of these choices, one expects that for large  $t$ , the probability distribution of  $A_t$  scales as

$$\text{Prob}(A_t \approx at) \sim \exp(-\pi(a)t). \quad (2)$$

This is an example of a large-deviation principle [1, 19]. The function  $\pi$  is known as the rate function, and satisfies  $\pi(a) \geq 0$ . Typically, there is a single value  $\bar{a}$  for which  $\pi(\bar{a}) = 0$ . In that case, one sees from (2) that as  $t \rightarrow \infty$ , the distribution of  $a_t = A_t/t$  concentrates on the single value  $\bar{a}$ , with the probability of any other value being suppressed exponentially in  $t$ . Of course, the validity of (2) depends on the system of interest and the observable  $A$  – here we consider irreducible Markov processes with finite (discrete) state spaces, for which (2) holds for a large set of observables  $A_t$ : see eg [19].

The general aims of rare-event sampling methods in this context are (i) to estimate the function  $\pi(a)$ , which gives the probability of rare events (with  $a_t \neq \bar{a}$ ); and (ii) to characterise the

rare events themselves: what trajectories lead to these rare values? To achieve these aims, it is convenient to introduce a *biasing field* – denoted by  $s$  – which allows access to the relevant rare events. We use  $\Theta$  to denote a trajectory of the system, during the time interval  $[0, t]$ , and let  $P(\Theta)$  be the probability of trajectory  $\Theta$  as defined, for example in [19]. The distribution  $P(\Theta)$  depends on the initial condition of the model. The results of the following large-deviation analysis are independent of the initial condition, but we assume for concreteness that the initial condition is taken from the steady-state probability distribution of the model, so that  $P(\Theta)$  corresponds to the steady state.

Now define a new probability distribution

$$\tilde{P}_t(\Theta, s) = \frac{P(\Theta) \exp[-sA_t(\Theta)]}{Z(s, t)}, \quad (3)$$

where  $A_t(\Theta)$  is the value of  $A_t$  associated with trajectory  $\Theta$ , and  $Z(s, t) = \langle e^{-sA_t} \rangle_0$  is a dynamical partition function (normalisation constant). Here and throughout, the notation  $\langle \cdot \rangle_0$  indicates an average with respect to  $P(\Theta)$ . The average of an observable with respect to  $\tilde{P}_t$  is denoted by

$$\langle \mathcal{O} \rangle_s = \int \mathcal{O}(\Theta) \tilde{P}_t(\Theta, s) d\Theta. \quad (4)$$

It is useful to define

$$\psi(s) = \lim_{t \rightarrow \infty} \frac{1}{t} \ln Z(s, t). \quad (5)$$

This limit certainly exists if the LDP (2) holds.

The distribution  $\tilde{P}$  is parameterised by the field  $s$ . For  $s = 0$  we recover the original distribution  $P$ ; for  $s > 0$  trajectories with large values of  $A_t$  are suppressed. The advantage of introducing the field  $s$  is that averages of the form  $\langle \mathcal{O} \rangle_s$  can often be evaluated by some numerical or analytical method. In the absence of dynamical phase transitions, one may then obtain the rate function in (2), as  $\pi(a) = \max_s [-sa - \psi(s)]$ . Moreover, if the value of  $s$  that achieves this minimum is  $s_a^*$  then trajectories obtained from the distribution  $\tilde{P}(\Theta, s_a^*)$  are representative trajectories associated with the rare event  $a_t = a$  discussed above [41]. Thus, computational analysis of  $\tilde{P}$  can achieve the two aims (i) and (ii) above, to estimate  $\pi$  and to characterise the trajectories that realise these rare events.

The function  $\psi$  is a scaled cumulant generating function [19]: one has

$$\lim_{t \rightarrow \infty} \langle A_t/t \rangle_s = -\psi'(s) \quad (6)$$

where the prime denotes a derivative. There is also an associated susceptibility (scaled variance)

$$\lim_{t \rightarrow \infty} \frac{1}{t} \left[ \langle A_t^2 \rangle_s - \langle A_t \rangle_s^2 \right] = \psi''(s). \quad (7)$$

## 2.2. Modified Dynamics

Large deviations are hard to analyse computationally because the associated events are rare – this means that averages such as  $\langle e^{-sA_t} \rangle_0$  are dominated by trajectories that are not at all typical of the system at equilibrium. To analyse such events, it is often convenient to modify the dynamics of the model, so that the relevant trajectories become less rare [12, 40, 35]. Consider a general system with modified (or “controlled”) dynamics, and let its path probability distribution be  $\hat{P}(\Theta)$ , which is analogous to the distribution  $P(\Theta)$  for the original system. For the modifications that we consider, it is possible to relate these two distributions, as

$$P(\Theta) = \hat{P}(\Theta) e^{-\hat{U}_t(\Theta)} \quad (8)$$

where  $\hat{U}_t(\Theta)$  is a weight function that depends on the trajectory  $\Theta$ . (A specific example will be considered in Sec. 3.2, below.) Hence one has also

$$\tilde{P}_t(\Theta, s) = \frac{1}{Z(s, t)} \hat{P}(\Theta) e^{-[\hat{U}_t(\Theta) + sA_t(\Theta)]}. \quad (9)$$

The significance of this result is that the distribution  $\tilde{P}_t$  can be analysed in many different ways: either directly as in (3) or by simultaneously modifying the dynamics of the system,  $P(\Theta) \rightarrow \hat{P}(\Theta)$ , and at the same time modifying the weighting factor as  $sA_t(\Theta) \rightarrow [sA_t(\Theta) + \hat{U}_t(\Theta)]$ . This freedom to modify the dynamics is very useful when designing computational algorithms. Finally, we define

$$\Upsilon_t(\Theta) = \exp[-\hat{U}_t(\Theta) - sA_t(\Theta)], \quad (10)$$

which is the weight that should be associated with trajectory  $\Theta$ , in order to obtain the distribution (3) by importance sampling from the distribution  $\hat{P}$ . In particular, we have

$$Z(s, t) = \int \Upsilon_t(\Theta) \hat{P}(\Theta) d\Theta. \quad (11)$$

### 2.3. Cloning algorithm

The results of this paper use the cloning (or population dynamics) algorithm that was proposed by Giardina, Kurchan and Peliti [12] as a method for studying large deviations. As noted above, this method draws on earlier work by Grassberger [31] as well as Diffusion Quantum Monte Carlo methods [32, 38]. We outline this algorithm here, further details are provided in Section 4, below. The method is based on simulations of a population of  $n_c$  copies (or clones) of the system, evolving over a total observation time  $t_{\text{obs}}$ . The field  $s$  is a fixed parameter: to obtain accurate estimates of  $\psi(s)$  one requires a limit of large  $n_c$  and  $t_{\text{obs}}$ . The dependence of the results of the algorithm on  $n_c$  and  $t_{\text{obs}}$  will be discussed in Section 5 below. In our implementation, the population size is held strictly constant, although modified algorithms with variable populations are also possible [29].

Within the algorithm, the total time  $t_{\text{obs}}$  is split into intervals of length  $\Delta t$ , so the number of such intervals is  $M = t_{\text{obs}}/\Delta t$ . Within each step of the algorithm, each clone evolves independently for a time  $\Delta t$ . Then, some clones are deleted and others copied, in order to bias the system towards the rare events of interest (this is a form of importance sampling). In the following we refer to these two sub-steps (or stages) as the dynamical stage and the cloning stage of the algorithm. The full algorithmic step – dynamics followed by cloning – is repeated  $M$  times. The parameter  $\Delta t$  can be chosen according to the problem of interest: as  $n_c \rightarrow \infty$  (with fixed  $t_{\text{obs}}$ ) then the results are independent of  $\Delta t$ . However,  $\Delta t$  has significant effects on the accuracy of the results obtained: this is discussed in Section 4.3 below.

We index the time intervals by  $\beta = 1, 2, \dots, M$  and define  $t_\beta = \beta\Delta t$ . Then the cloning method rests on the fact that for any trajectory  $\Theta$ , one may write

$$A_{t_{\text{obs}}}(\Theta) = \sum_{\beta=1}^M A^\beta(\Theta)$$

where  $A^\beta(\Theta)$  is the contribution to  $A_t(\Theta)$  from the time interval  $[t_{\beta-1}, t_\beta]$ . Similarly

$$\hat{U}_{t_{\text{obs}}}(\Theta) = \sum_{\beta=1}^M \hat{U}^\beta(\Theta). \quad (12)$$

We also index the clones of the system by  $i = 1, 2, \dots, n_c$ . Then, define a weighting factor for clone  $i$  associated with time-interval  $\beta$  as

$$\Upsilon^\beta(\Theta_i) = \exp[-\hat{U}^\beta(\Theta_i) - sA^\beta(\Theta_i)]. \quad (13)$$

where  $\Theta_i$  is the trajectory followed by clone  $i$ . This weighting factor plays two roles within the algorithm. First, in the importance sampling step that takes place at time  $t_\beta$ , the average number of times that clone  $i$  is copied is proportional to  $\Upsilon^\beta(\Theta_i)$ . Second, based on (5,11), one may estimate  $\psi(s)$  as

$$\hat{\psi}(s) = \frac{1}{t_{\text{obs}}} \sum_{\beta=1}^M \ln \left( \frac{1}{n_c} \sum_{i=1}^{n_c} \Upsilon^\beta(\Theta_i) \right). \quad (14)$$

For a given computation, this estimator is subject to both systematic and random errors. However, both these errors vanish as  $n_c, t_{\text{obs}} \rightarrow \infty$ , and the estimator becomes exact.

#### 2.4. Estimating averages with respect to $\tilde{P}$ .

To estimate averages of the form  $\langle \mathcal{O} \rangle_s$ , one starts by considering the population of clones at the final time  $t_{\text{obs}}$ . For each clone  $i$  in that population, one follows its trajectory backwards in time: this trajectory is denoted by  $\hat{\Theta}_i$ . Note that many members of the final clone population may be descended from a single clone at some earlier time. Hence, the trajectories  $\hat{\Theta}_i$  are not all independent samples from  $\tilde{P}$ . However, one may estimate the general expectation value (4) as

$$\langle \mathcal{O} \rangle_s \approx \frac{1}{n_c} \sum_{i=1}^{n_c} \mathcal{O}(\hat{\Theta}_i) \quad (15)$$

where  $\mathcal{O}(\Theta)$  is the value of observable  $\mathcal{O}$  in trajectory  $\Theta$ . The approximate equality becomes exact [36] as  $n_c \rightarrow \infty$ .

The observable  $\mathcal{O}$  may depend in general on all aspects of the trajectory  $\Theta$ . It is also useful to consider a specific class of time-dependent observables: let  $F_t$  be a function that depends on the state of the system at time  $t$ , such as the number of particles on a particular lattice site. The average of such an observable is  $\langle F_t \rangle_s$ , which may be evaluated for any time  $t$  between 0 and  $t_{\text{obs}}$ . For  $s = 0$ , the probability distribution  $P$  is time-translation invariant (TTI), which means that  $\langle F_t \rangle_0$  does not depend on  $t$ . However, for  $s \neq 0$ , the average  $\langle F_t \rangle_s$  depends on  $t$ : there are initial and final transient regimes for small  $t$  and for  $t \approx t_{\text{obs}}$ , with an intermediate time-translation invariant regime. That is,

$$\langle F_t \rangle_s = \begin{cases} F_i(t), & t \lesssim \tau \\ F_f(t_{\text{obs}} - t), & (t_{\text{obs}} - t) \lesssim \tau \\ F_\infty, & \text{otherwise} \end{cases} \quad (16)$$

where  $\tau$  is a characteristic time scale for the transient regimes, and  $F_i, F_f$  are functions describing the transients, which decay to the asymptotic value  $F_\infty$  as their arguments get large [42, 40].

It will be useful in the following to consider the probability distribution of  $F_t$ , and not just its mean value. This distribution is defined as

$$p_{t_{\text{obs}}}(F, t) = \langle \delta(F - F_t) \rangle_s,$$

which is the probability (density) to observe the value  $F$  for the observable  $F_t$ , given trajectories of length  $t_{\text{obs}}$  with distribution  $\tilde{P}$ . To characterise the TTI regime, we define

$$p_{\text{ave}}(F) = \lim_{t_{\text{obs}} \rightarrow \infty} p_{t_{\text{obs}}}(F, \alpha t_{\text{obs}}) \quad (17)$$

with  $0 < \alpha < 1$ . The result is independent of  $\alpha$  because we have both  $\alpha t_{\text{obs}} \rightarrow \infty$  and  $(t_{\text{obs}} - \alpha t_{\text{obs}}) \rightarrow \infty$ . Hence,  $\alpha t$  is a time in the TTI regime,  $\alpha t, (t_{\text{obs}} - \alpha t) \gg \tau$  and it follows that  $F_\infty = \int F p_{\text{ave}}(F) dF$ . We also define

$$p_{\text{end}}(F) = \lim_{t_{\text{obs}} \rightarrow \infty} p_{t_{\text{obs}}}(F, t_{\text{obs}}) \quad (18)$$

which is the distribution of  $F$  at time  $t = t_{\text{obs}}$ .

Note that we have focussed here on instantaneous observables:  $F_t$  depends only on the configuration of the system at time  $t$ . However, the definitions of  $p_{\text{ave}}, p_{\text{end}}$  can be straightforwardly generalised to observables that depend on the trajectory of the system, within a finite time window  $[t, t + \Delta t]$ .

Within the cloning algorithm, the relevance of  $p_{\text{ave}}$  and  $p_{\text{end}}$  is that the clone population just after the importance sampling step is distributed as  $p_{\text{end}}$ . On the other hand, the distribution  $p_{\text{ave}}$  characterises the ‘‘ancestral population’’: this is the distribution that one obtains by constructing the trajectories  $\hat{\Theta}$  from the current population by following their histories backwards in time. For a detailed discussion see Ref [40].

### 3. Dynamical phase transition in the symmetric simple exclusion process

In the following, we apply the cloning algorithm to a model system: the symmetric simple exclusion process (SSEP).




Figure 1: Illustration of the SSEP on a one-dimensional lattice of 8 sites with periodic boundary conditions and  $N = 4$  particles. Each hop is attempted with rate 1; all possibilities for the attempted hop are indicated with arrows.

### 3.1. Model and choice of dynamical observable

We consider a one-dimensional lattice of  $L$  sites with periodic boundaries. The lattice is occupied by  $N$  particles with at most one particle per site. Each particle attempts to hop with rate 1 to each of its neighbouring sites as in Figure 1. The attempted hop is successful if the neighbouring site is unoccupied. Let the occupancy of site  $i$  be  $n_i$ . The model obeys detailed balance which means that in its equilibrium (steady) state, the occupancy of each site is independent:  $n_i = 1$  with probability  $\rho = N/L$  and  $n_i = 0$  with probability  $1 - \rho$ .

For any trajectory  $\Theta$ , define the activity  $K_t(\Theta)$  as the total number of (successful) hops in the time-interval  $[0, t]$ . Within the steady state of the model, one has  $\langle K_t/t \rangle_0 = 2L\rho(1 - \rho)$ , since the rate of attempted hops is  $2N = 2\rho L$  and the expected fraction of successful hops is equal to the probability  $(1 - \rho)$  that the destination site is unoccupied. We consider the distribution  $\tilde{P}$  defined as in (3), with  $A_t = K_t$ : from (1) one has

$$\tilde{P}_{t_{\text{obs}}}(\Theta, s) = \frac{1}{Z(s, t_{\text{obs}})} \left[ \prod_{k=0}^{K-1} W(\mathcal{C}_k \rightarrow \mathcal{C}_{k+1}) e^{-s} \right] \cdot \exp \left[ - \sum_{k=0}^K r(\mathcal{C}_k) (t_{k+1} - t_k) \right].$$

### 3.2. Modified dynamics

To improve computational efficiency when sampling from  $\tilde{P}$ , it is useful to adopt a simple modification to the dynamics, as described in Section 2.2. In this modification, all transition rates are rescaled by a factor  $e^{-s}$ .

From (1), the resulting probability distribution is

$$\hat{P}(\Theta) = \left[ \prod_{k=0}^{K-1} W(\mathcal{C}_k \rightarrow \mathcal{C}_{k+1}) e^{-s} \right] \cdot \exp \left[ - \sum_{k=0}^K r(\mathcal{C}_k) e^{-s} (t_{k+1} - t_k) \right].$$

Hence, from (10),

$$\Upsilon_{t_{\text{obs}}}(\Theta) = \exp \left[ - \sum_{k=0}^K r(\mathcal{C}_k) (1 - e^{-s}) (t_{k+1} - t_k) \right]. \quad (19)$$

In the following, we use these modified dynamics and the weight factors  $\Upsilon$  within our implementation of the cloning algorithm. This modification to the dynamics is useful because the biasing field  $s$  has the effect of suppressing all transitions, so as to reduce  $K$ . Since the modified dynamics incorporates this (trivial) effect, the resulting trajectories are closer to the biased trajectories of interest than one would get by simulating the SSEP directly.

### 3.3. Dynamical phase transition

Our motivation for studying large deviations of  $K_t$  in the SSEP is twofold. First, the model is simple enough for a precise numerical characterisation, and is a useful test of the numerical method. Second, there is a dynamical phase transition that takes place in the model [27, 26, 24], which reveals interesting physical effects, and also leads to numerical challenges that we use in later sections to test the cloning method.




Figure 2: Sample trajectories of SSEP with  $L = 50$  and  $t_{\text{obs}} = 10^4$ . (a)  $s = 0$ , the equilibrium state; (b)  $s = 0.008$ , showing evidence of transient clusters; (c)  $s = 0.012$ , in which a single large cluster has formed; (d)  $s = 0.020$ , with most of the particles in a well-defined single cluster. The corresponding values of  $\lambda$  are 0, 20, 30, 50.

The phase transition appears if one fixes the density  $\rho = N/L$  and takes the lattice size  $L \rightarrow \infty$ . [We take this limit after the large- $t$  limit associated with the large deviation principle (2).] For finite  $L$ , the rate function  $\pi(a)$  and the cumulant generating function  $\psi(s)$  in (5) are both analytic functions, and there is no phase transition. However, on taking  $L \rightarrow \infty$ , a singular response to the field  $s$  can be observed, just as conventional phase transitions can be observed on taking the thermodynamic limit. Specifically, one considers  $\psi_*(s) = \lim_{L \rightarrow \infty} [\psi(s)/L]$ , which is analogous to a thermodynamic free energy density, whose derivatives show singular behaviour at phase transitions.

The existence of this phase transition can be easily inferred from the argument of [19]. For  $N$  particles on  $L$  sites, one may construct a configuration  $\mathcal{C}^*$  in which all particles are on adjacent sites, in which case the escape rate is  $r(\mathcal{C}^*) = 2$ . For large deviations of the activity  $K_t$ , one has quite generally [19] that  $\psi(s) \geq -\min_{\mathcal{C}} r(\mathcal{C})$ . Hence  $\psi_*(s) \geq \lim_{L \rightarrow \infty} r(\mathcal{C}^*)/L = 0$ . Moreover,  $\psi_*(0) = 0$  and  $\psi_*$  is non-increasing in  $s$ , so  $\psi_*(s) = 0$  for all  $s > 0$ . Also  $\psi'(0) = 2L\rho(1 - \rho)$ , from (6). It follows that the derivative  $(d\psi_*/ds)$  is singular at  $s = 0$ : this corresponds to a first-order dynamical phase transition.

The physical signature of this transition is shown in Figure 2 where there is a transition from a homogeneous state at  $s = 0$  (particles are distributed evenly throughout the system) to an inhomogeneous (“phase-separated”) state for  $s > 0$ , in which case the particles are segregated into a dense and a dilute region.

To quantify the particle clustering that takes place for  $s > 0$ , it is useful to consider the Fourier transform of the density field:

$$\delta\rho_n = \frac{1}{\sqrt{L}} \sum_{j=1}^N e^{-2\pi i n X_j / L} \quad (20)$$

where  $X_j$  is the index of the site occupied by particle  $j$ , and  $n = 0, 1, \dots, L - 1$ . We focus on the wavenumber that corresponds to the longest wavelength fluctuations, that is  $n = 1$ .

Figure 3 shows that the magnitude of this Fourier component grows as the system becomes inhomogeneous.

### 3.4. Scaling at the dynamical phase transition

To investigate this phase transition in more detail, it is useful to “zoom in” on the crossover from the homogeneous to the inhomogeneous case. To achieve this, let  $s = \lambda/L^2$  and consider the limit



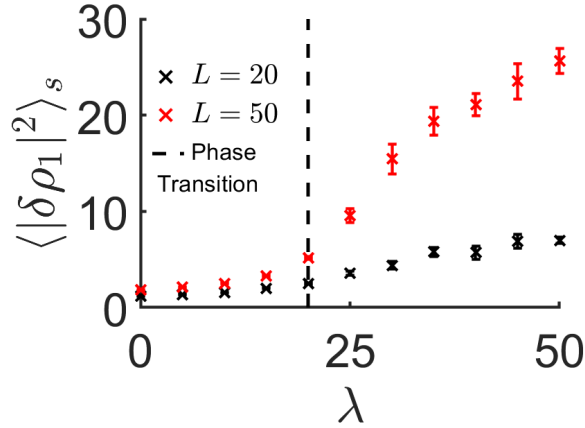


Figure 3: The mean square value of the first Fourier component of the density,  $\langle |\delta\rho_1|^2 \rangle_s$ , measured at  $t = t_{\text{obs}}/2$  with  $t_{\text{obs}} = 10^4$ ,  $n_c = 10^5$ , for various  $L, \lambda$ . The phase transition occurs at  $\lambda_c = 2\pi^2$ : for  $\lambda > \lambda_c$  one expects the system to become inhomogeneous, so that  $\langle |\delta\rho_1|^2 \rangle_s \propto L$ , consistent with the data.

of large  $L$  at fixed  $\lambda$ . This is analogous to finite-size scaling in equilibrium systems. We focus in this work on the crossover function

$$\mathcal{K}_L(\lambda) = k(\lambda/L^2), \quad \text{where} \quad k(s) = L^{-1} \lim_{t \rightarrow \infty} \langle K_t/t \rangle_s = -\psi'(s)/L. \quad (21)$$

As  $L \rightarrow \infty$ , the function  $\mathcal{K}_L$  converges to a limiting form  $\mathcal{K}_*$ , which can be computed using macroscopic fluctuation theory [24]. This function has a singularity at  $\lambda = \lambda_c = 2\pi^2$ . For  $\lambda < \lambda_c$ , one has a constant value  $\mathcal{K}_* = 2\rho(1-\rho)$ . For  $\lambda > \lambda_c$ , the function  $\mathcal{K}_*$  decreases with  $\lambda$ , converging to zero as  $\lambda \rightarrow \infty$ . This behaviour is illustrated in Figure 4, which also shows the rescaled “free energy”

$$\phi_L(\lambda) = L\psi(\lambda/L^2)$$

To show the singularity that appears at the phase transition, we also show two measures of susceptibility

$$\chi_L(s) = L^{-1}\psi''(s), \quad \mathcal{X}_L(\lambda) = -\mathcal{K}'_L(\lambda) = L^{-2}\chi_L(\lambda/L^2). \quad (22)$$

The bare susceptibility  $\chi_L$  corresponds to the scaled variance of  $K_t$  (recall (7)), and  $\lim_{L \rightarrow \infty} \chi_L(s)$  should have a finite value in a system that is away from any phase transition. Figure 4 shows a divergence in  $\chi$ , consistent with the existence of a phase transition. On the other hand, the function  $\mathcal{X}_L(\lambda)$  is predicted by MFT to have a finite limit as  $L \rightarrow \infty$ . Our numerics are consistent with this prediction but they show that measuring this limiting function requires large system sizes.

Note that the scenario illustrated is different from classical finite-size scaling in Figure 4 and from other first-order dynamical phase transitions [35]. Comparing with the classical case, note that we take  $t_{\text{obs}} \rightarrow \infty$  and then later  $L \rightarrow \infty$ . As discussed in Ref. [35], this is equivalent to thermodynamic finite-size scaling in a cylindrical geometry, with the length of the cylinder being much longer than its perimeter. In that case, one possibility is that the susceptibility  $\chi_L$  grows exponentially in the system size, due to the a distribution of domains along the cylinder, with “domain walls” perpendicular to the long axis of the cylinder. However, the results presented are very different from that case: one reason is that the coexisting phases at the transition have different densities, but the number of particles in the SSEP is equal at every time. As a result, the “domain walls” between dense and dilute regions in Figure 2 are constrained to lie parallel to the time axis. In general,  $\chi_L$  can be interpreted as a “correlation volume” in space-time. From (22), and noting that  $\mathcal{X}(\lambda)$  is finite (or zero) for all  $\lambda$  [24], one sees that  $\chi_L$  diverges as  $L^2$  in the vicinity of the transition. This factor arises from the characteristic time scale proportional to  $L^2$  that is associated with density fluctuations on length scale  $L$ . Note also that while  $k(s)$  exhibits a jump at  $s = 0$ , corresponding to a first-order phase transition, one may also consider the




Figure 4: Results illustrating the dynamical phase transition in the SSEP. The number of clones is  $n_c = 10^5$  for  $L = 20, 40, 50$  and  $n_c = 10^6$  for  $L = 80$ . The cloning interval is  $\Delta t = 10$  and all results are averaged across 10 independent computations. The vertical dashed line shows the position of the phase transition ( $\lambda_c = 2\pi^2$ ). The theoretical predictions are obtained by using finite differences to take derivatives of  $\phi(\lambda)$ , which is calculated as described in [24]. For  $\lambda < \lambda_c$  one has  $\mathcal{X} \rightarrow 0$  as  $L \rightarrow \infty$ : the leading behaviour for large  $L$  is  $\mathcal{X} = O(1/L)$ , and was computed in [23]. Our numerical results are consistent with that prediction (data not shown).

behaviour of the system as a function of  $\lambda$ , in which case  $\mathcal{K}(\lambda)$  is continuous at the transition but has a discontinuous first derivative: when viewed on this scale, the transition has some features of a continuous phase transition [43, 44].

#### 4. Implementation of algorithm and cloning stage

We outlined the cloning algorithm in Section 2.3. Here we provide some extra detail on its application to the SSEP. The modified SSEP dynamics are implemented using the Bortz-Kalos-Lebowitz (continuous time Monte Carlo) algorithm [45]. All possible particle hops have the same rate  $e^{-s}$ . These dynamics take place over the time interval  $[t_{\beta-1}, t_\beta]$ . During this time period, the factor  $\Upsilon^\beta(\Theta_i)$  is calculated: based on (19,13) one may write

$$\Upsilon^\beta(\Theta_i) = \exp\left(-\int_{t_{\beta-1}}^{t_\beta} (1 - e^{-s})r(\mathcal{C}_t^i)dt\right)$$

where  $\mathcal{C}_t^i$  is the configuration of clone  $i$  at time  $t$  and  $r(\mathcal{C})$  is the escape rate for configuration  $\mathcal{C}$  under the original (unmodified) dynamics. The factors  $\Upsilon$  appear in the estimate (14) for  $\psi(s)$ .




Figure 5: Example of the number line associated with the cloning factors  $\Upsilon^\beta(\Theta_i)$  used in the clone selection methods, as described in the main text. We take  $n_c = 1000$ . The red arrows indicate some of the points  $\alpha_j$  that determine which clone is selected for inclusion in the new population. The  $\alpha$  may be chosen independently ([iid] method, top line) or equally spaced ([eq] method, lower line).

#### 4.1. Copying and deletion of clones: [eq] and [iid] methods.

The next step is to clone and delete systems, according to their values of  $\Upsilon$ , so as to produce a new population. There is some freedom as to how this is implemented within the algorithm: we require that for a large population, the average number of descendants of clone  $i$  approaches  $\Upsilon^\beta(\Theta_i)/\Upsilon_T^\beta$  with  $\Upsilon_T = \sum_{i=1}^{n_c} \Upsilon^\beta(\Theta_i)$ . We consider two methods for selecting the clones that will form the new population, both of which are consistent with this requirement.

The first clone selection method is denoted by [eq]: the reason for this will be explained below. In this method, for each clone  $j$  of the new population (with  $1 \leq j \leq n_c$ ), we define  $\alpha_j = (j + d - 1)\Upsilon_T/n_c$ , where  $d$  is a random number in  $(0, 1)$ , which is equal for each clone. Thus, the  $\alpha_j$  are equally spaced on  $[0, \Upsilon_T]$ , with  $\alpha_1 = (d\Upsilon_T/n_c)$  and  $\alpha_{n_c} = \Upsilon_T - [(1 - d)\Upsilon_T/n_c]$ . Then the state of clone  $j$  of the new population is pulled from clone  $k$  of the old population, where  $k$  satisfies

$$\sum_{i=1}^k \Upsilon^\beta(\Theta_i) \leq \alpha_j < \sum_{i=1}^{k+1} \Upsilon^\beta(\Theta_i). \quad (23)$$

This clone selection method is equivalent to constructing a line segment of total length  $\Upsilon_T$  that is composed of contributions from each clone of the old population, with clone  $k$  contributing a length  $\Upsilon^\beta(\Theta_k)$ . Then clone  $j$  of the new population is selected by selecting the interval of the original line that contains the point  $\alpha_j$  as in Figure 5. On average, the number of copies of clone  $k$  in the new population is then  $\Upsilon^\beta(\Theta_k)/\Upsilon_T$ , as required. The label [eq] is used because the  $\alpha_j$  are equally spaced on the interval  $[0, \Upsilon_T]$ .

The second clone selection method that we use is denoted by [iid]. In this case, the  $\alpha_j$  are identically and independently distributed (uniformly) on  $[0, \Upsilon_T]$ , hence the label [iid]. However, this choice is less efficient, as we discuss below (Section 4.3). As noted in Sec. 2.3, the two sub-steps, of independent dynamical evolution followed by cloning, are each repeated  $M$  times, so that the total simulation time for each clone is  $t_{\text{obs}}$ . The final step is a cloning step.

#### 4.2. Test case

As a stringent test of algorithmic performance, we focus on the observable  $\mathcal{K}_L(\lambda)$  defined in (21). We consider values of  $\lambda$  in the vicinity of the phase transition.

One sees from Figure 4 that this function has a feature at  $\lambda = 2\pi^2$  that depends strongly on system size. In the following, we will test the ability of the algorithm to capture the finite-size scaling of this feature. Our numerical estimator for  $\mathcal{K}_L$  is




Figure 6: Dependence of results on the cloning interval  $\tau$  and the clone selection mechanism ([eq] or [iid]) at the representative state point  $\lambda = 28$  (see main text). (a) Analysis of the two clone selection mechanisms, showing that if the simulation is converged, the estimator  $\bar{k}$  is (almost) independent of  $\Delta t$ , and of the clone selection method (as required). (b) The statistical errors are heavily dependent on the clone selection method and vary with  $\Delta t$ . (c) The computational time is almost independent of the clone selection method. It decreases with  $\Delta t$  because the amount of communication between nodes decreases, but this has a cost in terms of accuracy (see panel (b)).

$$\hat{k}_L(\lambda) = \frac{1}{Lt_{\text{obs}}} \sum_{i=1}^{n_c} \sum_{\beta=1}^M K^\beta(\hat{\Theta}_i). \quad (24)$$

see also (14). Note, this requires that the trajectories  $\hat{\Theta}$  can be obtained by extrapolation backwards in time: this is easily achieved by including with its clone its accumulated value of  $\sum_{\beta} K^\beta(\hat{\Theta}_i)$ , which is copied along with the system's configuration when the clone state is copied, during the cloning stage of the algorithm (see [46]). We use the method to measure  $\mathcal{K}$  since this is a more revealing measure of algorithmic errors than  $\psi(s)$  or  $\phi(\lambda)$ , as may be seen already in Figure 4. An alternative estimator of  $\mathcal{K}$  can be obtained by using finite differences to estimate the derivative  $\phi'(\lambda)$ , but we prefer the direct measurement (24), which avoids uncertainties arising from the finite differencing.

We particularly focus on systematic errors: we average  $\hat{k}$  over  $R$  independent computations (indexed by  $r = 1, 2, \dots, R$ ) and we denote the average of the estimator by

$$\bar{k}_L(\lambda) = \frac{1}{R} \sum_{r=1}^R \hat{k}_L^r(\lambda),$$

and its variance by

$$[\Delta k_L(\lambda)^2] = \frac{1}{R} \sum_{r=1}^R \left[ \hat{k}_L^r(\lambda)^2 - \bar{k}_L(\lambda) \right]^2.$$

Consequently, its standard deviation is  $\sigma(\hat{k}) = [\Delta k_L(\lambda)^2]^{1/2}$ . For large  $R$ , the systematic error of the method can be obtained as the difference between  $\bar{k}_L(\lambda)$  and  $\mathcal{K}_L(\lambda)$ , and the size of the random error is determined by the variance  $[\Delta k_L(\lambda)^2]$ .

#### 4.3. Effect of cloning method and choice of $\Delta t$

Given a large enough number of clones, the results of the algorithm are independent of the parameter  $\Delta t$ , and they are also independent of whether the  $\alpha_j$  are chosen to be equally spaced [eq] or uniformly at random [iid]. In this section, we show the effects that these choices have on the results obtained with finite  $n_c$ . As a representative state point (unless otherwise stated) we focus on  $\lambda = 28$  and  $t_{\text{obs}} = 10^4$  with  $\Delta t = 10$ , but the results are qualitatively similar for other parameters too. (Note that in this section we show data for both [eq] and [iid] clone selection methods, to illustrate that the performance of the [eq] method is better. All other sections use the [eq] method, unless stated otherwise.)

Results are shown in Figure 6. Panel (a) shows that on using sufficiently many clones and averaging over many computational realisations, the results are independent of  $\Delta t$ , as expected. Panel (b) shows how the variance of the estimate of  $\mathcal{K}_L$  depends on  $\Delta t$ : the smaller is this variance, the less computational realisations are required to get accurate results, and hence the algorithm is more efficient. For the [eq] method, the variance is always smaller than for the [iid] method. Moreover, for [eq] method, the variance is monotonically increasing as a function of  $\Delta t$ . This implies that using a small value of  $\Delta t$  is desirable, from the point of view of accuracy. In practice, computational requirements mean that  $\Delta t$  should not be too small (Figure 6(c)), since the cloning stage of the algorithm incurs an overhead. For the [eq] method, this leads to a simple trade-off between accuracy (better for small  $\Delta t$ ) and computational efficiency (better for large  $\Delta t$ ). For the [iid] method, this trade-off is more complicated since accuracy may also be reduced by choosing smaller  $\Delta t$ .

To see why the [eq] method is more efficient, it is useful to consider the operation of the algorithm with  $s = 0$ . In this case, the most efficient algorithm is clearly to simulate each copy independently, so as to obtain  $n_c$  independent samples. Choosing equally-spaced  $\alpha_j$  ensures that this does indeed happen (all the  $\Upsilon^\beta(\Theta_k) = 1$  so one  $\alpha_j$  lies in each interval). However, choosing the  $\alpha_j$  independently means that some copies of the system are copied several times and some are deleted, because of the randomness inherent in the choice of  $\alpha_j$ . The deletion of clones reduces the number of independent samples in the system and tends to increase the errors. The effect is the same for non-zero  $s$  and the problem also occurs in the continuous-time cloning algorithm as described in [29]. Hence, our conclusion is that choosing equally-spaced  $\alpha_j$  seems to be the most accurate implementation of the cloning step of the algorithm, as described in Sec. 4.

More generally, it is desirable – while always maintaining algorithmic accuracy – to minimise the number of clones that are deleted, since each deletion results in an (irreversible) loss of information, reducing the number of independent samples from  $\bar{P}$ .

## 5. Dependence of the results on the parameters of the cloning algorithm

We have emphasised that the cloning algorithm gives accurate results only in the limit  $n_c \rightarrow \infty$ . Characterising large deviations of the activity also requires that  $t_{\text{obs}} \rightarrow \infty$ . Our ability to probe these limits depends on the available computational resource: it is therefore essential to characterise and understand the dependence of the algorithm's results on  $n_c$  and  $t_{\text{obs}}$ , in order to assess whether the algorithm gives reliable results. However, for finite  $n_c$ , a suitable choice of  $\Delta t$  is also essential to obtain accurate results. This section describes the dependence of the results on  $n_c$  and  $t_{\text{obs}}$ .

We will find that the convergence of the algorithm depends significantly on the value of  $\lambda$ , particularly whether the system is in the homogeneous regime  $\lambda < \lambda_c$  or the phase-separated regime  $\lambda > \lambda_c$ . The convergence also depends on the system size  $L$ , with large values of  $n_c, t_{\text{obs}}$  being required when  $L$  is larger.

### 5.1. Convergence with respect to the time $t_{\text{obs}}$ : effects of long time scales

To analyse the dependence of the method on the choice of  $t_{\text{obs}}$ , we consider the dependence of  $\bar{k}(\lambda)$  on the observation time  $t_{\text{obs}}$ . In this section, the number of clones used is sufficiently large that the results depend very weakly on  $n_c$ : see Section 5.2. Results are shown in Figure 7. For  $\lambda < \lambda_c$ , Figure 7(a) shows that results depend weakly on  $t_{\text{obs}}$  as long as  $t_{\text{obs}} \gtrsim 10^3$ . On the other hand, for  $\lambda > \lambda_c$  there is a systematic dependence on  $t_{\text{obs}}$  even for  $t_{\text{obs}} > 10^3$ . This message is confirmed by Figure 7(b) which shows that systematic errors from finite  $t_{\text{obs}}$  are larger for larger systems, but  $t_{\text{obs}} = 10^4$  seems to be large enough to achieve convergence even for  $L = 80$ .

To understand this effect, recall (16-18,24) and also the discussion of Section 2.4. The trajectories sampled by the algorithm have transient regimes at initial and final times, whose typical time scale is denoted by  $\tau$ . In terms of (24), this means that for values of  $\beta$  that are outside the transient regimes, the variables  $K^\beta$  are all distributed according to the distribution  $p_{\text{ave}}$ : in this case the average of  $K^\beta$  is independent of  $\beta$  and (assuming that  $n_c$  is large enough) this average is equal to  $\Delta t L \mathcal{K}_L(\lambda)$ . If these terms dominate the sum in (24) then  $\bar{k}(\lambda) = \mathcal{K}_L(\lambda)$ , as required. However, when  $\beta$  is close to 1 or  $M$ , the average value of  $K^\beta$  depends on  $\beta$  and is not equal to




Figure 7: Estimates  $\bar{k}(\lambda)$  of the activity  $\mathcal{K}_L(\lambda)$  as a function of  $t_{\text{obs}}$  for various system sizes  $L$ . (a)  $L = 50$  for various  $\lambda, t_{\text{obs}}$ . (b)  $\lambda = 28$  for various  $t_{\text{obs}}, L$ . The “theory” line is  $\mathcal{K}_* = \lim_{L \rightarrow \infty} \mathcal{K}_L(\lambda)$ . We took  $R = 10$ , with  $n_c = 10^5$  for  $L = 20, 50$  and  $n_c = 10^6$  for  $L = 80$  (these  $n_c$  are large enough that results depend very weakly on  $n_c$ .)




Figure 8: Time-dependent averages of the form  $\langle F_t \rangle_s$  with  $F$  being either  $|\delta\rho_1|^2$  (the squared modulus of the first fourier component of the density) or  $K^\beta$  (which is the number of particle hops between times  $t = t_\beta$  and  $t_\beta + \Delta t$ ). Parameters are  $\lambda = 28$ ,  $n_c = 10^6$ ,  $t_{\text{obs}} = 10^4$ ,  $\Delta t = 10$ .

$\Delta t \mathcal{K}_L(\lambda)$ . The inclusion of these terms in (24) means that  $\bar{k}(\lambda) \neq \mathcal{K}_L(\lambda)$  in general: there are corrections of the order of  $\tau/t$  coming from the transient regimes.

Figure 8 shows the transient regime that occurs for  $(t_{\text{obs}} - t) \lesssim \tau$ , for two different observables. Two features are important here: the time scale associated with the transient regime, and the difference between  $\langle F_{t_{\text{obs}}} \rangle_s$  and  $F_\infty$  (which is the value when  $t_{\text{obs}} - t \gg \tau$ , recall (16)). For both observables, the order of magnitude of the transient time scale is  $\tau \simeq 100$  for  $L = 50$  and  $\tau \simeq 300$  for  $L = 80$ . Obtaining an accurate estimate of  $\mathcal{K}$  from (24) requires that  $t_{\text{obs}} \gg \tau$ , in order that the sum is dominated by terms that are outside the transient regime. This is consistent with the observation that  $t_{\text{obs}} \approx 10^4$  is required in order to obtain results with high-accuracy for  $L = 80$ . In terms of the values of  $\langle F_t \rangle_s$ , one sees that  $|\rho_1|^2$  changes by more than a factor of 2, while the fractional changes in  $K^\beta$  are much smaller. This indicates that it is the long-wavelength density modes that respond most strongly to the field  $\lambda$ , consistent with Figure 3. This is as expected, because these slowly-relaxing long-wavelength modes the origin of the dynamical phase transition [25].

To understand these dynamical fluctuations in more detail, we also consider the autocorrelation function for observable  $F$ , defined as

$$C_F(t, t') = \frac{\langle F_{t'} F_t \rangle_s - \langle F_{t'} \rangle_s \langle F_t \rangle_s}{c_F(t')} \quad (25)$$




Figure 9: Autocorrelation function  $C_F$ , as a function of the general time  $t$  and the final time  $t_{\text{obs}}$ , for systems of  $L = 50$  (left),  $80$  (right). The normalisations  $c(K_\beta, t')$  are  $946$  ( $L = 50$ ) and  $1450$  ( $L = 80$ ). Data are shown for the observable  $F$  being the escape rate  $r$ , the first Fourier component of density  $|\delta\rho_1|^2$ , the activity per cloning interval  $K_\beta$  and the cloning factor  $\Upsilon$ . Parameters are  $t_{\text{obs}} = 10^4$ ,  $n_c = 10^6$ ,  $\lambda = 28$ ,  $\Delta t = 10$ .

where  $c_F(t') = \langle F_{t'} F_{t'} \rangle_s - \langle F_{t'} \rangle_s \langle F_{t'} \rangle_s$  is a normalisation factor that ensures that  $C_F(t, t) = 1$ . Since the transient time  $\tau$  is comparable with the inverse of the spectral gap of the stochastic process [42], one expects  $C_F(t, t')$  to be small if  $t - t' \gg \tau$ . Figure 9 shows results for several different observables, always with  $t' = t_{\text{obs}}$  (this is the case for which good statistics are most easily obtained). For the fourier component  $F = |\delta\rho_1|^2$ , one sees that  $C_F(t, t_{\text{obs}})$  remains close to 1 until  $t_{\text{obs}} - t \approx \tau$ , after which point the density fluctuations at the two times decorrelate and the correlation decays. On the other hand, for other observables such as the activity  $K_\beta$ , the correlation  $C(t, t_{\text{obs}})$  is significantly less than 1 already for  $t_{\text{obs}} - t = 10$ , indicating that the activity fluctuations have a “fast component” whose correlations decay quickly, as well as a slow component that is (presumably) correlated with the slow decay of large density fluctuations.

We also note that the activity correlations play a special role in the theory: one has  $\chi(s) \sim L^{-1} c_F(t_{\text{obs}}) \int_0^{t_{\text{obs}}} C_{K_\beta}(t, t_{\text{obs}}) dt$  [19], so  $\chi$  can be large if either the prefactor  $c_F$  is large, or if the function  $C_{K_\beta}$  decays slowly to zero (so that the time integral becomes large). The data in Figure 9 (and its caption) show that the prefactor  $c_F$  scales roughly as  $L$  while the time  $\tau$  is expected to scale as  $L^2$  (the slowest time scale in the system is diffusive decay associated with wavelengths of order  $L$ ). Hence  $\chi \sim L^2$  which is consistent with  $\mathcal{X} = O(1)$  and  $\chi \sim L^2$  (these results are in the regime  $\lambda > \lambda_c$ , see also (22) and Figure 4).

The conclusions of this analysis are (i) that the longest relaxation time in the system controls the convergence with respect to  $t_{\text{obs}}$  and (ii) that these relaxation times can be revealed by explicitly computing transients and autocorrelation functions as in Figs. 8,9. We also emphasize that making these measurements are not only useful for verifying convergence of the algorithm: they also reveal the important physical effects at work in the biased trajectories of interest: for  $\lambda > \lambda_c$ , the dominant physical effect is that the density becomes inhomogeneous on the macroscopic scale, so that  $\langle |\delta\rho_1|^2 \rangle_s$  diverges with system size. There is a slow time scale associated with this macroscopic inhomogeneity, which scales as  $\tau \sim L^2$  and results in a large susceptibility  $\chi$ . This long time scale necessitates a large  $t_{\text{obs}}$  in the cloning algorithm, since accurate estimates of observables like  $\langle K \rangle_s$  require  $t_{\text{obs}} \gg \tau$ .

## 5.2. Population Size $n_c$ Convergence

This section shows the convergence of the algorithm as  $n_c$  increases. In all cases we take  $t_{\text{obs}} = 10^4$ , which is large enough that results depend weakly on  $t_{\text{obs}}$ . Fig. 10 shows results, for a range of system sizes and numbers of clones. From Fig. 10a, one sees that for the relatively large system size  $L = 80$ , a good estimate of  $\mathcal{K}$  is available already for  $n_c = 10^5$ , but increasing to  $n_c = 10^6$  clones and focussing on  $s > s^*$ , the results still depend on  $n_c$ , indicating that the large- $n_c$  limit is not fully converged for  $n_c = 10^5$ . (For  $s < s^*$ , the results do not agree perfectly with theoretical

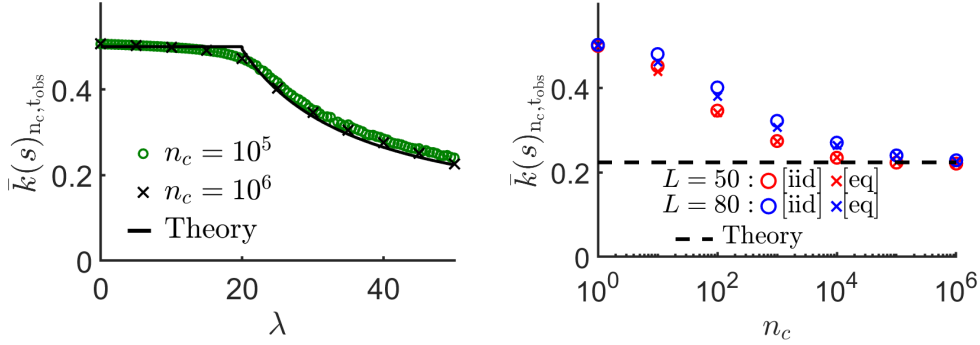


Figure 10: Estimates  $\bar{k}_L(\lambda)$  of the activity, as the number of clones is increased. (a) Data for  $L = 80$  and  $n_c = 10^5, 10^6$ , as  $\lambda$  is varied. The theory line is the result for  $L \rightarrow \infty$ . (b) Data for  $\lambda = 50$  and  $L = 50, 80$  as  $n_c$  is varied. In all cases,  $t_{\text{obs}} = 10^4$  and  $R = 10$ .




Figure 11: Distributions of the activity  $K^\beta(\Theta_i)$  at  $\beta = 970$  when  $t_{\text{obs}} = 10^4$  and  $\Delta t = 10$ . (a)  $L = 50$ ; (b)  $L = 80$ . The data points are measured across  $n_c = 10^6$  systems and gaussian fits produced by measuring the average  $\mu$  and variance  $\sigma^2$  across the population. All results are at  $\lambda = 50$ .

prediction shown in Fig. 10a: this is because the theory applies only in the large- $L$  limit, but these are finite systems. The important feature for convergence is not the agreement with the analytic theory, but whether the results depend significantly on  $n_c$ .) Fig. 10b shows results at the state point  $\lambda = sL^2 = 50$ , which is in the phase-separated regime (recall Fig. 2): one sees that the number of clones required to obtain accurate results is of order  $10^5 - 10^6$ , with the larger system requiring larger  $n_c$ .

To understand why these large numbers of clones are required, it is useful to return to the distributions  $p_{\text{ave}}$  and  $p_{\text{end}}$  defined in (17,18). We construct these distributions for the observable  $K^\beta$ , the number of particle hopping events in the time interval  $[t_\beta, t_{\beta+1}]$ . Results are shown in Figure 11. Recall that the distribution  $p_{\text{end}}$  is representative of the transient regime ( $t \approx t_{\text{obs}}$ ), while the distribution  $p_{\text{ave}}$  is representative of the time-translation-invariant (TTI) regime. Accurate characterisation of the TTI regime is essential for obtaining accurate estimates of  $\mathcal{K}_L$ .

However, we note that the operation of the cloning algorithm means that the distribution  $p_{\text{end}}$  is sampled directly, but  $p_{\text{ave}}$  is obtained by a form of importance sampling from the distribution  $p_{\text{end}}$ . In terms of Figure 11, this means that data is only available for  $p_{\text{ave}}$  over a restricted range, which is the range over which data is available for  $p_{\text{end}}$ . If these two distributions have a low overlap, then the algorithm will not sample  $p_{\text{ave}}$  correctly: this is an important source of the systematic errors in Figure 10.



To estimate the significance of this effect (see also [47]), suppose that both  $p_{\text{ave}}$  and  $p_{\text{end}}$  are described by (approximately) Gaussian distributions with mean values  $\mu_{\text{ave}}, \mu_{\text{end}}$  and the same variance  $\sigma^2$ . (Note  $\mu_{\text{ave}} = \mathcal{K}_L$ .) To be precise, define  $g(x, \mu, \sigma) = \exp(-(x - \mu)^2/2\sigma^2)/\sqrt{2\pi\sigma^2}$  and suppose  $p_{\text{end}}(K) \approx g(K, \mu_{\text{end}}, \sigma)$ , and similarly for  $p_{\text{ave}}$ . Also define a scaled version of the complementary error function as  $G_\sigma(x) = \int_x^\infty g(y, 0, \sigma)dy$ . Given  $n_c$  clones drawn independently from  $p_{\text{end}}$ , we expect to sample a range of  $K$ -values  $(\mu_{\text{end}} - A) < K < (\mu_{\text{end}} + A)$  with  $G_\sigma(A) = (1/n_c)$ . We further assume that the measured  $p_{\text{ave}}$  is close to the true  $p_{\text{ave}}$  over the range over which data is available – this is certainly an approximation but it allows an estimate of the effect of the unsampled range on the results of the algorithm. With this approximation, the average value of  $K$  with respect to  $p_{\text{ave}}$  can be estimated as

$$\bar{k} \approx \frac{\int_{\mu_{\text{end}}-A}^{\mu_{\text{end}}+A} K g(K, \mu_{\text{ave}}, \sigma) dK}{\int_{\mu_{\text{end}}-A}^{\mu_{\text{end}}+A} g(K, \mu_{\text{ave}}, \sigma) dK} \quad (26)$$

Assuming as in Figure 11 that  $\mu_{\text{end}} > \mu_{\text{ave}}$  (and  $\sigma$  is not too large), one may replace the upper limits in (26) by infinity. Writing  $\Delta\mu = (\mu_{\text{end}} - \mu_{\text{ave}})$ , this yields

$$(\bar{k} - \mu_{\text{ave}}) \approx \sigma/\sqrt{2\pi} \cdot \frac{\exp(-(A - \Delta\mu)^2/2\sigma^2)}{G_\sigma(\Delta\mu - A)} \quad (27)$$

This error converges to zero as  $A \rightarrow +\infty$ , as it should do. However, for large  $n_c$  one has the scaling relation  $A \sim \sigma\sqrt{\log n_c}$  so  $A$  grows slowly with  $n_c$ . The relevant dimensionless parameter for convergence is  $X = (A - \Delta\mu)/\sigma$ : this parameter is positive if the peak of the  $p_{\text{ave}}$  distribution is within the range of the sampled data: accurate results require  $X \gtrsim 1$ . This requires

$$n_c \gtrsim \exp [(\Delta\mu/\sigma)^2]. \quad (28)$$

In general, one expects  $\Delta\mu$  to increase as the biasing field increases; on the other hand  $\sigma$  is expected to depend weakly on the bias but (since  $K$  is extensive in space) one expects  $\sigma \sim 1/L$  in large systems. Hence one expects that the number of clones required to achieve convergence increases exponentially in the system size and in the distance from equilibrium. For a similar argument, see [47, Section 5]. Based on the Gaussian fits in Figure 11 the lower bounds on  $n_c$  are  $2.16 \times 10^3$  ( $L = 50$ ) and  $4.93 \times 10^4$  ( $L = 80$ ): for high-accuracy estimates at these state points, we require  $n_c = 10^5$  and  $n_c = 10^6$  respectively, so high accuracy requires parameters significantly in excess of these lower bounds, consistent with this argument.

## 6. Computation

In this section, we discuss the performance of our computational implementation of the cloning algorithm. Given the large number of clones evolving independently, it is natural to use high-performance (parallel) computing methods to speed up the computations. However, the step where clones are copied and deleted involves a significant communication overhead, which can substantially reduce efficiency. Close to the phase transition ( $\lambda = \lambda_c \approx 20$ ) the variance in activity is close to its maximum. This suggests that there should be a large variance in the  $\Upsilon^\beta(\Theta_i)$ , and a large amount of cloning activity. This makes the phase transition a useful place to test the algorithm.

We compare results obtained using MPI with simple serial computations and a shared-memory (OpenMP) implementation. We define the speed-up factor of an implementation  $m$  as

$$\mathcal{S}_m = \text{RT}_s/\text{RT}_m \quad (29)$$

where  $\text{RT}_m$  is the run time of implementation  $m$  and  $\text{RT}_s$  is the run time for a simple serial code. For a computation that uses  $n_t$  threads, we define the efficiency as

$$\mathcal{E}_m = \mathcal{S}_m/n_t \quad (30)$$

Both the speed-up and the efficiency are depend in general on  $n_t$ . We describe an OpenMP implementation that is 99% efficient (but does not scale beyond a single computational node).




Figure 12: OpenMP implementation of the algorithm code run on one node (ellipse) and multiple threads (grey circles).

Method	Speed-Up	Efficiency (%)
One clone per message	48.3	75.5
Many clones per message	55.2	86.2
Reduced communication	57.2	89.3

Table 1: Speed-ups and efficiencies relative to our serial implementation obtained by MPI implementations when using the equal spacing [eq] clone selection method with different amounts of packing when running on 64 processors distributed across 4 nodes. The results are obtained at  $\lambda = 20$  and  $L = 50$  with  $t_{\text{obs}} = 10^4$ ,  $n_c = 10^5$  and  $\Delta t = 10$ .

For performance on multiple nodes, we use a method based on the message-passing interface (MPI) protocol, which achieves an efficiency of around 90% on computations distributed over 64 processors. We measure weak scaling performance up to 128 processors. All codes were written in C++ and will be available shortly after publication.

All results were obtained on (8-core) Intel Xeon E5-2650V2 Ivybridge processors running at 2.6 GHz. As a representative test case, we consider a system of size  $L = 50$  at parameters at which the activity estimator  $\hat{k}$  has converged,  $t_{\text{obs}} = 10^4$  and  $n_c = 10^5$ . We average all results over 10 independent computations. This computation takes 5500 seconds to run on a single core. We present results for efficiency relative to this reference point. The shared memory implementation uses OpenMP with 16 threads on a dual-socket node. For the MPI implementation, we typically use four such nodes connected using Infiniband QDR with a total of 64 threads.

### 6.1. Serial and OpenMP Implementations

The simplest implementation is a serial code, which simulates the dynamics for each clone in turn. The cloning step involves copying and deletion of the different clones. This is achieved by each clone in the new population “pulling” its state (that is, the site occupancies, particle positions, etc) from a member of the old population. Determining which clone to copy is based on (23), and implemented using a binary search tree. Random number generation uses the Mersenne Twister algorithm [48], with one instance of the random number generator for each clone. Except where otherwise stated (see below), this means that the results are fully reproducible, independent of whether the serial or parallel code is used.

In the OpenMP (shared memory) implementation, the clones are distributed equally over 16 threads, on a single computational node. The dynamics are simulated as in the serial case. For the cloning step, the binary search key is constructed on a single thread, but the copying of new clone states is done in parallel. A diagram summarising this process is shown in Figure 12. For our test case ( $L = 50$ ,  $t_{\text{obs}} = 10^4$ ,  $n_c = 10^5$ ,  $\lambda = 20$ ), the speed-up of the OpenMP implementation with respect to the serial code is 15.8, corresponding to an efficiency of 99%. Efficiencies at other state points are similar.

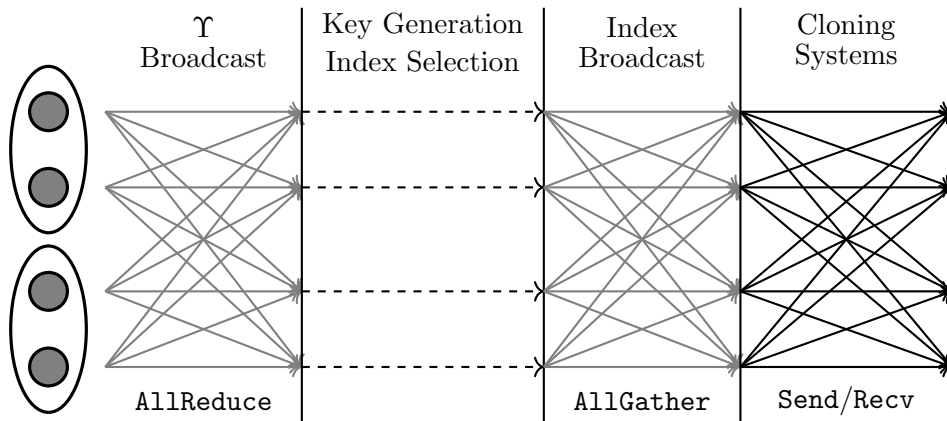


Figure 13: MPI implementation of the algorithm code run on two nodes (ellipses) and multiple threads (grey circles). Arrows are MPI communications, black arrows are pointwise and grey arrows are collective.

## 6.2. MPI Implementation

In order to exploit higher levels of parallelisation, we also use a distributed memory (MPI) implementation, which is illustrated in Figure 13. The direct generalisation of the OpenMP code gives rise to an implementation that we call “simple communication”. In this case, the MPI method produces identical output to the OpenMP method. We also describe a method with reduced communication, in which case the results are statistically equivalent (but not identical), because the systems being simulated are distributed differently among the relevant threads, and hence the instance of the random number generator that is used to simulate a particular system is different.

**6.2.1. Simple communication.** In the dynamical stage of the algorithm, the clones are distributed evenly across the threads. In the cloning step, each thread sends the  $\Upsilon^\beta(\Theta_i)$  for its clones to all other threads (this is an MPI `AllReduce` communication). Each thread constructs the binary key independently (this computation is therefore duplicated for each thread, but this is more efficient than calculating it on one thread and broadcasting it to all others). With the binary key in place, each clone in the new population uses (23) to decide from which clone to pull its state (the value of  $d$  in (23) is produced by thread 0 and broadcast to all threads after the key generation). Next, the indices of the clones that are required by each thread are communicated to all other threads (MPI `AllGather` operation); based on that information, the clone states are sent between the nodes, as required (using many MPI `Send` operations). In the simplest case, the state of each clone is sent using a single MPI message (we call this “one-clone-per-message”). Alternatively, all information to be exchanged between each pair of threads can be packed into a single message (“many-clones-per-message”).

From Table 1, these two methods achieve efficiencies of 75-86% on systems of  $n_t = 64$  threads, with the many-clones-per-message method being more efficient. We emphasise that the amount of information being exchanged between each thread is identical in this case, so the increased efficiency comes from packing the same information into a smaller number of messages. For one-clone-per-message, the total number of messages is  $n_c = 10^5$ ; for many-clones-per-message, each thread receives at most one other message from each other thread, so the total number of messages is at most  $n_t(n_t - 1) = 4032$  (see Sec. 6.2.3, below). This is a reduction by more than an order of magnitude in the number of messages, which significantly improves performance.

**6.2.2. Reduced communication.** The simple communication method is somewhat inefficient because in some situations, a particular thread ( $A$ ) may send several clones to some other thread ( $B$ ), but thread  $B$  also sends several clones back to thread  $A$ . To avoid this redundancy, we use




Figure 14: The number of systems sent between each pair of processors using MPI when evolving  $n_c = 10^5$ ,  $t_{\text{obs}} = 10^4$  and  $L = 50$ . Simple communications (*left*) and reduced communications (*right*) using the independent [iid] clone selection method. These communications are in the 500<sup>th</sup> cloning interval ( $t = 5000$ ) and the cloning intervals are  $dt = 10$  units of time. The algorithm runs at a bias  $\lambda = 20$ . *Bottom Row*: The number of systems sent by each processor using MPI (blue) and those whose states are simply copied (orange).

a cancellation procedure where each thread preferentially pulls states from clones that are on the same thread. (This is achieved by pairwise cancellation of clones that have been designated by (23) to be sent between threads; the resulting new population is identical but its partitioning among the threads is different.) With this “reduced communication” method, clones are sent either from  $A$  to  $B$  or from  $B$  to  $A$  (instead of sending one message in each direction). This typically reduces the number of messages by almost a factor of 2, and the size of the resulting messages by a factor of 9. Table 1 shows that this leads to an improvement in efficiency.

**6.2.3. Communication Patterns** To illustrate the communications pattern, Figure 14 shows the number of messages sent and received by each thread in a typical cloning step at  $t = t_{\text{obs}}/2 = 5000$ . In the simple implementation, each thread typically sends around 30 clones to each other thread and receives a similar number. In the reduced communication implementation, there are less than half as many MPI communications and those that occur contain about 5 systems.

**6.2.4. Weak Scaling** For a fixed number of threads  $n_t$ , the run time for each algorithm scales linearly with the number of clones  $n_c$ . A *weak scaling* analysis is one in which the problem size (in this case  $n_c$ ) is increased, with a proportionate increase in the computational resource (the number of nodes  $n_N = n_t/16$ ).

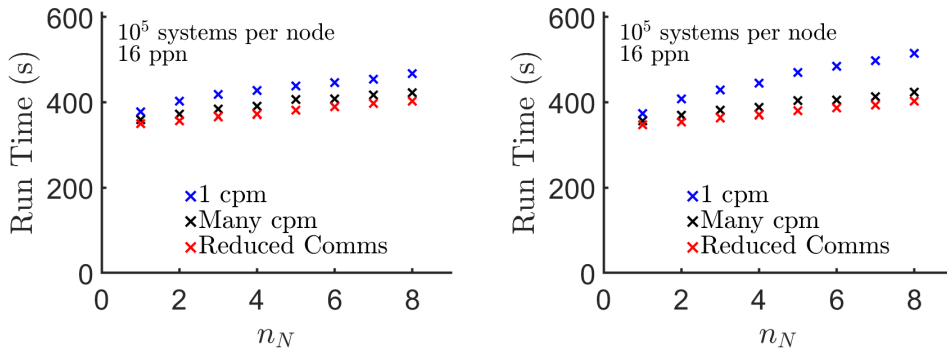


Figure 15: The weak scaling of the run time of the code for both clone selection methods [iid] (left) and [eq] (right). We fix  $n_c/n_N = 10^5$  and the code is run at  $t_{\text{obs}} = 10^4$ ,  $L = 50$ ,  $\lambda = 20$  and  $dt = 10$ . We vary the number of clones per message (cpm) and reduce the communications.

In the best (most efficient) case, this leads to a run-time that remains constant as the problem-size  $n_c$  increases. In general, one expects the run-time to increase with problem size, with a smaller increase corresponding to a more efficient algorithm.

In Figure 15 we measure the weak-scaling of a system at  $L = 50$  with enough clones ( $n_c/n_N = 10^5$ ) that the algorithm is converged when the code is run on one node. The value of the bias ( $\lambda = 20$ ) that we consider is very close to the phase transition and the value of the estimator  $\hat{k}$  depends only very weakly on  $n_c$ . We show results for both the [eq] and [iid] clone selection methods, but the results are similar in both cases.

As the problem size increases, one observes a gradual increase in run time: for the most efficient algorithm, this increases by approximately 5% while the number of clones has increased by a factor of 8. This shows that parallel-computing platforms significantly reduce the wall-time required to perform large computations, although (as expected) the method is most efficient when run on just a few nodes, since the communication overhead is lower in that case. The gradual increase in run-time with  $n_N$  does not indicate any optimum value for  $n_N$  in this case nor any cutoff beyond which the method becomes inefficient: the code is performing well across this range of  $n_N$  (which corresponds to at most 128 parallel threads).

## 7. Conclusion

We have analysed the finite-size scaling of the dynamical phase transition in the SSEP. Our results show that the cloning algorithm can demonstrate convergence of the limit of infinite system-size  $L \rightarrow \infty$  in this problem, although of order  $10^6$  clones are required to achieve this. By analysis of the  $p_{\text{ave}}$  and  $p_{\text{end}}$  distributions of the activity we can understand the number of clones required for convergence. We also analysed the time  $t_{\text{obs}}$  required for convergence of the limit  $t_{\text{obs}} \rightarrow \infty$ , by measuring of the transient decay of several observables, and their associated autocorrelation functions.

We have achieved computational efficiencies of 99% using an OpenMP implementation of the algorithm, although this requires shared memory for all threads and so can be used on a single node. We also used an MPI implementation, so that the algorithm can be scaled across multiple nodes. Optimising MPI communication to use as few messages as possible significantly increases the algorithmic performance. We have achieved this by point-wise reduction of the MPI communications between pairs of processors and packing multiple systems into each MPI message. By measuring the weak-scaling of run time across multiple processors we find that there is a slow drop off in efficiency.

The implementation of the cloning stage of the algorithm can affect the errors in the method. We have made a comparison of two methods for selecting which systems are cloned, one of these methods reduces the number of systems that are deleted. Whilst both methods produce similar

systematic errors and require similar computation time, we have found that the method which reduces the number of system deletions incurs much smaller statistical errors.

As the cloning algorithm becomes more widely used [35, 38, 43], we hope that these methods for error analysis and parallel computing implementation will be useful for the growing effort into research on large deviations of time-averaged quantities.

## Acknowledgments

RLJ thanks Takahiro Nemoto and Vivien Lecomte for useful discussions about the population dynamics algorithm. Computations were performed on the HPC service at the University of Bath.

## 8. References

- [1] Touchette H 2009 *Phys. Rep.* **478** 1–69
- [2] Gretener P 1967 *AAPG Bulletin* **51** 2197–2206
- [3] Easterling D R, Meehl G A, Parmesan C, Changnon S A, Karl T R and Mearns L O 2000 *Science* **289** 2068–2074
- [4] Derrida B 2007 *J. Stat. Mech.* **2007** P07023
- [5] Garrahan J P, Jack R L, Lecomte V, Pitard E, van Duijvendijk K and van Wijland F 2007 *Phys. Rev. Lett.* **98** 195702
- [6] Hedges L O, Jack R L, Garrahan J P and Chandler D 2009 *Science* **323** 1309–1313
- [7] Speck T, Malins A and Royall C P 2012 *Phys. Rev. Lett.* **109**(19) 195703
- [8] Weber J K, Jack R L and Pande V S 2013 *Journal of the American Chemical Society* **135** 5501–5504
- [9] Pinchaipat R, Campo M, Turci F, Hallett J E, Speck T and Royall C P 2017 *Phys. Rev. Lett.* **119**(2) 028004
- [10] Bolhuis P G, Chandler D, Dellago C and Geissler P L 2002 *Ann. Rev. Phys. Chem.* **53** 291–318
- [11] Allen R J, Warren P B and ten Wolde P R 2005 *Phys. Rev. Lett.* **94** 018104
- [12] Gardina C, Kurchan J and Peliti L 2006 *Phys. Rev. Lett.* **96** 120603
- [13] Huber G A and Kim S 1996 *Biophysical Journal* **70** 97–110
- [14] Zhang B W, Jasnow D and Zuckerman D M 2010 *J. Chem. Phys.* **132** 054107
- [15] Bouchet F and Reygnier J 2016 *Annales Henri Poincaré* **17** 3499–3532
- [16] Johnson T H, Elliott T, Clark S R and Jaksch D 2015 *Phys. Rev. Lett.* **114** 090602
- [17] Johnson T H, Clark S R and Jaksch D 2010 *Phys. Rev. E* **82**(3) 036702
- [18] Jack R L, Garrahan J P and Chandler D 2006 *The Journal of chemical physics* **125** 184509
- [19] Garrahan J P, Jack R L, Lecomte V, Pitard E, van Duijvendijk K and van Wijland F 2009 *Journal of Physics A: Mathematical and Theoretical* **42** 075007
- [20] Weber J K, Jack R L, Schwantes C R and Pande V S 2014 *Biophysical journal* **107** 974–982
- [21] Mey A S, Geissler P L and Garrahan J P 2014 *Physical Review E* **89** 032109
- [22] Tailleur J and Kurchan J 2007 *Nature Physics* **3** 203–207
- [23] Appert-Rolland C, Derrida B, Lecomte V and van Wijland F 2008 *Phys. Rev. E* **78** 021122
- [24] Lecomte V, Garrahan J P *et al.* 2012 *Journal of Physics A: Mathematical and Theoretical* **45** 175001
- [25] Thompson I R and Jack R L 2015 *Physical Review E* **92** 052115
- [26] Bodineau T and Derrida B 2005 *Physical Review E* **72** 066110
- [27] Bodineau T and Derrida B 2007 *Comptes Rendus Physique* **8** 540–555
- [28] Bodineau T, Lecomte V and Toninelli C 2012 *Journal of Statistical Physics* **147** 1–17
- [29] Lecomte V and Tailleur J 2007 *Journal of Statistical Mechanics: Theory and Experiment* **2007** P03004
- [30] Grassberger P and Procaccia I 1983 *Phys. Rev. A* **28**(4) 2591–2593
- [31] Grassberger P 2002 *Computer Physics Communications* **147** 64–70
- [32] Anderson J B 1975 *J. Chem. Phys.* **63** 1499–1503
- [33] Pitard E, Lecomte V and van Wijland F 2011 *EPL* **96** 56002
- [34] Hurtado P I, Espigares C P, del Pozo J J and Garrido P L 2014 *J. Stat. Phys.* **154** 214–264
- [35] Nemoto T, Jack R L and Lecomte V 2017 *Phys. Rev. Lett.* **118**(11) 115702
- [36] Nemoto T, Guevara Hidalgo E and Lecomte V 2017 *Phys. Rev. E* **95**(1) 012102
- [37] Guevara Hidalgo E, Nemoto T and Lecomte V 2017 *Phys. Rev. E* **95**(6) 062134
- [38] Ray U, Kin-Lic Chan G and T Limmer D 2017 arXiv:1708.00459
- [39] Ray U, Kin-Lic Chan G and T Limmer D 2017 arXiv:1708.09482
- [40] Nemoto T, Bouchet F, Jack R L and Lecomte V 2016 *Physical Review E* **93** 062123
- [41] Chetrite R and Touchette H 2015 *Ann. Henri Poincaré* **16** 2005–2057
- [42] Jack R L and Sollich P 2010 *Progress of Theoretical Physics Supplement* **184** 304–317
- [43] Carollo F, Garrahan J P, Lesanovsky I and Pérez-Espigares C 2017 *Phys. Rev. E* **96**(5) 052118
- [44] Baek Y, Kafri Y and Lecomte V 2017 *Phys. Rev. Lett.* **118**(3) 030604
- [45] Bortz A B, Kalos M H and Lebowitz J L 1975 *J. Comp. Phys.* **17** 10–18
- [46] Tailleur J and Lecomte V 2009 *AIP Conference Proceedings* vol 1091 p 212
- [47] Hurtado P I and Garrido P L 2009 *J. Stat. Mech.* **2009** P02032
- [48] Matsumoto M and Nishimura T 1998 *ACM Trans. Model. Comput. Simul.* **8** 3–30