

THE WANG-LANDAU ALGORITHM REACHES THE FLAT HISTOGRAM CRITERION IN FINITE TIME

BY PIERRE E. JACOB*, ROBIN J. RYDER^{†‡}

Department of Statistics & Applied Probability – National University of Singapore

Centre de Recherche en Économie et Statistique and CEREMADE – Université Paris Dauphine

stajpe@nus.edu.sg

ryder@ceremade.dauphine.fr

The Wang-Landau algorithm aims at sampling from a probability distribution, while penalizing some regions of the state space and favouring others. It is widely used, but its convergence properties are still unknown. We show that for some variations of the algorithm, the Wang-Landau algorithm reaches the so-called Flat Histogram criterion in finite time, and that this criterion can be never reached for other variations. The arguments are shown on a simple context – compact spaces, density functions bounded from both sides – for the sake of clarity, and could be extended to more general contexts.

1. Introduction and notations. Consider the problem of sampling from a probability distribution π defined on a measure space $(\mathcal{X}, \Sigma, \mu)$. We suppose that we can compute the probability density function of π at any point $x \in \mathcal{X}$, up to a multiplicative constant. Given a proposal kernel $Q(\cdot, \cdot)$ we define a Metropolis–Hastings (MH) (Hastings, 1970; Tierney, 1998) transition kernel targeting π , denoted by $K(\cdot, \cdot)$, as follows:

$$\forall x, y \in \mathcal{X} \quad K(x, y) = Q(x, y)\rho(x, y) + \delta_x(y)(1 - r(x))$$

with $\rho(x, y)$ defined by $\rho(x, y) := 1 \wedge \frac{\pi(y)Q(y, x)}{\pi(x)Q(x, y)}$ and $r(x)$ defined by:

$$r(x) := \int_{\mathcal{X}} \rho(x, y)Q(x, y)dy$$

*Department of Statistics & Applied Probability, National University of Singapore. Partially funded by AXA Research.

[†]Centre de Recherche en Économie et Statistique and CEREMADE. Partially funded by a GIS Sciences de la Décision postdoc scholarship.

[‡]Both authors contributed equally to this work.

AMS 2000 subject classifications: Primary 65C05, 60J22; secondary 60J05

Keywords and phrases: Markov Chain, Monte Carlo, Wang-Landau algorithm, Flat Histogram

Here the delta function $\delta_a(b)$ takes value 1 when $a = b$ and 0 otherwise. Under some conditions on the proposal Q and the target π , the MH kernel defines an algorithm to generate a Markov chain with stationary distribution π (Robert and Casella, 2004).

Let us consider a partition of the state space \mathcal{X} into d disjoint sets $\mathcal{X}_1, \dots, \mathcal{X}_d$:

$$\mathcal{X} = \bigcup_{i=1}^d \mathcal{X}_i$$

If we have a sample X_1, \dots, X_t independent and identically distributed from π , then for any $i \in [1, d]$:

$$\frac{1}{t} \sum_{n=1}^t \mathbb{I}_{\mathcal{X}_i}(X_n) \xrightarrow[t \rightarrow \infty]{\mathbb{P}} \int_{\mathcal{X}_i} \pi(x) dx =: \psi_i$$

where we denote by $\mathbb{I}_{\mathcal{X}_i}(x)$ the indicator function that is equal to 1 when $x \in \mathcal{X}_i$ and 0 otherwise. Similar convergence is obtained when X_1, \dots, X_t is an ergodic chain such as the one generated by the MH algorithm. The purpose of the Wang-Landau algorithm (Wang and Landau, 2001a,b; Liang, 2005; Atchadé and Liu, 2010) is to obtain a sample

- such that for any $i \in [1, d]$ the subsample

$$\{X_n \text{ for } n \in [1, t] \text{ s.t. } X_n \in \mathcal{X}_i\}$$

is distributed according to the restriction of π to \mathcal{X}_i , and

- such that for any $i \in [1, d]$

$$\frac{1}{t} \sum_{n=1}^t \mathbb{I}_{\mathcal{X}_i}(X_n) \xrightarrow[t \rightarrow \infty]{\mathbb{P}} \phi_i$$

where $\phi = (\phi_1, \dots, \phi_d)$ is chosen by the user, and could be any vector in $]0, 1[^d$ such that $\sum_{i=1}^d \phi_i = 1$.

A typical use of this algorithm is to sample from multimodal distributions, by penalizing already-visited regions and favouring the exploration of regions between modes, in an attempt to recover all the modes. In many applications, ϕ is set to $\forall i, \phi_i = 1/d$; however, other choices are possible, as exemplified by the adaptive binning scheme of Bornn et al. (2011).

This algorithm, in the class of Markov Chain Monte Carlo (MCMC) algorithms (Robert and Casella, 2004), therefore allows to learn about π while “forcing” the proportions of visits ϕ_i of the generated chain to any of the

sets \mathcal{X}_i , which are typically also chosen by the user. The vector ϕ_1, \dots, ϕ_d might be referred to as the “desired frequencies”, and the sets \mathcal{X}_i are called the “bins”. In a typical situation, the mass of π over bin \mathcal{X}_i , which we denote by ψ_i , is unknown, and hence one cannot easily guess how much to “penalize” or to “favour” a bin \mathcal{X}_i in order to obtain the desired frequency ϕ_i . The Wang-Landau algorithm introduces a vector $\theta_t = (\theta_t(1), \dots, \theta_t(d))$, referred to as “penalties” at time t , which is updated at every iteration t , and which acts like an approximation of the ratios $\psi_1/\phi_1, \dots, \psi_d/\phi_d$, up to a multiplicative constant.

For a distribution π and a vector of penalties $\theta = (\theta(1), \dots, \theta(d))$, we define the penalized distribution π_θ :

$$\pi_\theta(x) \propto \pi(x) \times \sum_{i=1}^d \frac{\mathbb{1}_{\mathcal{X}_i}(x)}{\theta(i)}$$

To be more concise we define a function $J : \mathcal{X} \mapsto \{1, \dots, d\}$ that takes a state $x \in \mathcal{X}$ and returns the index i of the bin \mathcal{X}_i such that $x \in \mathcal{X}_i$. We can now write: $\pi_\theta(x) \propto \pi(x)/\theta(J(x))$. We will denote by K_θ the MH kernel targeting π_θ .

The Wang-Landau algorithm, described in the next section, alternates between generating a sample by targeting π_θ using K_θ , and updating θ using the generated sample. In this sense it is an adaptive MCMC algorithm (past samples are used to update the kernel at a given iteration), using an auxiliary chain (θ_t) , and therefore the behaviour of the sample is not obvious.

The Wang-Landau algorithm is widely used in the Physics community (Silva, Caparica and Plascak, 2006; Malakis, Kalozoumis and Tyraskis, 2006; Cunha Netto et al., 2006). In particular, many practitioners use flavours of the algorithm with a “Flat-Histogram” criterion. However, its convergence properties are still partially unknown. We show that this criterion is reached in finite time for some variations of the algorithm. This result is all that was missing to apply results on adaptive algorithms with diminishing adaptation (Fort et al., 2011).

In Section 2, we define variations of the Wang-Landau algorithm. We then introduce ratios of penalties and argue for their convenience in studying the properties of the algorithm. We prove in Sections 3 and 4 that under certain conditions, the Flat Histogram criterion is met in finite time, for the cases $d = 2$ and $d > 2$ respectively. The result is illustrated in Section 5, and in Section 6, we hint at how our assumptions might be relaxed.

2. Wang-Landau algorithms: different flavours. There are several versions of the Wang-Landau algorithm. We describe the general version

introduced by [Atchadé and Liu \(2010\)](#), both in its deterministic form and with a stochastic schedule.

2.1. *A first version with deterministic schedule.* Let $(\gamma_t)_{t \in \mathbb{N}}$ (referred to as a schedule or a temperature) be a sequence of positive real numbers such that:

$$\begin{cases} \sum_{t \geq 0} \gamma_t & = \infty \\ \sum_{t \geq 0} \gamma_t^2 & < \infty \end{cases}$$

A typical choice is $\gamma_t := t^{-\alpha}$ with $\alpha \in]0.5, 1[$. The Wang-Landau algorithm is described in pseudo-code in [Algorithm 1](#). In this form, the schedule γ_t decreases at each iteration, and is therefore called “deterministic”.

Algorithm 1 Wang-Landau with deterministic schedule

- 1: Init $\forall i \in \{1, \dots, d\}$ set $\theta_0(i) \leftarrow 1/d$.
 - 2: Init $X_0 \in \mathcal{X}$.
 - 3: **for** $t = 1$ to T **do**
 - 4: Sample X_t from $K_{\theta_{t-1}}(X_{t-1}, \cdot)$, MH kernel targeting $\pi_{\theta_{t-1}}$.
 - 5: Update the penalties: $\log \theta_t(i) \leftarrow \log \theta_{t-1}(i) + f(\mathbb{I}_{\mathcal{X}_i}(X_t), \phi_i, \gamma_t)$.
 - 6: **end for**
-

Step 5 of [Algorithm 1](#) updates the penalties from θ_{t-1} to θ_t , by increasing it if the corresponding bin has been visited by the chain at the current iteration, and by decreasing it otherwise. This rationale seems natural, however we did not find any article arguing for a particular choice of update, among the infinite number of updates that would also follow the same rationale. In other words, it is not obvious how to choose the function f , except that it should be such that it is positive when $X_t \in \mathcal{X}_i$ and such that it is closer to 0 when γ_t decreases, to ensure that the penalties converge. Some practitioners use the following update:

$$(1) \quad \log \theta_t(i) \leftarrow \log \theta_{t-1}(i) + \gamma_t (\mathbb{I}_{\mathcal{X}_i}(X_t) - \phi_i)$$

while others use:

$$(2) \quad \log \theta_t(i) \leftarrow \log \theta_{t-1}(i) + \log [1 + \gamma_t (\mathbb{I}_{\mathcal{X}_i}(X_t) - \phi_i)]$$

Since γ_t converges to 0 when t increases, and since update (1) is the first-order Taylor expansion of update (2), one legitimately expects both updates to result in similar performance in practice. We shall see in [Section 3](#) that this is not necessarily the case.

Some convergence results have been proven about [Algorithm 1](#): the deterministic schedule ensures that θ_t changes less and less along the iterations of the algorithm, and consequently the kernels K_{θ_t} change less and

less as well. The study of the algorithm hence falls into the realm of adaptive MCMC where the *diminishing adaptation* condition holds (Andrieu and Thoms, 2008; Atchadé et al., 2009; Fort et al., 2011), although it is original in the sense that the target distribution (π_{θ_t}) is adaptive but not necessarily the proposal distribution Q . See also the literature on stochastic approximation (Andrieu, Moulines and Priouret, 2006).

In this article we are especially interested in a more sophisticated version of the Wang-Landau algorithm that uses a stochastic schedule, and for which, as we shall see in the following, the two updates result in different performance.

2.2. A sophisticated version with stochastic schedule. A remarkable improvement has been made over Algorithm 1: the use of a “Flat Histogram” (FH) criterion to decrease the schedule only at certain random times. Let us introduce $\nu_t(i)$, the number of generated points at iteration t that are in \mathcal{X}_i :

$$\nu_t(i) := \sum_{n=1}^t \mathbb{1}_{\mathcal{X}_i}(X_n)$$

For some predefined precision threshold c , we say that (FH) is met at iteration t if:

$$\max_{i \in \{1, \dots, d\}} \left| \frac{\nu_t(i)}{t} - \phi_i \right| < c$$

Intuitively, this criterion is met if the observed proportion of visits to each bin is not far from ϕ , the desired proportion. The name “Flat Histogram” comes from the observation that if the desired proportions are all equal to $1/d$, this criterion is verified when the histogram of visits is approximately flat. The threshold c could possibly decrease along the iterations, to get an always finer precision.

The Wang-Landau with Flat Histogram (Algorithm 2) is similar to the previous algorithm, with a single difference: the schedule γ does not decrease at each step anymore, but only when (FH) is met. To know whether it is met or not, a counter ν_t of visits to each bin is updated at each iteration, and when (FH) is met, the schedule decreases and the counter is reset to 0.

Note the difference between Algorithms 1 and 2: γ is indexed by κ instead of t , and κ is a random variable. As with Algorithm 1, the update of penalties (step 13 of Algorithm 2) can be either update (1) or update (2), or possibly something else. Interestingly in this case, it is not obvious anymore that both updates will give similar results. Indeed, for γ_κ to go to 0, we need (FH) to be reached in finite time, so that κ regularly increases.

Algorithm 2 Wang-Landau with Flat Histogram

```

1: Init  $\forall i \in \{1, \dots, d\}$  set  $\theta_0(i) \leftarrow 1/d$ .
2: Init  $X_0 \in \mathcal{X}$ .
3: Init  $\kappa = 0$ , the number of (FH) criteria already reached.
4: Init the counter  $\forall i \in \{1, \dots, d\}$   $\nu_1(i) \leftarrow 0$ 
5: for  $t = 1$  to  $T$  do
6:   Sample  $X_t$  from  $K_{\theta_{t-1}}(X_{t-1}, \cdot)$  targeting  $\pi_{\theta_{t-1}}$ .
7:   Update  $\nu_t$ :  $\forall i \in \{1, \dots, d\}$   $\nu_t(i) \leftarrow \nu_{t-1}(i) + \mathbb{1}_{\mathcal{X}_i}(X_t)$ 
8:   Check whether (FH) is met.
9:   if (FH) is met then
10:      $\kappa \leftarrow \kappa + 1$ 
11:      $\forall i \in \{1, \dots, d\}$   $\nu_t(i) \leftarrow 0$ 
12:   end if
13:   Update the bias:  $\log \theta_t(i) \leftarrow \log \theta_{t-1}(i) + f(\mathbb{1}_{\mathcal{X}_i}(X_t), \phi_i, \gamma_\kappa)$ .
14: end for

```

This flavour of the Wang-Landau algorithm is widely used in the Physics literature (Cunha Netto et al., 2006; Silva, Caparica and Plascak, 2006; Malakis, Kalozoumis and Tyraskis, 2006; Ngo and Diep, 2008).

Our contribution is to show in a simple context that update (1) is such that (FH) is met in finite time, while (2) is not so. Hence only using update (1) can one expect the convergence properties of Algorithm 1 to still hold for Algorithm 2, since if (FH) is met in finite time a sort of *diminishing adaptation* condition would still hold.

To underline the difficulty of knowing whether (FH) is met in finite time or not, let us recall that between two (FH) occurrences, the schedule is constant (equal to some $\gamma_\kappa > 0$), hence the penalties (θ_t) change at a constant scale and *diminishing adaptation* does not directly hold. Other adaptive algorithms share this lack of *diminishing adaptation*, as e.g. the Accelerated Stochastic Approximation algorithm (Kesten, 1958), in which the adaptation of some process diminishes only if its increments change sign. In our case, (FH) will be reached if the chain (X_t) lands with frequency ϕ_i in each bin \mathcal{X}_i (see Corollary 2).

Note that in the implementation of the algorithm, the penalties θ_t need only be defined up to a normalizing constant, since they only appear in ratios of the form $\theta_t(i)/\theta_t(j)$. We therefore introduce the following notation:

$$\forall i, j \in \{1, \dots, d\} \text{ such that } i \neq j \quad Z_t^{(i,j)} = \log \frac{\theta_t(i)}{\theta_t(j)}$$

and we note Z_t the collection of all the $Z_t^{(i,j)}$. Some intuition behind the study of such ratios comes from considering update (1). With this update, assume that for each i , $\mathbb{E}[\mathbb{1}_{\mathcal{X}_i}(X_t)] = \phi_i$. Then we could easily check that

for each pair (i, j) , $\mathbb{E}[Z_t^{(i,j)} | Z_{t-1}^{(i,j)}] = Z_{t-1}^{(i,j)}$, so this process would be constant on average. The remainder of this paper hinges on two facts: that we can control (Z_t) , in the sense that $Z_t^{(i,j)}/t \rightarrow 0$; and that if we control (Z_t) , then we control the frequencies of visits (ν_t/t) .

More generally, notice that with fixed γ , the pair (X_t, Z_t) forms a homogeneous Markov chain. We would like to prove that its proportion of visits to the set $\mathcal{X}_i \times \mathbb{R}^{d(d-1)}$ converges to some value in $[0, 1]$; one way to prove this is to show that the chain is irreducible. We would then need to check that the limit is indeed the desired frequency ϕ_i for all i . Unfortunately, properties of the joint chain (X_t, Z_t) are difficult to establish due to the complexity of its transition kernel. Finding a so-called drift function for the joint Markov chain is also typically difficult. In general, we are not able to show that the chain is irreducible. In section 3, we prove directly that $Z_t^{1,2}/t \rightarrow 0$ in the special case $d = 2$, under some assumptions. In section 4, we make more restrictive assumptions which imply irreducibility. In both cases, we show the implication of this convergence on the frequencies of visits.

3. Proof when $d=2$. In the following we consider a simple context with only two bins: $d = 2$ and X_t can therefore only be either in \mathcal{X}_1 or in \mathcal{X}_2 . Suppose the current schedule is at $\gamma > 0$, and we want to know whether (FH) is going to be met in finite time (hence γ is fixed here). To simplify notation, in this section we note

$$Z_t = Z_t^{(1,2)} = \log \theta_t(1) - \log \theta_t(2)$$

Using the definition of the penalties (θ_t) and of the counts (ν_t) , we obtain

$$\begin{aligned} Z_t &= Z_0 + [\nu_t(1)f(1, \phi_1, \gamma) + (t - \nu_t(1))f(0, \phi_1, \gamma)] \\ &\quad - [\nu_t(2)f(1, \phi_2, \gamma) + (t - \nu_t(2))f(0, \phi_2, \gamma)] \\ &= Z_0 + \nu_t(1) [f(1, \phi_1, \gamma) - f(0, \phi_1, \gamma)] + tf(0, \phi_1, \gamma) \\ &\quad - [(t - \nu_t(1))f(1, \phi_2, \gamma) + \nu_t(1)f(0, \phi_2, \gamma)] \\ &= Z_0 + \nu_t(1) [f(1, \phi_1, \gamma) - f(0, \phi_1, \gamma) + f(1, \phi_2, \gamma) - f(0, \phi_2, \gamma)] \\ &\quad + t(f(0, \phi_1, \gamma) - f(1, \phi_2, \gamma)) \end{aligned}$$

If we prove that Z_t/t goes to 0 (for instance in mean), this will imply the following convergence of the proportion of visits:

$$\frac{\nu_t(1)}{t} \xrightarrow[t \rightarrow \infty]{} \frac{f(1, \phi_2, \gamma) - f(0, \phi_1, \gamma)}{f(1, \phi_1, \gamma) - f(0, \phi_1, \gamma) + f(1, \phi_2, \gamma) - f(0, \phi_2, \gamma)}$$

(also in mean). Since we want (FH) to be reached in finite time for any precision threshold $c > 0$, we need the proportions of visits to \mathcal{X}_i to converge

to ϕ_i . Hence we want:

$$(3) \quad \frac{f(1, \phi_2, \gamma) - f(0, \phi_1, \gamma)}{f(1, \phi_1, \gamma) - f(0, \phi_1, \gamma) + f(1, \phi_2, \gamma) - f(0, \phi_2, \gamma)} = \phi_1$$

Using the specific forms of $f(\mathbb{1}_{\mathcal{X}_i}(X_t), \phi_i, \gamma)$ for both updates, we can easily see that

- update (1) satisfies equation (3) for any ϕ and γ ;
- in general, update (2) does not satisfy equation (3), except in the special case where $\phi_1 = \phi_2 = 1/2$.

Note that the second point states that the proportions of visits converge to some vector $\tilde{\phi}$ different than ϕ . The vector $\tilde{\phi}$ can be expressed as a function of ϕ : $\tilde{\phi} = g(\phi)$. By numerically or analytically inverting this function g , one can plug $g^{-1}(\phi)$ as an algorithmic parameter, so that the limiting proportions of visits converge to $g(g^{-1}(\phi)) = \phi$. Hence one can use update (2) and get the desired proportions of visits, by plugging $g^{-1}(\phi)$ instead of ϕ in the update.

The rest of the paper is devoted to the proof that Z_t/t goes to 0 under some assumptions. More formally, we state in Theorem 1 what we shall prove in the remainder of this section. This theorem holds for both updates.

THEOREM 1. *Consider the sequence of penalties (θ_t) introduced in Algorithm 2. We define:*

$$Z_t = \log \theta_t(1) - \log \theta_t(2)$$

Then:

$$\frac{Z_t}{t} \xrightarrow[t \rightarrow \infty]{L_1} 0$$

As a consequence, the long run proportion of visits to each bin converges to the desired frequency ϕ for update (1), and not necessarily for update (2). Corollary 2 clarifies the consequence of Theorem 1 on the validity of Algorithm 2.

COROLLARY 2. *When the proportions of visits converge in mean to the desired proportions, the Flat Histogram criterion is reached in finite time for any precision threshold c .*

We already made the simplification of considering the simple case $d = 2$. We make the following assumptions:

ASSUMPTION 1. *The bins are not empty with respect to μ and π :*

$$\forall i \in \{1, 2\} \quad \mu(\mathcal{X}_i) > 0 \text{ and } \pi(\mathcal{X}_i) > 0$$

ASSUMPTION 2. *The state space \mathcal{X} is compact.*

ASSUMPTION 3. *The proposition distribution $Q(x, y)$ is such that:*

$$\exists q_{min} > 0 \quad \forall x \in \mathcal{X} \quad \forall y \in \mathcal{X} \quad Q(x, y) > q_{min}$$

ASSUMPTION 4. *The MH acceptance ratio is bounded from both sides:*

$$\exists m > 0 \quad \exists M > 0 \quad \forall x \in \mathcal{X} \quad \forall y \in \mathcal{X} \quad m < \frac{\pi(y) Q(y, x)}{\pi(x) Q(x, y)} < M$$

Assumption 1 guarantees that the bins are well designed, and if it was not verified, the algorithm would never reach (FH), regardless of the other assumptions. Assumptions 2-4 are for example verified by a gaussian random walk proposal over a compact space, where there is a lower bound on π . We believe that these assumptions can be relaxed to cover the most general Wang-Landau algorithm. Making these four assumptions allows to propose a clearer proof, and we propose hints on how to relax them in Section 6.

We denote by U_t the increment of Z_t , such that for any t :

$$Z_{t+1} = Z_t + U_t = Z_t + f(\mathbb{1}_{\mathcal{X}_1}(X_t), \phi_1, \gamma) - f(\mathbb{1}_{\mathcal{X}_2}(X_t), \phi_2, \gamma)$$

Here with only two bins, the increments U_t can take two different values, $+a$ or $-b$, for some $a > 0$ and $b > 0$ that depend on ϕ and γ . For example, with update (1) :

$$\begin{cases} a &= 2\gamma(1 - \phi_1) > 0 \\ b &= 2\gamma\phi_1 > 0 \end{cases}$$

whereas with update (2) :

$$\begin{cases} a &= \log\left(\frac{1+\gamma(1-\phi_1)}{1-\gamma(1-\phi_1)}\right) > 0 \\ b &= \log\left(\frac{1+\gamma\phi_1}{1-\gamma\phi_1}\right) > 0 \end{cases}$$

and in both cases, if $X_t \in \mathcal{X}_1$ then $U_t = +a$, otherwise $U_t = -b$.

We want to prove that Z_t/t goes to 0, and we are going to prove a stronger result that states, in words, that when Z_t leaves a fixed interval $[\bar{Z}^{lo}, \bar{Z}^{hi}]$, it returns to it in a finite time.

3.1. *Behaviour of (Z_t) outside an interval.* First, lemma 3 states that if Z_t goes above a value \bar{Z}^{hi} , it has a strictly positive probability of starting to decrease, and that when that happens, it keeps on decreasing with a high probability.

LEMMA 3. *With the introduced processes Z_t and U_t , there exists $\epsilon > 0$ such that for all $\eta > 0$, there exists \bar{Z}^{hi} such that, if $Z_t \geq \bar{Z}^{hi}$, we have the following two inequalities:*

$$\begin{aligned} P[U_{t+1} = -b | U_t = +a, Z_t] &> \epsilon \\ P[U_{t+1} = -b | U_t = -b, Z_t] &> 1 - \eta. \end{aligned}$$

PROOF OF LEMMA 3. We start with the first inequality. Let q_{min} be like in Assumption 3.

In terms of events $\{U_t = +a\}$ is equivalent to $\{X_t \in \mathcal{X}_1\}$, by definition. If $X_t \in \mathcal{X}_1$ and $\pi(X_t) > 0$ then:

$$\begin{aligned} K_{\theta_t}(X_t, \mathcal{X}_2) &= \int_{\mathcal{X}_2} K_{\theta_t}(X_t, y) dy \\ &= \int_{\mathcal{X}_2} Q(X_t, y) \rho_{\theta_t}(X_t, y) dy \\ &= \int_{\mathcal{X}_2} Q(X_t, y) \left(1 \wedge \frac{\pi(y)}{\pi(X_t)} \frac{Q(y, X_t)}{Q(X_t, y)} \frac{\theta_t(J(X_t))}{\theta_t(J(y))} \right) dy \\ &= \int_{\mathcal{X}_2} Q(X_t, y) \left(1 \wedge \frac{\pi(y)}{\pi(X_t)} \frac{Q(y, X_t)}{Q(X_t, y)} e^{Z_t} \right) dy \end{aligned}$$

Using Assumption 4, $\frac{\pi(y)}{\pi(x)} \frac{Q(y, x)}{Q(x, y)}$ is bounded from below, hence there exists K_1 such that:

$$\forall k \geq K_1 \quad \forall x, y \in \mathcal{X} \quad \frac{\pi(y)}{\pi(x)} \frac{Q(y, x)}{Q(x, y)} e^k \geq 1.$$

If $Z_t \geq K_1$ and $X_t \in \mathcal{X}_1$, then:

$$K_{\theta_t}(X_t, \mathcal{X}_2) = \int_{\mathcal{X}_2} Q(X_t, y) dy > q_{min} \mu(\mathcal{X}_2).$$

Hence if $Z_t \geq K_1$:

$$\begin{aligned} P[U_{t+1} = -b | U_t = +a, Z_t] &= P[X_{t+1} \in \mathcal{X}_2 | X_t \in \mathcal{X}_1, Z_t] \\ &> q_{min} \mu(\mathcal{X}_2). \end{aligned}$$

We now prove the second inequality. Let us show that for any $\eta > 0$ there exists K_2 such that, provided $Z_t > K_2$:

$$P[U_{t+1} = -b | U_t = -b, Z_t] > 1 - \eta.$$

We have

$$P[U_{t+1} = -b | U_t = -b, Z_t] = P[X_{t+1} \in \mathcal{X}_2 | X_t \in \mathcal{X}_2, Z_t].$$

Again let us first work for a fixed $X_t \in \mathcal{X}_2$.

$$\begin{aligned} K_{\theta_t}(X_t, \mathcal{X}_2) &= 1 - K_{\theta_t}(X_t, \mathcal{X}_1) \\ &= 1 - \left[\int_{\mathcal{X}_1} Q(X_t, y) \rho_{\theta_t}(X_t, y) dy \right] \\ &= 1 - \left[\int_{\mathcal{X}_1} Q(X_t, y) \left(1 \wedge \frac{\pi(y)}{\pi(X_t)} \frac{Q(y, X_t)}{Q(X_t, y)} e^{-Z_t} \right) dy \right] \end{aligned}$$

Using the assumption that the MH ratio $\frac{\pi(y) Q(y, x)}{\pi(x) Q(x, y)}$ is bounded from above, there exists K_2 such that:

$$\forall k \geq K_2 \quad \forall x, y \in \mathcal{X} \quad \frac{\pi(y) Q(y, x)}{\pi(x) Q(x, y)} e^{-k} \leq 1.$$

And hence for $Z_t > K_2$:

$$\begin{aligned} K_{\theta_t}(X_t, \mathcal{X}_2) &= 1 - e^{-Z_t} \int_{\mathcal{X}_1} Q(X_t, y) \frac{\pi(y)}{\pi(X_t)} \frac{Q(y, X_t)}{Q(X_t, y)} dy \\ &> 1 - e^{-Z_t} \int_{\mathcal{X}_1} Q(X_t, y) M dy \\ &> 1 - M e^{-Z_t} \end{aligned}$$

and hence for any η , there is a K_3 greater than K_2 such that for all $Z_t \geq K_3$:

$$K_{\theta_t}(X_t, \mathcal{X}_2) > 1 - \eta$$

We thus obtain:

$$P[U_{t+1} = -b | U_t = -b, Z_t] > 1 - \eta.$$

To conclude we finally define $\epsilon = q_{\min} \mu(\mathcal{X}_2)$ and then for any $\eta > 0$, by taking any \bar{Z}^{hi} greater than $K_1 \vee K_3$ we have both inequalities. \square

Considering the symmetry of the problem, we instantly have the following corollary result. It states that if Z_t goes too low, it has a strictly positive probability of starting to increase, and when that happens, it keeps on increasing with a high probability.

LEMMA 4. *With the introduced processes Z_t and U_t , there exists $\epsilon > 0$ such that for all $\eta > 0$, there exists \bar{Z}^{lo} such that, if $Z_t \leq \bar{Z}^{lo}$, we have the following two inequalities:*

$$\begin{aligned} P[U_{t+1} = +a | U_t = -b, Z_t] &> \epsilon \\ P[U_{t+1} = +a | U_t = +a, Z_t] &> 1 - \eta. \end{aligned}$$




FIG 1. Trajectory of Z (full line with dots) and of \tilde{Z} (dotted line), when these processes go above some level \bar{Z}^{hi} indicated by a horizontal full line. Z goes above the level at time s , and returns below it at time $s+T$, whereas \tilde{Z} stays above the level until time $s+\tilde{T}$, with $T \leq \tilde{T}$.

3.2. *A new process that bounds (Z_t) outside the set.* In this section, the proof introduces a new sequence of increments \tilde{U}_t that bounds U_t , and such that the sequence \tilde{Z}_t using \tilde{U}_t as increments:

$$\tilde{Z}_{t+1} = \tilde{Z}_t + \tilde{U}_t$$

returns to $[\bar{Z}^{lo}, \bar{Z}^{hi}]$ in a finite time whenever it leaves it. It will imply that Z_t also returns to $[\bar{Z}^{lo}, \bar{Z}^{hi}]$ in finite time whenever it leaves it. Figure 1 might help to visualize the proof.

First let us use Lemma 3. We can take $\epsilon < 1/2$ and $\eta < \min(1/2, \epsilon b/a)$. The Lemma gives the existence of an integer K such that if $Z_t \geq K$, we have the following two inequalities:

$$(4) \quad P[U_{t+1} = -b | U_t = +a, Z_t] > \epsilon$$

$$(5) \quad P[U_{t+1} = -b | U_t = -b, Z_t] > 1 - \eta.$$

Suppose that there is some time s such that $Z_{s-1} \leq K$ and $Z_s \geq K$. Note that necessarily $Z_s \in [K, K+a]$. Then we define $\tilde{Z}_s = Z_s$, a new process starting at time s . Let $s+T$ be the first time after s such that $Z_{s+T} \leq K$. We wish to show that $E[T] < \infty$.

We define the sequence of random variables $(Z_t)_{t \geq s}$ defined by $\tilde{Z}_s = Z_s$ and $\tilde{Z}_{t+1} = \tilde{Z}_t + \tilde{U}_t$ for $t > s$, where $(\tilde{U}_t)_{t \geq s}$ is a sequence of random variables taking the values $+a$ or $-b$.

For $s \leq t < T$, \tilde{U}_t is defined as follows:

- if $U_{t+1} = +a$ then $\tilde{U}_{t+1} = +a$;
- if $U_{t+1} = -b$, $U_t = -b$ and $\tilde{U}_t = -b$ then $\tilde{U}_{t+1} = -b$ with probability $p_1 = (1 - \eta)/P[U_{t+1} = -b|U_t = -b, Z_t]$ and $\tilde{U}_{t+1} = +a$ otherwise;
- if $U_{t+1} = -b$, $U_t = +a$ and $\tilde{U}_t = +a$, then $\tilde{U}_{t+1} = -b$ with probability $p_2 = \epsilon/P[U_{t+1} = -b|U_t = +a, Z_t]$ and $\tilde{U}_{t+1} = +a$ otherwise;
- if $U_{t+1} = -b$, $U_t = -b$ and $\tilde{U}_t = +a$, then $\tilde{U}_{t+1} = -b$ with probability $p_3 = \epsilon(1 + P[U_{t+1} = +a|U_t = -b, Z_t])/P[U_{t+1} = -b|U_t = -b, Z_t]$ and $\tilde{U}_{t+1} = +a$ otherwise.

For times $t \geq T$, \tilde{U}_t is a Markov chain independent of U_t and Z_t , with transition matrix

$$\begin{pmatrix} 1 - \epsilon & \epsilon \\ \eta & 1 - \eta \end{pmatrix}$$

where the first state corresponds to $+a$ and the second state to $-b$.

First, let us check that all these probabilities are indeed less than 1. For p_1 , it follows from inequality (5). For p_2 , it follows from inequality (4). For p_3 , we have

$$\epsilon \left(1 + \frac{P[U_{t+1} = +a|U_t = -b, Z_t]}{P[U_{t+1} = -b|U_t = -b, Z_t]} \right) \leq \epsilon \left(1 + \frac{\eta}{1 - \eta} \right) \leq 2\epsilon \leq 1$$

where we used the conditions $\eta < 1/2$ and $\epsilon < 1/2$. Hence (\tilde{U}_t) is well defined.

LEMMA 5. (\tilde{U}_t) is a Markov chain over the space $\{+a, -b\}$ with transition matrix

$$\begin{pmatrix} 1 - \epsilon & \epsilon \\ \eta & 1 - \eta \end{pmatrix}$$

where the first state corresponds to $\{+a\}$ and the second state to $\{-b\}$.

PROOF OF LEMMA 5. We only need to check this for times $t \leq T$. The events $\{\tilde{U}_t = -b\}$ and $\{\tilde{U}_t = -b, U_t = -b\}$ are identical, hence:

$$\begin{aligned} P[\tilde{U}_{t+1} = -b|\tilde{U}_t = -b, Z_t] &= P[\tilde{U}_{t+1} = -b|\tilde{U}_t = -b, U_t = -b, Z_t] \\ &= P[\tilde{U}_{t+1} = -b|U_{t+1} = -b, \tilde{U}_t = -b, U_t = -b, Z_t] \\ &\quad \times P[U_{t+1} = -b|\tilde{U}_t = -b, U_t = -b, Z_t] \\ &= \frac{(1 - \eta)P[U_{t+1} = -b|U_t = -b, Z_t]}{P[U_{t+1} = -b|U_t = -b, Z_t]} \\ &= 1 - \eta. \end{aligned}$$

Note that this does not depend on Z_t .

Similarly:

$$\begin{aligned}
P[\tilde{U}_{t+1} = -b | \tilde{U}_t = +a, U_t = +a, Z_t] \\
&= P[\tilde{U}_{t+1} = -b | U_{t+1} = -b, \tilde{U}_t = +a, U_t = +a, Z_t] \\
&\quad \times P[U_{t+1} = -b | \tilde{U}_t = +a, U_t = +a, Z_t] \\
&= \frac{\epsilon P[U_{t+1} = -b | U_t = +a, Z_t]}{P[U_{t+1} = -b | U_t = +a, Z_t]} \\
&= \epsilon.
\end{aligned}$$

And:

$$\begin{aligned}
P[\tilde{U}_{t+1} = -b | \tilde{U}_t = +a, U_t = -b, Z_t] \\
&= P[\tilde{U}_{t+1} = -b | U_{t+1} = -b, \tilde{U}_t = +a, U_t = -b, Z_t] \\
&\quad \times P[U_{t+1} = -b | \tilde{U}_t = +a, U_t = -b, Z_t] \\
&= \epsilon \left(1 + \frac{P[U_{t+1} = +a | U_t = -b, Z_t]}{P[U_{t+1} = -b | U_t = -b, Z_t]} \right) \\
&\quad \times P[U_{t+1} = -b | U_t = -b, Z_t] \\
&= \epsilon (P[U_{t+1} = -b | U_t = -b, Z_t] + P[U_{t+1} = +a | U_t = -b, Z_t]) \\
&= \epsilon.
\end{aligned}$$

These last two calculations result in:

$$P[\tilde{U}_{t+1} = -b | \tilde{U}_t = +a] = \epsilon$$

with no dependence on Z_t (or U_t). □

The previous lemma is central to the proof, and especially the lack of dependence on Z_t . We always have $\tilde{U}_s = +a$, since $U_s = +a$. Hence for each $t \geq s$, the distribution of \tilde{U}_t depends only on η and ϵ , and implicitly on the threshold K , but not on the value of Z_s . Hence (\tilde{U}_t) has the same law, every time the process (Z_t) goes above K .

3.3. Conclusion: proof of Theorem 1 and Corollary 2. Let us now use the bounding process (\tilde{Z}_t) to control the time spent by (Z_t) above K .

LEMMA 6. *There exists $\tau \in \mathbb{R}$ such that, for all times s such that $Z_{s-1} \leq K$ and $Z_s \geq K$, and defining T by $T = \inf_{d \geq 0} \{Z_{s+d} \leq K\}$, then:*

$$\mathbb{E}[T] \leq \tau$$

PROOF OF LEMMA 6. The Markov chain (\tilde{U}_t) admits the following stationary distribution:

$$\pi_{\tilde{U}} = \left(\frac{\eta}{\epsilon + \eta}, \frac{\epsilon}{\epsilon + \eta} \right)$$

Let us denote by \tilde{T} the time spent by (\tilde{Z}_t) over K , that is:

$$\tilde{T} = \inf_{d \geq 0} \left\{ \sum_{t=s+1}^{s+d} \tilde{U}_t \leq -a \right\}$$

Remember that $\tilde{Z}_s = Z_s \in [K, K+a]$, hence $\tilde{Z}_{s+\tilde{T}} \leq K$ (whatever the value of Z_s). Now, our choice of η results in $a\eta < b\epsilon$ which implies $E[\tilde{T}] < \infty$ (Norris, 1998). Let $\tau = E[\tilde{T}]$. Note that since the law of (\tilde{U}_t) does not depend on the value of Z_s , τ does not depend on Z_s .

Since, for $t \leq T$, we impose that “if $U_{t+1} = +a$ then $\tilde{U}_{t+1} = +a$ ”, it follows that $\forall t \leq T, U_t \leq \tilde{U}_t$. Consequently $\forall t \leq T, Z_t \leq \tilde{Z}_t$ and hence $T \leq \tilde{T}$. Note that (the distribution of) T depends on the exact value of Z_s , but that \tilde{T} as we have defined it has a fixed distribution. We have $E[T] \leq \tau$ (whatever the value Z_s). \square

PROOF OF THEOREM 1. Let us define the following sequence of indices:

$$S_1 = \inf_{s \geq 0} \{Z_{s-1} \leq K \text{ and } Z_s \geq K\}; S_k = \inf_{s \geq S_{k-1}} \{Z_{s-1} \leq K \text{ and } Z_s \geq K\}$$

The sequence (S_k) represents the times at which the process (Z_t) goes above K . Moreover let us introduce the sequence of time spent above K :

$$T_k = \inf_{s \geq 0} \{Z_{S_k+s-1} \geq K \text{ and } Z_{S_k+s} \leq K\}$$

We have $Z_{S_k} \in [K, K+a]$. Define $k(t)$ such that $S_{k(t)} \leq t < S_{k(t)+1}$. Either $Z_t \leq K$ or $Z_t > K$. In the latter case, $Z_t \leq Z_{S_{k(t)}} + aT_{k(t)}$. Clearly in any case:

$$(6) \quad \mathbb{E}[Z_t] \leq (K+a) + a\tau.$$

A similar reasoning on the the lower bound leads to K' and $\tau' < \infty$ such that

$$(7) \quad \mathbb{E}[Z_t] \geq (K' - b) - b\tau'.$$

Inequalities (7) and (6) imply

$$\mathbb{E} \left[\frac{Z_t}{t} \right] \rightarrow 0.$$

\square

As stated at the beginning of the section, for update (1) the convergence $Z_t/t \rightarrow 0$ (in mean) implies the convergence of the proportions (ν_t/t) to ϕ (also in mean). We now show that this ensures that the Flat Histogram is reached in finite time.

PROOF OF COROLLARY 2. For a fixed threshold c , recall that (FH) being reached at time t corresponds to the event:

$$FH_t = \{\forall i \in \{1, \dots, d\} \quad \left| \frac{\nu_t(i)}{t} - \phi_i \right| < c\}$$

We will only use the convergence in probability of the proportions to ϕ for all i :

$$\frac{\nu_t(i)}{t} \xrightarrow[t \rightarrow \infty]{\mathbb{P}} \phi_i$$

which implies:

$$\forall \varepsilon > 0 \exists N \in \mathbb{N} \forall t \geq N \quad \mathbb{P}(FH_t) \geq 1 - \varepsilon$$

We can hence define a stopping time T^{FH} corresponding to the first (FH) being reached:

$$T^{\text{FH}} = \inf_{t \geq 0} \{FH_t\}$$

and some $\varepsilon > 0$ such that:

$$\exists N \in \mathbb{N} \forall n \geq N \quad \mathbb{P}(T^{\text{FH}} \leq N + n) \geq \varepsilon$$

Using Lemma 10.11 of [Williams \(1991\)](#), the expectation of T^{FH} is then finite. \square

4. Proof when $d \geq 2$. In this section we extend the proof to the more general case $d \geq 2$. Having proved that for $d = 2$, only update (1) is valid, we now focus on this update and omit update (2).

We consider the log penalties defined for update (1) by:

$$\log \theta_t(i) = \nu_t(i)(1 - \phi_i) - (t - \nu_t(i))\phi_i = \nu_t(i) - t\phi_i$$

where $\nu_t(i)$ is the number of visits of (X_t) in \mathcal{X}_i . We assume without loss of generality that $\log \theta_0 = 0$. Then $(X_t, \log \theta_t)$ is a Markov chain, by definition of the WL algorithm. We first prove that $(X_t, \log \theta_t)$ is λ -irreducible, for a sigma-finite measure λ . We will require the following additional assumption on the desired frequencies ϕ .

ASSUMPTION 5. *The desired frequencies are rational numbers:*

$$\phi = (\phi_1, \dots, \phi_d) \in \mathbb{Q}^d.$$

LEMMA 7. *Let Θ be the following subset of \mathbb{R}^d :*

$$\Theta = \{z \in \mathbb{R}^d : \exists(n_1, \dots, n_d) \in \mathbb{N}^d \ z_i = n_i - \phi_i S_n \text{ where } S_n = \sum_{j=1}^d n_j\}$$

Then denoting by λ the product of the Lebesgue measure μ on \mathcal{X} and of the counting measure on Θ , $(X_t, \log \theta_t)$ is λ -irreducible.

PROOF. The proof essentially comes from Bézout's lemma, and is detailed in the Appendix. Note however that it relies on Assumption 5, that was not required for the case $d = 2$. Although not a very satisfying assumption, which is likely not to be necessary for proving the occurrence of (FH) in finite time, it seems to be necessary for the irreducibility of $(X_t, \log \theta_t)$, at least with respect to a standard sigma-finite measure. In any case, this assumption is not restrictive in practice. \square

Since this chain is λ -irreducible, the proportion of visits to any λ -measurable set of $\mathcal{X} \times \Theta$ converges to a limit in $[0, 1]$. This implies that the vector $(\nu_t(i)/t)$ converges to some vector (p_i) . The following is a *reductio ad absurdum*.

Suppose that for some $i \in \{1, \dots, d\}$, $p_i \neq \phi_i$. Since the vectors p and ϕ both sum to 1, this means that for some i , $p_i < \phi_i$: such a state i is visited less than the desired frequency.

Let $\{i_1, i_2, \dots\} = \operatorname{argmin}_{1 \leq j \leq d} (p_j - \phi_j)$. Then for any i_k and for $j \notin \{i_1, i_2, \dots\}$, we have

$$Z_t^{j, i_k} = -\nu_t(i_k) + \nu_t(j) + t(\phi_{i_k} - \phi_j) \sim t(-p_{i_k} + \phi_{i_k} + p_j - \phi_j) \rightarrow \infty$$

This implies:

$$\forall K > 0 \exists T \in \mathbb{N} \forall t > T \quad Z_t^{j, i_k} > K$$

Now consider the stochastic process (U_t) such that

- $U_t = -b$ if $J(X_t) \in \{i_1, i_2, \dots\}$
- $U_t = +a$ otherwise

for some real numbers a and b . Recall that the function J is such that if $X_t \in \mathcal{X}_i$ then $J(X_t) = i$.

Let ϵ be such that when $X_t \notin \mathcal{X}_{i_1} \cup \mathcal{X}_{i_2} \cup \dots$, there is probability at least ϵ of proposing in $\mathcal{X}_{i_1} \cup \mathcal{X}_{i_2} \cup \dots$. For large enough K , these proposals

will always be accepted. As before, for large enough K , we can make the probability η of leaving $\mathcal{X}_{i_1} \cup \mathcal{X}_{i_2} \cup \dots$ as small as we wish.

Using the exact same reasoning as in Section 3, we can construct a process (\tilde{U}_t) which is a Markov chain with transition matrix

$$\begin{pmatrix} 1 - \epsilon & \epsilon \\ \eta & 1 - \eta \end{pmatrix}$$

and with $U_t < \tilde{U}_t$ almost surely. Therefore for $t > T$, (U_t) decreases on average, hence (Z_t^{j,i_k}) decreases on average, which contradicts the assumption that it goes to infinity. Hence for all i , $p_i = \phi_i$.

5. Illustration of Theorem 1 on a toy example. Let us show the consequences of Theorem 1 on a simple example. We consider as the target distribution the standard normal distribution truncated to the set $\mathcal{X} = [-10, 10]$. We use a Gaussian random walk proposal, with unit standard deviation. Finally we arbitrarily split the state space in $\mathcal{X}_1 = [-10, 0]$ and $\mathcal{X}_2 =]0, 10]$, and we set the desired frequencies to be $\phi = (0.75, 0.25)$. Figure 2 shows the results of the Wang-Landau algorithm. Using update (1) and 200,000 iterations, we obtain the histogram of Figure 2(a). Figure 2(b) shows the convergence of the proportions of visits to each bin, using update (1). The dotted horizontal lines indicate ϕ , and we can check that the observed proportions of visits converge towards it.

Figure 2(c) shows a similar plot, this time using update (2). Again, the desired frequencies are represented by dotted lines. Using the left hand side of equation (3), we can calculate the theoretical limit of the observed proportion of visits in each bin, which for $\gamma = 1$ and $\phi = (0.75, 0.25)$, is approximately equal to $(0.79, 0.21)$. Hence for a precision threshold c equal to e.g. 1%, the occurrence of (FH) is not likely to occur if one uses update (2).

As expected, update (1) leads to convergence to the desired frequencies but update (2) does not.

6. Discussion. As seen in Theorem 1 and Corollary 2 of Section 3, update (1) is valid, in the sense that the frequencies of visits of the chain (X_t) converges towards ϕ . Consequently (FH) is met in finite time, for any threshold $c > 0$.

Regarding the proof of Theorem 1 in the case $d > 2$, we assume that the desired frequencies ϕ are rationals (Assumption 5), which allows to prove that the Markov chain generated by the algorithm (X_t, Z_t) is λ -irreducible for some sigma-finite measure λ . However, our proof requires mainly that the proportions of visits of (X_t) to any bin \mathcal{X}_i converge, which is equivalent to




FIG 2. Results of the Wang-Landau algorithm using two different updates of the penalties. Histogram of the generated sample using update (1), with a vertical line showing the binning (left). Convergence of the proportions of visits to each bin, using update (1) (middle) and using update (2) (right). The dotted horizontal lines represent the desired frequencies.

the convergence of (Z_t/t) . We believe that results on Random Walks in Random Environments (Zeitouni, 2006) would allow to remove the rationality assumption.

Assumptions 2–4 could be relaxed by using the well-known properties of the Metropolis–Hastings algorithm, from which we did not take advantage here. More precisely, note that the Wang-Landau transition kernel differs from the Metropolis–Hastings only when the proposed points, generated through $Q(\cdot, \cdot)$, land in a different bin than the current position of the chain. Otherwise, the kernel behaves like a Metropolis–Hastings targeting π . Hence under some weaker assumptions than the one we have formulated here, it has recurrence properties.

To conclude, we have shown that for fixed γ , the Flat Histogram criterion is reached in finite time for certain updates. For other updates, the observed frequencies do not converge to the desired frequencies, and so there is a non-zero probability that the Flat Histogram criterion will never be verified. Note that we do not make any claims about the distribution of the sample inside each of the bins \mathcal{X}_i at fixed γ .

Acknowledgements. The authors thank Luke Bornn, Arnaud Doucet, Anthony Lee, Éric Moulines, Christian P. Robert and an anonymous reviewer for helpful comments.

References.

- ANDRIEU, C., MOULINES, E. and PRIOURET, P. (2006). Stability of stochastic approximation under verifiable conditions. *SIAM Journal on control and optimization* **44** 283–312.
- ANDRIEU, C. and THOMS, J. (2008). A tutorial on adaptive MCMC. *Statistics and Computing* **18** 343–373.
- ATCHADÉ, Y. F. and LIU, J. (2010). The Wang-Landau algorithm in general state spaces: applications and convergence analysis. *Statistica Sinica* **20** 209–233.
- ATCHADÉ, Y., FORT, G., MOULINES, E. and PRIOURET, P. (2009). Adaptive Markov Chain Monte Carlo: Theory and Methods.
- BORNN, L., JACOB, P., DEL MORAL, P. and DOUCET, A. (2011). An adaptive interacting Wang-Landau algorithm for automatic density exploration. *arXiv preprint arXiv:1109.3829*.
- CUNHA NETTO, A., SILVA, C., CAPARICA, A. and DICKMAN, R. (2006). Wang-Landau sampling in three-dimensional polymers. *Brazilian journal of physics* **36** 619–622.
- FORT, G., MOULINES, E., PRIOURET, P. and VANDEKERKHOVE, P. (2011). A central limit theorem for adaptive and interacting Markov chains. *ArXiv e-prints*.
- HASTINGS, W. K. (1970). Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* **57** 97–109.
- KESTEN, H. (1958). Accelerated Stochastic Approximation. *Annals of Mathematical Statistics* **29** 41–59.
- LIANG, F. (2005). A Generalized Wang-Landau Algorithm for Monte Carlo Computation. *Journal of the American Statistical Association* **100** 1311–1327.
- MALAKIS, A., KALOZOUNIS, P. and TYRASKIS, N. (2006). Monte Carlo studies of the square Ising model with next-nearest-neighbor interactions. *The European Physical Journal B-Condensed Matter and Complex Systems* **50** 63–67.
- NGO, V. T. and DIEP, H. (2008). Phase transition in Heisenberg stacked triangular antiferromagnets: End of a controversy. *Physical Review E* **78** 031119.
- NORRIS, J. R. (1998). *Markov Chains*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, UK.
- ROBERT, C. P. and CASELLA, G. (2004). *Monte Carlo Statistical Methods*. Springer.
- SILVA, C., CAPARICA, A. and PLASCAK, J. (2006). Wang-Landau Monte Carlo simulation of the Blume-Capel model. *Physical Review E* **73** 036702.
- TIERNEY, L. (1998). A note on Metropolis-Hastings kernels for general state spaces. *The Annals of Applied Probability* **8** 1–9.
- WANG, F. and LANDAU, D. (2001a). Determining the density of states for classical statistical models: A random walk algorithm to produce a flat histogram. *Physical Review E* **64** 56101.
- WANG, F. and LANDAU, D. (2001b). Efficient, multiple-range random walk algorithm to calculate the density of states. *Physical Review Letters* **86** 2050–2053.
- WILLIAMS, D. (1991). *Probability With Martingales*. Cambridge University Press.
- ZEITOUNI, O. (2006). Random walks in random environments. *Journal of Physics A: Mathematical and General* **39** R433.

APPENDIX A: PROOF OF LEMMA 6

Let $\Theta \subset \mathbb{R}^d$ be the set of possibly reachable values of the process $(\log \theta_t)$. We define it by:

$$\Theta = \left\{ z \in \mathbb{R}^d : \exists (n_1, \dots, n_d) \in \mathbb{N}^d \quad z_i = n_i - \phi_i S_n \text{ where } S_n = \sum_{j=1}^d n_j \right\}$$

We want to prove the existence of a measure λ on $\mathcal{X} \times \Theta$ such that the Markov chain $(X_t, \log \theta_t)$ is λ -irreducible. Denote by μ the Lebesgue measure on \mathcal{X} and let $A \in \mathcal{B}(\mathcal{X})$ such that $\mu(A) > 0$, and let $z^* \in \Theta$. Let us show that for any time s at which $X_s = x_s \in \mathcal{X}$ and $\log \theta_s = z_s \in \Theta$, there exists $t > 0$ such that $X_{s+t} \in A$ and $\log \theta_{s+t} = z^*$ with strictly positive probability. This will prove the λ -irreducibility of $(X_t, \log \theta_t)$ where λ is the product of the Lebesgue measure μ on \mathcal{X} and the counting measure on Θ .

Note first that for any $n = (n_1, \dots, n_d) \in \mathbb{N}^d$, the process (X_t) can visit exactly n_i times each set \mathcal{X}_i (for all i) between some time $s+1$ and some time $s + \sum_{i=1}^d n_i$, since there is always a non-zero probability of X_{t+1} visiting any \mathcal{X}_i given X_t and $\log \theta_t$ (using Assumptions 3 on the proposal distribution and the form of the MH kernel). More formally, given any $n \in \mathbb{N}^d$ and any time s , denoting $S_n = \sum_{i=1}^d n_i$:

$$(8) \quad \mathbb{P} \left(\forall i \in \{1, \dots, d\} \quad \sum_{t=s+1}^{s+S_n} \mathbb{I}_{\mathcal{X}_i}(X_t) = n_i \mid X_s, \log \theta_s \right) > 0$$

Furthermore since $\mu(A) > 0$ and since $(\mathcal{X}_i)_{i=1}^d$ is a partition of \mathcal{X} (satisfying Assumption 1 on non-empty bins), there exists $B \subset A$ such that $B \subset \mathcal{X}_{i^*}$ for some $i^* \in \{1, \dots, d\}$ and $\mu(B) > 0$. We are going to prove the following statement, which means that there is a “path” between any pair of points in Θ :

LEMMA 8.

$$\forall z^1, z^2 \in \Theta \exists n \in \mathbb{N}^d \forall i \in \{1, \dots, d\} \quad z_i^1 + n_i - \left(\sum_{j=1}^d n_j \right) \phi_i = z_i^2$$

Then we will conclude as follows: the Markov chain can go from any (x_s, z_s) to some (x_{s+t-1}, z_{s+t-1}) where z_{s+t-1} can be anywhere in Θ , and then in one final step to (x_{s+t}, z_{s+t}) such that $x_{s+t} \in B$ and $z_{s+t} = z^*$, since z_{s+t-1} can be chosen such that $z_{s+t} = z^*$ when $x_{s+t} \in B \subset \mathcal{X}_{i^*}$.

PROOF OF LEMMA 8. The structure of the proof is the following: we prove that $(\log \theta_t)$ can go from 0 to 0, then from any $z \in \Theta$ to 0, and the possibility of going from 0 to any $z \in \Theta$ comes from the definition of Θ .

Suppose that $\log \theta_0 = (0, \dots, 0)$, and let us prove that the process $(\log \theta_t)$ can go back to 0, ie let us find a vector $n = (n_1, \dots, n_d) \in \mathbb{N}^d$ such that

$$\forall i \in \{1, \dots, d\} \quad 0 = n_i - \phi_i S_n \quad \text{where} \quad S_n = \sum_{j=1}^d n_j$$

Under the rationality assumption on ϕ (Assumption 5), there exists $(a_1, \dots, a_d) \in \mathbb{N}^d$ and $b \in \mathbb{N}$ such that $\phi_i = a_i/b$ for all i . Now define $n \in \mathbb{N}^d$ as follows:

$$\forall i \in \{1, \dots, d\} \quad n_i = k \prod_{j=1, j \neq i}^d \frac{1}{a_j}$$

where $k \in \mathbb{N}$ is such that $n_i \in \mathbb{N}$ for all i . Then, using $\sum_{j=1}^d \phi_j = 1$ one can readily check that:

$$\forall i \in \{1, \dots, d\} \quad n_i - \phi_i \left(\sum_{j=1}^d n_j \right) = 0$$

Hence the vector n defines a possible path for $(\log \theta_t)$ between 0 and 0, in $S_n = \sum_{j=1}^d n_j$ steps, with a strictly positive probability (using Equation (8)).

A similar reasoning allows to find a possible path from any $z \in \Theta$ to 0. For such a $z \in \Theta$, there exists $(m_1, \dots, m_d) \in \mathbb{N}^d$ such that

$$(9) \quad \forall i \in \{1, \dots, d\} \quad z_i = m_i - S_m a_i / b \text{ where } S_m = \sum_{j=1}^d m_j$$

We wish to show that there exists $(k_1, \dots, k_d) \in \mathbb{N}^d$ such that $k_i - S_k a_i / b = -z_i$ for all i , where $S_k = \sum_{j=1}^d k_j$. To construct (k_1, \dots, k_d) , we use the already introduced vector (n_1, \dots, n_d) such that $n_i - S_n a_i / b = 0$ for all i , where $S_n = \sum_{j=1}^d n_j$. Putting this together with (9), we get for any $C \in \mathbb{N}$:

$$(10) \quad -z_i + C * 0 = -m_i + C * n_i - \frac{a_i}{b} (C S_n - S_m)$$

For C large enough, for all i , $C n_i - m_i > 0$. We simply take $k_i = C n_i - m_i$ for all i . This proves that starting from a point $z \in \Theta$ (by definition reachable from 0), $(\log \theta_t)$ can reach 0 again. \square

CEREMADE – UNIVERSITÉ PARIS DAUPHINE
PLACE DU MARÉCHAL DE LATTRE DE TASSIGNY
75016 PARIS, FRANCE