

Efficient approximation of branching random walk Gibbs measures*

Fu-Hsuan Ho^{†‡} Pascal Maillard^{†‡}

Abstract

Disordered systems such as spin glasses have been used extensively as models for high-dimensional random landscapes and studied from the perspective of optimization algorithms. In a recent paper by L. Addario-Berry and the second author, the *continuous random energy model (CREM)* was proposed as a simple toy model to study the efficiency of such algorithms. The following question was raised in that paper: what is the threshold β_G , at which sampling (approximately) from the Gibbs measure at inverse temperature β becomes algorithmically hard?

This paper is a first step towards answering this question. We consider the branching random walk, a time-homogeneous version of the continuous random energy model. We show that a simple greedy search on a renormalized tree yields a linear-time algorithm which approximately samples from the Gibbs measure, for every $\beta < \beta_c$, the (static) critical point. More precisely, we show that for every $\varepsilon > 0$, there exists such an algorithm such that the specific relative entropy between the law sampled by the algorithm and the Gibbs measure of inverse temperature β is less than ε with high probability.

In the supercritical regime $\beta > \beta_c$, we provide the following hardness result. Under a mild regularity condition, for every $\delta > 0$, there exists $z > 0$ such that the running time of any given algorithm approximating the Gibbs measure stochastically dominates a geometric random variable with parameter $e^{-z\sqrt{N}}$ on an event with probability at least $1 - \delta$.

Keywords: algorithmic hardness; branching random walk; Gibbs measure; Kullback–Leibler divergence; sampling algorithm.

MSC2020 subject classifications: 68Q17; 82D30; 60K35; 60J80.

Submitted to EJP on July 22, 2021, final version accepted on May 18, 2022.

*Supported in part by grants ANR-20-CE92-0010-01 and ANR-11-LABX-0040 (ANR program “Investissements d’Avenir”).

[†]Institut de Mathématiques de Toulouse, CNRS UMR5219.

[‡]Postal address: Institut de Mathématiques de Toulouse, Université Toulouse 3 Paul Sabatier, 118 Route de Narbonne, 31062 Toulouse Cedex 9, France. E-mail: fu-hsuan.ho@math.univ-toulouse.fr, pascal.maillard@math.univ-toulouse.fr

1 Introduction

We consider the following family of branching random walks. An initial particle is located at the origin. It gives birth to d child particles, $d \geq 2$, scattering on the real line, and each of the child particles produces d child particles again, and so on. The displacement of each particle is independent of the past of the process and of the displacements of its siblings. The genealogy of the particles can be represented by a d -ary tree \mathbb{T}_N , identifying particles with the vertices of the tree. We denote by X_v the location of a particle $v \in \mathbb{T}_N$.

Addario-Berry and Maillard [1] considered algorithms that, for a given $x > 0$, find a leaf v of \mathbb{T}_N such that $X_v \geq xN$ in the framework of the continuous random energy model (CREM). This is a binary time-inhomogeneous branching random walk with Gaussian displacements. More precisely, the CREM is a Gaussian process whose covariance function is given by

$$\mathbb{E}[X_v X_w] = A\left(\frac{|v \wedge w|}{N}\right), \quad \forall v, w \in \mathbb{T}_n$$

where $A : [0, 1] \rightarrow [0, 1]$ is an increasing function with $A(0) = 0$ and $A(1) = 1$ and $|v \wedge w|$ is the depth of the most recent common ancestor of v and w . The authors proved the existence of a threshold x_* such that the following holds: a) for every $x < x_*$, there exists a polynomial-time algorithm that can accomplish the task with high probability, b) for every $x > x_*$, every such algorithm has a running time which is at least exponential in N with high probability. The authors also raised the question of the complexity of sampling a *typical* vertex of value roughly xN , which can be interpreted as sampling a vertex according to a Gibbs measure with a certain parameter β depending on x .

The present work attacks this problem in the simpler setting of the (homogeneous) branching random walk, corresponding to the case $A(x) = x$ of the CREM in the special case of Gaussian displacements. The Gibbs measure is a probability measure on the leaves v of \mathbb{T}_N with weight proportional to $e^{\beta X_v}$, where $\beta > 0$ is a given parameter called the *inverse temperature*. We show that there exists a threshold $\beta_c > 0$ such that the following holds: a) in the subcritical regime $\beta < \beta_c$, there exists a linear-time algorithm such that with high probability, the specific relative entropy between the law sampled by the algorithm and the Gibbs measure of inverse temperature β is arbitrarily small, b) in the supercritical regime $\beta > \beta_c$, under a mild regularity condition, we show that with high probability, the running time of any given algorithm approximating the Gibbs measure in this sense is at least stretched exponential in N .

1.1 Notations and main results

Let \mathbb{T} be a rooted d -ary tree, where $d \geq 2$. The depth of a vertex $v \in \mathbb{T}$ is denoted by $|v|$. We denote the root by \emptyset , and any vertex v with depth $n \geq 1$ is indexed by a string $v_1 \cdots v_n \in \{1, \dots, d\}^n$. For any $v, w \in \mathbb{T}$, we write $v \leq w$ if v is a prefix of w and write $v < w$ if v is a prefix of w strictly shorter than w . We denote \mathbb{T}_n to be the subtree of \mathbb{T} containing vertices of depth less or equal to n and $\partial\mathbb{T}_n$ to be the leaves of \mathbb{T}_n .

Let $\mathbf{Y} = (Y_0, \dots, Y_{d-1})$ be a d -dimensional random vector. Let $(\mathbf{Y}^v)_{v \in \mathbb{T}}$ be iid copies of \mathbf{Y} where $\mathbf{Y}^v = (Y_{v0}, \dots, Y_{v(d-1)})$ – this uniquely defines Y_u for every $u \in \mathbb{T} \setminus \{\emptyset\}$. Define the process $\mathbf{X} = (X_v)_{v \in \mathbb{T}}$ by

$$\begin{cases} X_\emptyset = 0, \\ X_v = \sum_{\emptyset < w \leq v} Y_w, \quad |v| \geq 1. \end{cases}$$

The process \mathbf{X} is called the *branching random walk with increments \mathbf{Y}* .

It is well-known that \mathbf{X} has the branching property: let $\mathcal{F} = (\mathcal{F}_n)_{n \geq 0}$ be its natural filtration. For any $v \in \mathbb{T}$ with $|v| = n$, define

$$\mathbf{X}^v = (X_w^v)_{|w| \geq 0} = (X_{vw} - X_v)_{|w| \geq 0}.$$

Then $(\mathbf{X}^v)_{|v|=n}$ are iid copies of \mathbf{X} and independent of \mathcal{F}_n .

Gibbs measures Define the following function

$$\varphi(\beta) = \log \left(\mathbb{E} \left[\sum_{i=0}^{d-1} e^{\beta Y_i} \right] \right) \in (-\infty, \infty], \quad \beta \in \mathbb{R},$$

and set $\mathcal{D}(\varphi) = \{\beta \in \mathbb{R} : \varphi(\beta) < \infty\}$. It is well known that φ is convex and that it is smooth on $\mathcal{D}(\varphi)^\circ$, the interior of $\mathcal{D}(\varphi)$. Throughout the article, we assume the following.

Assumption 1.1. $0 \in \mathcal{D}(\varphi)^\circ$.

For $n \in \mathbb{N}$, define the following (normalized) partition functions

$$\begin{aligned} W_{\beta,n} &= \sum_{|v|=n} e^{\beta X_v - \varphi(\beta)n}, \quad n \geq 0, \\ W_{\beta,n}^u &= \sum_{|w|=n} e^{\beta X_w^u - \varphi(\beta)n}, \quad u \in \mathbb{T} \text{ and } n \geq 0. \end{aligned}$$

Note that for every $m \leq n$,

$$W_{\beta,n} = \sum_{|u|=m} e^{\beta X_u - \varphi(\beta)m} \cdot W_{\beta,n-m}^u. \tag{1.1}$$

For $\beta \in \mathbb{R}$ and $n \in \mathbb{N}$, we now define the Gibbs measure of parameter β on \mathbb{T}_n to be

$$\mu_{\beta,n}(u) = e^{\beta X_u - \varphi(\beta)m} \cdot \frac{W_{\beta,n-m}^u}{W_{\beta,n}}, \quad |u| = m \leq n. \tag{1.2}$$

Note that $\mu_{\beta,n}$ is usually defined on $\partial\mathbb{T}_n$ only, but it will be helpful to define it on the whole tree \mathbb{T}_n . By (1.1), for every $m \leq n$, the restriction of $\mu_{\beta,n}$ to $\partial\mathbb{T}_m$ is a probability measure. Similarly, we can define

$$\mu_{\beta,n}^v(u) = e^{\beta X_u^v - \varphi(\beta)m} \cdot \frac{W_{\beta,n-m}^{vu}}{W_{\beta,n}^v}, \quad v \in \mathbb{T} \text{ and } |u| = m \leq n. \tag{1.3}$$

The free energy of the branching random walk has been calculated by Derrida and Spohn [16] (and can also be deduced from Biggins [10])

$$\lim_{n \rightarrow \infty} \frac{1}{\beta n} \log \sum_{|v|=n} e^{\beta X_v} = \begin{cases} \frac{1}{\beta} \varphi(\beta) & \text{if } \beta \in (0, \beta_c) \\ \frac{1}{\beta_c} \varphi(\beta_c) & \text{if } \beta \geq \beta_c, \end{cases}$$

where the limit is meant to be in probability and where the critical inverse temperature β_c is defined by

$$\beta_c = \sup\{\beta \in \mathcal{D}(\varphi)^\circ \mid \beta \varphi'(\beta) < \varphi(\beta)\} \in (0, \infty].$$

We will mostly be interested in the phase $\beta < \beta_c$. In this phase, we recall the following result.

Fact 1.2 ([21, 11, 24]). *If $\beta < \beta_c$, then the martingale $(W_{\beta,n})_{n \geq 0}$ is uniformly integrable. In fact, there exists a (strictly) positive random variable $W_{\beta,\infty}$ with $\mathbb{E}[W_{\beta,\infty}] = 1$ and such that $W_{\beta,n} \rightarrow W_{\beta,\infty}$ almost surely and in L^1 as $n \rightarrow \infty$.*

Algorithms We define an algorithmic model similar to the ones in [28, 1]. Let $N \in \mathbb{N}$. A random sequence $v = (v(k))_{k \geq 0}$ taking values in \mathbb{T}_N is called a (*randomized*) *algorithm* if $v(0) = \emptyset$ and $v(k+1)$ is $\tilde{\mathcal{F}}_k$ -measurable for every $k \geq 0$. Here,

$$\tilde{\mathcal{F}}_k = \sigma(v(1), \dots, v(k); X_{v(1)}, \dots, X_{v(k)}; U_1, \dots, U_{k+1})$$

where $(U_k)_{k \geq 1}$ is a sequence of iid uniform random variables on $[0, 1]$, independent of the branching random walk \mathbf{X} . Roughly speaking, the filtration $\tilde{\mathcal{F}} = (\tilde{\mathcal{F}}_k)_{k \geq 0}$ contains all information about everything we have queried so far, as well as the additional randomness needed to choose the next vertex. We further suppose that there exists a stopping time τ with respect to the filtration $\tilde{\mathcal{F}}$ and such that $v(\tau) \in \partial\mathbb{T}_N$. We call τ the *running time* and $v(\tau)$ the *output* of the algorithm. The *law of the output* is the (random) distribution of $v(\tau)$, conditioned on the branching random walk.

Often, we will consider a family of algorithms indexed by N , which we also call an algorithm by abuse of notation. We say that the algorithm is a *polynomial-time* algorithm if there exists a (deterministic) polynomial $P(N)$ such that $\tau \leq P(N)$ almost surely.

Fix $M \in \mathbb{N}$. Given a configuration of the branching random walk of depth N , consider the following algorithm (See Figure 1):

Algorithm 1: Recursive sampling on M -renormalized tree

```

set  $v = \emptyset$ ;
while  $|v| < N$  do
    choose  $w$  with  $|w| = M \wedge (N - |v|)$  according to the Gibbs measure  $\mu_{\beta, M \wedge (N - |v|)}^v$ ;
    replace  $v$  with  $vw$ ;
output  $v$ 
    
```




Figure 1: Schematic illustration of Algorithm 1 with $N = 5$ and $M = 2$ on a given configuration of the binary branching random walk. The squares are the vertices that the algorithm chooses at each step, and the thick line is the path from the root to the sampled vertex. The bar charts represent the Gibbs measures from which the algorithm samples at each step, and the check mark indicates which vertex is chosen. Note that the histograms are not drawn to scale.

Remark 1.3. It is easy to see that Algorithm 1 can be formally written as a randomized algorithm according to the above algorithmic model. Furthermore, its running time is deterministic and bounded by $\lceil N/M \rceil 2^M$. The law of its output is a random probability measure $\mu_{\beta, M, N}$ on $\partial\mathbb{T}_N$ that can be recursively defined as follows:

$$\begin{aligned} \mu_{\beta, M, 0}(\emptyset) &= 1 \\ \mu_{\beta, M, N \wedge (K+1)M}(vw) &= \mu_{\beta, M, KM}(v) \cdot \mu_{\beta, M \wedge (N - KM)}^v(w) \end{aligned} \tag{1.4}$$

for all $|v| = KM$, $|w| = M \wedge (N - KM)$ and $0 \leq K \leq \lfloor \frac{N}{M} \rfloor$.

Approximation and threshold Given two probability measures P and Q defined on a discrete space Ω , the entropy of Q and the Kullback–Leibler divergence from Q to P are respectively defined by

$$H(Q) = \sum_{\omega \in \Omega} Q(\omega) \cdot \log \left(\frac{1}{Q(\omega)} \right) \tag{1.5}$$

$$\mathbf{d}(P \parallel Q) = \sum_{\omega \in \Omega} P(\omega) \cdot \log \left(\frac{P(\omega)}{Q(\omega)} \right). \tag{1.6}$$

Note that by Jensen’s inequality, the entropy and the Kullback–Leibler divergence are non-negative. In what follows, we will often take P and/or Q to be a Gibbs measure $\mu_{\beta,N}$ for some N . In that case, we set $\Omega = \partial\mathbb{T}_n$, where n is the largest number such that both P and Q are defined on $\partial\mathbb{T}_n$.

The following lemma is folklore. For completeness, we provide a proof in Section 6.

- Lemma 1.4.** 1. If $\beta \in [0, \beta_c)$, then $H(\mu_{\beta,N})/N$ converges in probability to a positive constant as $N \rightarrow \infty$.
 2. If $\beta > \beta_c$ and if $\beta_c \in \mathcal{D}(\varphi)^\circ$, then $H(\mu_{\beta,N}) = O(1)$ in probability, as $N \rightarrow \infty$. In other words, the sequence of random variables $(H(\mu_{\beta,N}))_{N \geq 1}$ is tight.

We are interested in the *algorithmically efficient approximation* of the Gibbs measure $\mu_{\beta,N}$. The notion of approximation we will use is the following.

Definition 1.5. Let $\beta \geq 0$. We say that a sequence of random probability measures $(\tilde{\mu}_{\beta,N})_{N \geq 1}$ approximates the Gibbs measure $\mu_{\beta,N}$ if

$$\frac{\mathbf{d}(\tilde{\mu}_{\beta,N} \parallel \mu_{\beta,N})}{H(\mu_{\beta,N})} \rightarrow 0 \text{ in probability as } N \rightarrow \infty. \tag{1.7}$$

Remark 1.6. More generally, assume that we are given two sequences of probability measures $(P_N)_{N \geq 1}$ and $(Q_N)_{N \geq 1}$ satisfying $H(Q_N)/N \rightarrow C \in (0, \infty)$ and

$$\frac{1}{N} \mathbf{d}(P_N \parallel Q_N) \rightarrow 0 \text{ as } N \rightarrow \infty. \tag{1.8}$$

This has been called *measure equivalence* or *equivalence in the sense of specific relative entropy* in the physics literature [31]. Mathematically, Equation (1.8) implies the following: if $(A_N)_{N \geq 1}$ is a sequence of sets such that $Q_N(A_N)$ converges to 0 exponentially fast as $N \rightarrow \infty$, then we also have $P_N(A_N) \rightarrow 0$. Indeed, this is an easy consequence of Birgé’s inequality (see e.g. Theorem 4.20 in [13]).

Main results We now state the main theorem of this paper.

Theorem 1.7 (Approximation bounds). Let $N \in \mathbb{N}$, $M \in \llbracket 1, N \rrbracket$, and $\beta \in [0, \beta_c)$. Then for all $p \geq 1$, there exists a constant $C_1(p) > 0$ such that

$$\|\mathbf{d}(\mu_{\beta,M,N} \parallel \mu_{\beta,N})\|_p \leq C_1(p) \cdot \left\lfloor \frac{N}{M} \right\rfloor. \tag{1.9}$$

Moreover, for all $p \geq 1$, there exists a constant $C_1(p) > 0$ such that

$$\|\mathbf{d}(\mu_{\beta,M,N} \parallel \mu_{\beta,N}) - \mathbb{E}[\mathbf{d}(\mu_{\beta,M,N} \parallel \mu_{\beta,N})]\|_p \leq C_1(p). \tag{1.10}$$

By Theorem 1.7, we can derive the following corollary which states the existence of an algorithm that approximates the Gibbs measure $\mu_{\beta,N}$ efficiently.

Corollary 1.8 (Complexity upper bound). *If $\beta \in [0, \beta_c)$, then there exists a polynomial-time algorithm such that for every $p > 1$, denoting by $\tilde{\mu}_{\beta,N}$ the law of its output,*

$$\frac{1}{N} \|\mathbf{d}(\tilde{\mu}_{\beta,N} \parallel \mu_{\beta,N})\|_p \rightarrow 0 \tag{1.11}$$

as $N \rightarrow \infty$. In particular, $\tilde{\mu}_{\beta,N}$ approximates the Gibbs measure $\mu_{\beta,N}$ in the sense of Definition 1.5.

Proof. Let $M = M(N)$ be a sequence that goes to infinity as $N \rightarrow \infty$, and set $\tilde{\mu}_{\beta,N} = \mu_{\beta,M,N}$. Equation (1.11) then follows from (1.9) in Theorem 1.7. Assuming moreover that $M = O(\log N)$, Remark 1.3 implies that $\tilde{\mu}_{\beta,N}$ can be computed by a polynomial-time algorithm. The second statement follows from the first part of Lemma 1.4. \square

Finally, we also provide a hardness result, assuming a mild regularity condition.

Theorem 1.9 (Complexity lower bound). *Assume $\beta_c \in \mathcal{D}(\varphi)^\circ$ (in particular, $\beta_c < \infty$). Let $\beta > \beta_c$. Let $v = (v(k))_{k \geq 0}$ be an algorithm which outputs a vertex of law $\tilde{\mu}_N$ such that $\tilde{\mu}_N$ approximates the Gibbs measure $\mu_{\beta,N}$ in the sense of Definition 1.5. Let τ be the running time of the algorithm. Then for every $\delta > 0$, there exists $z > 0$, such that for large enough N ,*

$$\mathbb{P} \left(\tau \geq e^{z\sqrt{N}} \right) \geq 1 - \delta.$$

1.2 Related work

An early study on searching algorithms on the branching random walk can be found in Karp and Pearl [22]. They considered the binary branching random walk with Bernoulli increments $\text{Ber}(p)$, and they showed that for $p > 1/2$ and $p = 1/2$, they gave an algorithm that can find an exact maximal vertex in linear and quadratic expected time, respectively. While for $p < 1/2$, it is possible to find an approximate maximal vertex in linear time with high probability using a depth-first search on a renormalized tree. Aldous [5] gave a different algorithm and, among other things, extended the result of Karp and Pearl [22] to general increment distributions. A hardness result was obtained by Pemantle [28]. Among other things, he showed for the binary branching random walk with Bernoulli increments with mean $p < 1/2$ that any search algorithm which finds a vertex within a $(1 - \varepsilon)$ factor of the maximum with high probability needs at least $N \times \exp(\Theta(\varepsilon^{-1/2}))$ with high probability.¹

As mentioned above, Addario-Berry and Maillard [1] considered this optimization problem for the continuous random energy model (CREM), which is a binary time-inhomogeneous branching random walk with Gaussian displacements, proving the existence of an threshold x_* such that the following holds: a) for every $x < x_*$, there exists a polynomial-time algorithm that finds a vertex with $X_v \geq xN$ with high probability, b) for every $x > x_*$, every such algorithm has a running time which is at least exponential in N with high probability.

The CREM, introduced by Bovier and Kurkova [14] based on previous work by Derrida and Spohn [16], is a toy model of a disordered system in statistical physics, i.e. a model where the Hamiltonian – the function that assigns energies to the states of the system – is itself random. These systems have recently seen a lot of interest in the mathematical literature with regards to efficient algorithms for finding low-energy states. A key quantity of importance in these models is the so-called *overlap* between two states, a measure of their correlation. In the case of the CREM, it is equal to the depth of the most recent common ancestor of two vertices, divided by N . Then, for a given $\beta > 0$,

¹To be precise, this explicit bound relied on a conjecture on branching random walk killed at a linear space-time barrier (Conjecture 1 in [28]), which was subsequently proven to be true [9, 19].

the *overlap distribution* is the limiting law (as $N \rightarrow \infty$) of the overlap of two vertices sampled independently according to the Gibbs measure with inverse temperature β . A picture that has emerged is that the existence of a gap in the support of the overlap distribution, the so-called “overlap gap property”, is an obstruction to the existence of efficient algorithms finding approximate minimizers of the Hamiltonian. This has been rigorously proven for a certain class of algorithms by Gamarnik and Jagannath [18] in the case of the Ising p -spin model, with $p \geq 4$. On the other hand, Montanari [26] showed that for $p = 2$, the Sherrington–Kirkpatrick model, there exists a quadratic time algorithm that can find near optimal state with high probability, assuming a widely believed conjecture that this model does not exhibit an overlap gap. A similar result for spherical spin glass models (for which the overlap distribution is explicitly known) has been obtained by Subag [30].

The question of efficient *sampling* of the Gibbs measure of a disordered system seems to have been considered mostly under the angle of Glauber or Langevin dynamics. See e.g. [7, 20] for the spherical spin glass model. We restrict ourselves here to the case of the Sherrington–Kirkpatrick model. For this model, it has been recently obtained that fast mixing occurs for $\beta < 1/4$ [8, 17]. However, very recently, El Alaoui, Montanari and Sellke [3] have provided another algorithm which yields fast mixing for $\beta < 1/2$, and they conjecture that this in fact holds for all $\beta < 1$. They also provide a hardness result for $\beta > 1$ for a certain class of algorithms. Their algorithm for the $\beta < 1/2$ phase belongs to the class of *approximate message passing* algorithms, which is also the case for Montanari’s algorithm for the optimization problem [26]. This illustrates the fact that Glauber or Langevin dynamics may in general not be optimal sampling algorithms, and that algorithms which exploit the underlying tree structure of the model may be efficient in a wider range of the parameters. For a discussion of this question in the context of statistical inference problems, see e.g. [6]. Altogether, this motivates the study of tree-based models as a toy problem, such as the one from the present article.

Outline The paper is organized as follows. In Section 2, we prove that the Kullback–Leibler divergence can be decomposed into a weighted sum of the Kullback–Leibler divergences on subtrees. In Section 3, we give L^p bounds of the logarithm of Biggins’ martingales and L^p bounds of the Kullback–Leibler divergences between two Gibbs measures. Theorem 1.7 is proven in Section 4 and Theorem 1.9 in Section 5. Section 6 provides the proof of Lemma 1.4. Finally, we state in Section 7 some open questions that might interest the readers.

2 Decomposition of the Kullback–Leibler divergence

The main goal of this section is to prove Theorem 2.2. Before proving the theorem, we need the following lemma, which states that the weight of $u_1 u_2$ with respect to the Gibbs measure $\mu_{\beta,m}$ can be decomposed into the product of the weights of u_1 and u_2 with respect to another two Gibbs measures.

Lemma 2.1. *For any $|u_1| = m_1$ and $|u_2| = m_2$, we have the decomposition*

$$\mu_{\beta,m}(u_1 u_2) = \mu_{\beta,m}(u_1) \cdot \mu_{\beta,m-m_1}^{u_1}(u_2)$$

for all $m \geq m_1 + m_2$.

Proof. By computation, we have

$$\mu_{\beta,m}(u_1 u_2) = \frac{e^{\beta X_{u_1 u_2} - \varphi(\beta)m} \cdot W_{\beta,m-m_1-m_2}^{u_1 u_2}}{W_{\beta,m}}$$

$$\begin{aligned}
 &= \frac{e^{\beta X_{u_1} - \varphi(\beta)m_1} \cdot W_{\beta, m-m_1}^{u_1}}{W_{\beta, m}} \cdot \frac{e^{\beta X_{u_2}^{u_1} - \varphi(\beta)(m-m_1)} \cdot W_{\beta, m-m_1-m_2}^{u_1 u_2}}{W_{\beta, m-m_1}^{u_1}} \\
 &= \mu_{\beta, m}(u_1) \cdot \mu_{\beta, m-m_1}^{u_1}(u_2),
 \end{aligned}$$

where the last equality is by (1.2) and (1.3). □

Now we can decompose the Kullback–Leibler divergence as follows.

Theorem 2.2. *For any two M and N integers such that $M \leq N$, we have*

$$\mathbf{d}(\mu_{\beta, M, N} \parallel \mu_{\beta, N}) = \sum_{K=0}^{\lfloor \frac{N}{M} \rfloor - 1} \sum_{|u|=KM} \mu_{\beta, M, KM}(u) \cdot \mathbf{d}(\mu_{\beta, M}^u \parallel \mu_{\beta, N-KM}^u). \quad (2.1)$$

Proof. Denote $N' = \lfloor \frac{N}{M} \rfloor \cdot M$ for simplicity. By (1.4) and Lemma 2.1, for all $|u_1| = N'$ and $|u_2| = N - N'$, we have

$$\log \left(\frac{\mu_{\beta, M, N}(u_1 u_2)}{\mu_{\beta, N}(u_1 u_2)} \right) = \log \left(\frac{\mu_{\beta, M, N'}(u_1) \cdot \mu_{\beta, N-N'}^{u_1}(u_2)}{\mu_{\beta, N}(u_1) \cdot \mu_{\beta, N-N'}^{u_1}(u_2)} \right) = \log \left(\frac{\mu_{\beta, M, N'}(u_1)}{\mu_{\beta, M, N}(u_1)} \right). \quad (2.2)$$

Thus, the Kullback–Leibler divergence can be rewritten as

$$\begin{aligned}
 &\mathbf{d}(\mu_{\beta, M, N} \parallel \mu_{\beta, N}) \\
 &= \sum_{|u|=N} \mu_{\beta, M, N}(u) \cdot \log \left(\frac{\mu_{\beta, M, N}(u)}{\mu_{\beta, N}(u)} \right) \\
 &= \sum_{|u_1|=N'} \sum_{|u_2|=N-N'} \mu_{\beta, M, N}(u_1 u_2) \cdot \log \left(\frac{\mu_{\beta, M, N}(u_1 u_2)}{\mu_{\beta, N}(u_1 u_2)} \right) \\
 &= \sum_{|u_1|=N'} \sum_{|u_2|=N-N'} \mu_{\beta, M, N'}(u_1) \cdot \mu_{\beta, N-N'}^{u_1}(u_2) \cdot \log \left(\frac{\mu_{\beta, M, N'}(u_1)}{\mu_{\beta, M, N}(u_1)} \right) \quad (\text{by (1.4) and (2.2)}) \\
 &= \sum_{|u_1|=N'} \mu_{\beta, M, N'}(u_1) \cdot \log \left(\frac{\mu_{\beta, M, N'}(u_1)}{\mu_{\beta, M, N}(u_1)} \right) \\
 &= \mathbf{d}(\mu_{\beta, M, N'} \parallel \mu_{\beta, N}). \quad (2.3)
 \end{aligned}$$

Next, by (1.4) and Lemma 2.1, for all $|u_1| = KM$, $|u_2| = M$ and $0 \leq K \leq \lfloor \frac{N}{M} \rfloor - 1$, we have

$$\begin{aligned}
 \log \left(\frac{\mu_{\beta, M, (K+1)M}(u_1 u_2)}{\mu_{\beta, N}(u_1 u_2)} \right) &= \log \left(\frac{\mu_{\beta, M, KM}(u_1) \cdot \mu_{\beta, M}^{u_1}(u_2)}{\mu_{\beta, N}(u_1) \cdot \mu_{\beta, N-KM}^{u_1}(u_2)} \right) \\
 &= \log \left(\frac{\mu_{\beta, M, KM}(u_1)}{\mu_{\beta, N}(u_1)} \right) + \log \left(\frac{\mu_{\beta, M}^{u_1}(u_2)}{\mu_{\beta, N-KM}^{u_1}(u_2)} \right). \quad (2.4)
 \end{aligned}$$

Thus,

$$\begin{aligned}
 &\mathbf{d}(\mu_{\beta, M, (K+1)M} \parallel \mu_{\beta, N}) \\
 &= \sum_{|u|=(K+1)M} \mu_{\beta, M, (K+1)M}(u) \cdot \log \left(\frac{\mu_{\beta, M, (K+1)M}(u)}{\mu_{\beta, N}(u)} \right) \\
 &= \sum_{|u_1|=KM} \sum_{|u_2|=M} \mu_{\beta, M, KM}(u_1)
 \end{aligned}$$

$$\begin{aligned} & \cdot \mu_{\beta,M}^{u_1}(u_2) \cdot \left[\log \left(\frac{\mu_{\beta,M,KM}(u_1)}{\mu_{\beta,N}(u_1)} \right) + \log \left(\frac{\mu_{\beta,M}^{u_1}(u_2)}{\mu_{\beta,N-KM}^{u_1}(u_2)} \right) \right] && \text{(by (1.4) and (2.4))} \\ = & \mathbf{d}(\mu_{\beta,M,KM} \parallel \mu_{\beta,N}) + \sum_{|u_1|=KM} \mu_{\beta,M,KM}(u_1) \cdot \mathbf{d}(\mu_{\beta,M}^{u_1} \parallel \mu_{\beta,N-KM}^{u_1}). && (2.5) \end{aligned}$$

Finally, by (2.3), (2.5) and the fact that $\mathbf{d}(\mu_{\beta,M,0} \parallel \mu_{\beta,N}) = 0$, we derive (2.1). \square

Next, we show that the Kullback–Leibler divergence between two Gibbs measures can be written in terms of the logarithms of the partition functions.

Proposition 2.3. *For any two M and N integers such that $M \leq N$, we have*

$$\mathbf{d}(\mu_{\beta,M} \parallel \mu_{\beta,N}) = \log W_{\beta,N} - \log W_{\beta,M} - \sum_{|u|=M} \mu_{\beta,M}(u) \cdot \log W_{\beta,N-M}^u. \quad (2.6)$$

Proof. For any $|u| = M$, we have

$$\frac{\mu_{\beta,M}(u)}{\mu_{\beta,N}(u)} = \frac{e^{\beta X_u - \varphi(\beta)M} \cdot \frac{1}{W_{\beta,M}}}{e^{\beta X_u - \varphi(\beta)M} \cdot \frac{W_{\beta,N-M}^u}{W_{\beta,N}}} = \frac{W_{\beta,N}}{W_{\beta,M} \cdot W_{\beta,N-M}^u}.$$

Thus,

$$\begin{aligned} \mathbf{d}(\mu_{\beta,M} \parallel \mu_{\beta,N}) &= \sum_{|u|=M} \mu_{\beta,M}(u) \cdot \log \left(\frac{\mu_{\beta,M}(u)}{\mu_{\beta,N}(u)} \right) \\ &= \sum_{|u|=M} \mu_{\beta,M}(u) \cdot (\log W_{\beta,N} - \log W_{\beta,M} - \log W_{\beta,N-M}^u) \\ &= \log W_{\beta,N} - \log W_{\beta,M} - \sum_{|u|=M} \mu_{\beta,M}(u) \cdot \log W_{\beta,N-M}^u. \end{aligned}$$

This completes the proof. \square

3 Some L^p bounds

We first show that whenever $\beta \in [0, \beta_c)$, $(\log W_{\beta,n})_{n \geq 0}$ is bounded in L^p for all $p > 1$.

Lemma 3.1. *Let $\beta \in [0, \beta_c)$. Then $(\log W_{\beta,n})_{n \geq 0}$ is a supermartingale such that*

$$\sup_{n \geq 0} \|\log W_{\beta,n}\|_p < \infty, \quad p \geq 1.$$

Proof. The supermartingale property of $(\log W_{\beta,n})_{n \geq 0}$ follows from the fact that $(W_{\beta,n})_{n \geq 0}$ is a martingale and $x \mapsto \log x$ is a concave function.

Now let $p \geq 1$. Since $\beta < \beta_c$, by Fact 1.2, we derive that

$$\sup_{n \geq 0} \|W_{\beta,n}\|_1 < \infty$$

Furthermore, using Assumption 1.1 and Liu [23, Theorem 2.4], we have for some $s > 0$,

$$\sup_{n \geq 0} \|W_{\beta,n}^{-s}\|_1 < \infty.$$

Then by the fact that

$$|\log x|^p \leq C \left(|x| + |x|^{-s} \right)$$

for some constant $C > 0$, we have

$$\sup_{n \geq 0} \|\log W_{\beta,n}\|_p^p \leq C \cdot \sup_{n \geq 0} \left(\|W_{\beta,n}\|_1 + \|W_{\beta,n}^{-s}\|_1 \right) < \infty.$$

This proves the lemma. \square

Lemma 3.1 implies the following proposition about the boundedness of the Kullback-Leibler divergence between two Gibbs measures.

Proposition 3.2. *For any $p > 1$, for any two integers M and N such that $M \leq N$, there exists a constant $C(p) > 0$ such that*

$$\|\mathbf{d}(\mu_{\beta,M} \parallel \mu_{\beta,N})\|_p \leq C(p).$$

Proof. By Minkowski’s inequality and Proposition 2.3, we have

$$\|\mathbf{d}(\mu_{\beta,M} \parallel \mu_{\beta,N})\|_p \leq \|\log W_{\beta,N}\|_p + \|\log W_{\beta,M}\|_p + \left\| \sum_{|u|=M} \mu_{\beta,M}(u) \cdot \log W_{\beta,N-M}^u \right\|_p.$$

Since the first and the second term above are bounded by Lemma 3.1, it suffices to prove that the third term above is bounded. By the branching property and Jensen’s inequality, we have

$$\mathbb{E} \left[\left\| \sum_{|u|=M} \mu_{\beta,M}(u) \cdot \log W_{\beta,M}^u \right\|_p^p \middle| \mathcal{F}_M \right] \leq \mathbb{E} [|\log W_{\beta,N-M}|^p]. \tag{3.1}$$

Therefore, by (3.1),

$$\left\| \sum_{|u|=M} \mu_{\beta,M}(u) \cdot \log W_{\beta,N-M}^u \right\|_p \leq \|\log W_{\beta,N-M}\|_p. \tag{3.2}$$

Finally, we conclude by (3.2) and Lemma 3.1 that

$$\|\mathbf{d}(\mu_{\beta,M} \parallel \mu_{\beta,N})\|_p \leq 3 \cdot \sup_{n \geq 0} \|\log W_{\beta,n}\|_p < \infty,$$

and the proof is completed. □

4 Proof of Theorem 1.7

In this section, we prove (1.9) and (1.10) of Theorem 1.7. The proof of (1.9) relies essentially on the decomposition theorem of the Kullback–Leibler divergence (Theorem 2.2) and Proposition 3.2. The proof of (1.10) needs more precise moment estimates.

4.1 Proof of (1.9)

Let $p \geq 1$. By Theorem 2.2 and Minkowski’s inequality, we have

$$\begin{aligned} \|\mathbf{d}(\mu_{\beta,M,N} \parallel \mu_{\beta,N})\|_p &= \left\| \sum_{K=0}^{\lfloor \frac{N}{M} \rfloor - 1} \sum_{|u|=KM} \mu_{\beta,M,KM}(u) \cdot \mathbf{d}(\mu_{\beta,M}^u \parallel \mu_{\beta,N-KM}^u) \right\|_p \\ &\leq \sum_{K=0}^{\lfloor \frac{N}{M} \rfloor - 1} \left\| \sum_{|u|=KM} \mu_{\beta,M,KM}(u) \cdot \mathbf{d}(\mu_{\beta,M}^u \parallel \mu_{\beta,N-KM}^u) \right\|_p. \end{aligned} \tag{4.1}$$

Let $K \leq \lfloor N/M \rfloor - 1$. Applying Jensen’s inequality to $\mu_{\beta,M,KM}$, we have

$$\mathbb{E} \left[\left\| \sum_{|u|=KM} \mu_{\beta,M,KM}(u) \cdot \mathbf{d}(\mu_{\beta,M}^u \parallel \mu_{\beta,N-KM}^u) \right\|_p^p \right]$$

$$\leq \mathbb{E} \left[\sum_{|u|=KM} \mu_{\beta, M, KM}(u) \cdot |\mathbf{d}(\mu_{\beta, M}^u \parallel \mu_{\beta, N-KM}^u)|^p \right]. \tag{4.2}$$

Then by the law of iterated expectation and the branching property, (4.2) is equal to

$$\begin{aligned} & \mathbb{E} \left[\sum_{|u|=KM} \mu_{\beta, M, KM}(u) \cdot \mathbb{E} [|\mathbf{d}(\mu_{\beta, M}^u \parallel \mu_{\beta, N-KM}^u)|^p \mid \mathcal{F}_{KM}] \right] \\ &= \mathbb{E} \left[\sum_{|u|=KM} \mu_{\beta, M, KM}(u) \right] \cdot \mathbb{E} [|\mathbf{d}(\mu_{\beta, M} \parallel \mu_{\beta, N-KM})|^p] = \mathbb{E} [|\mathbf{d}(\mu_{\beta, M} \parallel \mu_{\beta, N-KM})|^p]. \end{aligned} \tag{4.3}$$

Combining (4.1), (4.3) and Proposition 3.2, we conclude that

$$\|\mathbf{d}(\mu_{\beta, M, N} \parallel \mu_{\beta, N})\|_p \leq \left\lfloor \frac{N}{M} \right\rfloor \cdot C(p).$$

4.2 Proof of (1.10)

In this section, we prove (1.10) which gives a tighter control on the Kullback–Leibler divergence between $\mu_{\beta, M, N}$ and the Gibbs measure $\mu_{\beta, N}$. We start with the following simple lemma.

Lemma 4.1. *Let $\beta \in [0, \beta_c)$. For all $M \in \mathbb{N}$, there exists $r \in (0, 1)$ independent of M such that,*

$$\mathbb{E} \left[\sum_{|u|=M} \mu_{\beta, M}(u)^2 \right] \leq r \tag{4.4}$$

Moreover, for all $K \in \mathbb{N}$,

$$\mathbb{E} \left[\sum_{|u|=KM} \mu_{\beta, M, KM}(u)^2 \right] \leq r^K. \tag{4.5}$$

Proof of Lemma 4.1. For $x \in (0, 1)$, $x^2 < x$. Thus, by the fact that $\mu_{\beta, M}(u) \in (0, 1)$ for all $|u| = M$, we derive that

$$\sum_{|u|=M} \mu_{\beta, M}(u)^2 < \sum_{|u|=M} \mu_{\beta, M}(u) = 1,$$

for every $M \in \mathbb{N}$. This shows that (4.4) holds for every fixed $M \in \mathbb{N}$ and with $r < 1$ possibly depending on M . Uniformity in M follows as soon as we show that

$$\limsup_{M \rightarrow \infty} \mathbb{E} \left[\sum_{|u|=M} \mu_{\beta, M}(u)^2 \right] < 1.$$

To this end, recall that $W_{\beta, M} \rightarrow W_{\beta, \infty}$ almost surely as $M \rightarrow \infty$ and that $\mathbb{E}[W_{\beta, \infty}] = 1$. Hence, there exist $a < 1$ and $M_0 \in \mathbb{N}$ such that

$$\forall M \geq M_0 : \mathbb{P}(W_{\beta, M} < 1/2) \leq a. \tag{4.6}$$

Now fix $p \in (1, 2]$ such that $\varphi(p\beta) < p\varphi(\beta)$, which exists because $\beta < \beta_c$. Then decompose:

$$\mathbb{E} \left[\sum_{|u|=M} \mu_{\beta, M}(u)^2 \right] \leq \mathbb{E} \left[\sum_{|u|=M} \mu_{\beta, M}(u)^p \right]$$

$$\begin{aligned} &\leq \mathbb{E} \left[\left(\sum_{|u|=M} \mu_{\beta,M}(u)^p \right) \mathbf{1}_{W_{\beta,M} < 1/2} \right] \\ &\quad + \mathbb{E} \left[\left(\sum_{|u|=M} \mu_{\beta,M}(u)^p \right) \mathbf{1}_{W_{\beta,M} \geq 1/2} \right] \\ &\leq \mathbb{P}(W_{\beta,M} < 1/2) + 2^p e^{(\varphi(p\beta) - p\varphi(\beta))M}, \end{aligned}$$

using the definition of $\mu_{\beta,M}(u)$ for the last inequality. This shows that

$$\limsup_{M \rightarrow \infty} \mathbb{E} \left[\sum_{|u|=M} \mu_{\beta,M}(u)^2 \right] \leq a,$$

which concludes the proof of (4.4).

We now show (4.5) by induction. The case $K = 1$ follows directly from (1.4) and (4.4). For $K > 1$, by (1.4) and branching property, we obtain that

$$\mathbb{E} \left[\sum_{|u|=KM} \mu_{\beta,M,KM}(u)^2 \middle| \mathcal{F}_{(K-1)M} \right] \tag{4.7}$$

$$\begin{aligned} &= \mathbb{E} \left[\sum_{|u_1|=(K-1)M} \sum_{|u_2|=M} \mu_{\beta,M,(K-1)M}(u_1)^2 \cdot \mu_{\beta,M}^{u_1}(u_2)^2 \middle| \mathcal{F}_{(K-1)M} \right] \\ &= \sum_{|u_1|=(K-1)M} \mu_{\beta,M,(K-1)M}(u_1)^2 \cdot \mathbb{E} \left[\sum_{|u_2|=M} \mu_{\beta,M}(u_2)^2 \right]. \end{aligned} \tag{4.8}$$

Taking expectations, we derive that

$$\mathbb{E} \left[\sum_{|u|=KM} \mu_{\beta,M,KM}(u)^2 \right] = \mathbb{E} \left[\sum_{|u|=(K-1)M} \mu_{\beta,M,(K-1)M}(u)^2 \right] \cdot \mathbb{E} \left[\sum_{|u|=M} \mu_{\beta,M}(u)^2 \right]. \tag{4.9}$$

The equation (4.5) then follows from the induction hypothesis. \square

We now proceed with the proof of (1.10). Without loss of generality, using the fact that $\|\cdot\|_p \leq \|\cdot\|_2$ for every $p \in [1, 2]$, we assume that $p \geq 2$. By Theorem 2.2 and Minkowski's inequality, we have

$$\begin{aligned} &\|\mathbf{d}(\mu_{\beta,M,N} \parallel \mu_{\beta,N}) - \mathbb{E}[\mathbf{d}(\mu_{\beta,M,N} \parallel \mu_{\beta,N})]\|_p \\ &\leq \sum_{K=0}^{\lfloor \frac{N}{M} \rfloor - 1} \left\| \sum_{|u|=KM} \mu_{\beta,M,KM}(u) \cdot (\mathbf{d}(\mu_{\beta,M}^u \parallel \mu_{\beta,N-KM}^u) - \mathbb{E}[\mathbf{d}(\mu_{\beta,M}^u \parallel \mu_{\beta,N-KM}^u)]) \right\|_p. \end{aligned} \tag{4.10}$$

We now introduce some notation. For all $0 \leq K \leq \lfloor \frac{N}{M} \rfloor - 1$, denote

$$d_K^u = \mathbf{d}(\mu_{\beta,M}^u \parallel \mu_{\beta,N-KM}^u) \quad \text{and} \quad d_K = \mathbf{d}(\mu_{\beta,M} \parallel \mu_{\beta,N-KM})$$

and

$$Z_u = \mu_{\beta,M,KM}(u) \cdot (d_K^u - \mathbb{E}[d_K^u]).$$

We claim that for all $p \geq 1$, the sequence

$$a_K = \left\| \sum_{|u|=KM} Z_u \right\|_p \tag{4.11}$$

is summable, with a bound independent of M . This will imply that the right-hand side of (4.10) is bounded by the same quantity, which completes proof. To prove this, first observe that $(Z_u)_{|u|=KM}$ is a sequence of iid random variables having zero mean and finite p -th moments, for any $p \geq 1$, with respect to $\mathbb{E}[\cdot | \mathcal{F}_{KM}]$, by Proposition 3.2 and the branching property. Denote by C_1, C_2, \dots some constants possibly depending on p (and β and the law of \mathbf{Y}). By Rosenthal’s inequality [29, Theorem 3], we have

$$\begin{aligned} & \mathbb{E} \left[\left| \sum_{|u|=KM} Z_u \right|^p \middle| \mathcal{F}_{KM} \right] \\ & \leq C_1 \cdot \left\{ \left(\sum_{|u|=KM} \mathbb{E} [|Z_u|^p | \mathcal{F}_{KM}] \right) + \left(\sum_{|u|=KM} \mathbb{E} [|Z_u|^2 | \mathcal{F}_{KM}] \right)^{p/2} \right\} \end{aligned} \tag{4.12}$$

By the branching property and Proposition 3.2, the first term of (4.12) can be bounded by

$$\begin{aligned} \sum_{|u|=KM} \mathbb{E} [|Z_u|^p | \mathcal{F}_{KM}] &= \sum_{|u|=KM} \mu_{\beta, M, KM}(u)^p \cdot \mathbb{E} [|d_K - \mathbb{E} [d_K]|^p] \\ &\leq \sum_{|u|=KM} \mu_{\beta, M, KM}(u)^p \cdot C_2 \\ &\leq \sum_{|u|=KM} \mu_{\beta, M, KM}(u)^2 \cdot C_2, \end{aligned} \tag{4.13}$$

using that $p \geq 2$ in the last line. Taking expectations and applying Lemma 4.1, we get

$$\mathbb{E} \left[\sum_{|u|=KM} \mathbb{E} [|Z_u|^p | \mathcal{F}_{KM}] \right] \leq \mathbb{E} \left[\sum_{|u|=KM} \mu_{\beta, M, KM}(u)^2 \right] \cdot C_2 \leq r^K \cdot C_2, \tag{4.14}$$

with $r < 1$ as in Lemma 4.1.

We now estimate the second term of (4.12). By the branching property and Lemma 3.1,

$$\begin{aligned} \left(\sum_{|u|=KM} \mathbb{E} [|Z_u|^2 | \mathcal{F}_{KM}] \right)^{p/2} &= \left(\sum_{|u|=KM} \mu_{\beta, M, KM}(u)^2 \cdot \mathbb{E} [|d_K - \mathbb{E} [d_K]|^2] \right)^{p/2} \\ &\leq \left(\sum_{|u|=KM} \mu_{\beta, M, KM}(u)^2 \right)^{p/2} \cdot C_3 \\ &\leq \left(\sum_{|u|=KM} \mu_{\beta, M, KM}(u)^2 \right) \cdot C_3. \end{aligned} \tag{4.15}$$

The inequality (4.15) is because

$$\sum_{|u|=KM} \mu_{\beta, M, KM}(u)^2 \in [0, 1]$$

and $x^{p/2} \leq x$, for all $x \in [0, 1]$ and $p \geq 2$. Taking expectations and applying Lemma 4.1, we get

$$\mathbb{E} \left[\left(\sum_{|u|=KM} \mathbb{E} \left[|Z_u|^2 \mid \mathcal{F}_{KM} \right] \right)^{p/2} \right] \leq \mathbb{E} \left[\sum_{|u|=KM} \mu_{\beta, M, KM}(u)^2 \right] \cdot C_3 \leq r^K \cdot C_3. \quad (4.16)$$

Combining (4.12), (4.14) and (4.16), we conclude that

$$\left\| \sum_{|u|=KM} Z_u \right\|_p \leq C_1(C_2 + C_3) \cdot r^K.$$

This implies that (4.11) is summable and finishes the proof.

5 Proof of hardness result (Theorem 1.9)

In this section, we prove Theorem 1.9. To do this, we recall two results of asymptotic behaviors of the maximal particle of a branching random walk.

Under the assumption of the theorem, it is known in Corollary (3.4) of [10] that

$$\frac{\max_{|u|=N} X_u}{N} \rightarrow m := \varphi'(\beta_c), \quad \text{a.s. as } N \rightarrow \infty. \quad (5.1)$$

Furthermore, we have the following tail estimate, which easily follows from a union bound together with Chernoff's bound (see e.g. the proof of Theorem 2 in [32]): there exists a constant $c > 0$ such that

$$\forall x \geq 0 : \mathbb{P} \left(\max_{|u|=N} X_u \geq mN + x \right) \leq e^{-cx}. \quad (5.2)$$

The key to Theorem 1.9 is the following observation:

Lemma 5.1. *Let $\beta > \beta_c$ and assume $\beta_c \in \mathcal{D}(\varphi)^\circ$. Let u be a particle sampled according to the Gibbs measure $\mu_{\beta, N}$ and let w be its ancestor at generation $\lfloor N/2 \rfloor$. Then there exists a positive random variable Z with continuous distribution function such that*

$$\frac{X_u - X_w - mN/2}{\sqrt{N}} \rightarrow Z, \quad \text{in law as } N \rightarrow \infty.$$

Proof. This is a consequence of a result by Chen, Madaule and Mallein [15]. These authors show the following fact: if u is sampled according to the Gibbs measure $\mu_{\beta, N}$, and $X_u(t)$ denotes the position of its ancestor at generation $\lfloor tN \rfloor$, and if we define

$$Z_t^N := \frac{mtN - X_u(t)}{\sqrt{N}}, \quad t \in [0, 1],$$

then $(Z_t^N)_{t \in [0, 1]}$ converges in law (w.r.t. Skorokhod's topology) to a multiple of a Brownian excursion as $N \rightarrow \infty$. Note that the assumptions in their article are implied by our hypothesis that $\beta_c \in \mathcal{D}(\varphi)^\circ$ and the fact that in our branching random walk, the number of offspring of a particle is deterministic. Now, we also have that $X_u - mN = O(\log N)$ in probability (see e.g. [2]), so that

$$\frac{X_u - mN}{\sqrt{N}} \xrightarrow{\text{law}} 0, \quad \text{as } N \rightarrow \infty.$$

Together, both results imply the lemma. □

Proof of Theorem 1.9. Assume that we are given an algorithm $(v(k))_{k \geq 1}$ that samples a vertex according to a random probability measure $\tilde{\mu}_{\beta,N}$ approximating the Gibbs measure $\mu_{\beta,N}$. By Lemma 1.4, it follows that $d(\tilde{\mu}_{\beta,N} || \mu_{\beta,N}) \rightarrow 0$ in probability as $N \rightarrow \infty$. Hence, by Pinsker's inequality (see e.g. Theorem 4.19 in [13]), the total variation distance between $\tilde{\mu}_{\beta,N}$ and $\mu_{\beta,N}$ goes to 0 as well in probability, as $N \rightarrow \infty$. It follows that Lemma 5.1 holds as well for u sampled according to $\tilde{\mu}_{\beta,N}$.

Let $\delta > 0$. For $z > 0$, call a vertex $w \in \partial\mathbb{T}_{\lfloor N/2 \rfloor}$ *z-exceptional* if it has a descendant $u \in \partial\mathbb{T}_N$ such that $X_u - X_w - mN/2 > z\sqrt{N}$. By the preceding paragraph, there exists $z > 0$ such that for large enough N , the algorithm finds a vertex u whose ancestor w at generation $\lfloor N/2 \rfloor$ is *z-exceptional* with probability at least $1 - \delta$. Hence, it is enough to show that any algorithm which solves the simpler problem of finding a *z-exceptional* vertex at generation $\lfloor N/2 \rfloor$ has a running time at least $e^{z'\sqrt{N}}$ with probability $1 - \delta$, for some $z' > 0$. We will now show that the statement of the theorem holds even for this simpler problem. For this, we use an argument similar to the one in Section 3 of [1]. We first present the argument in an informal way.

Denote by E_w the event that a given vertex $w \in \partial\mathbb{T}_{\lfloor N/2 \rfloor}$ is *z-exceptional*. Note that this event only depends on the displacements of the descendants of w . Hence, the events $(E_w)_{w \in \partial\mathbb{T}_{\lfloor N/2 \rfloor}}$ are independent by the branching property. Furthermore, by (5.2), for each $w \in \partial\mathbb{T}_{\lfloor N/2 \rfloor}$, we have $\mathbb{P}(E_w) \leq e^{-cz\sqrt{N/2}}$ for some $c > 0$. Finally, in order to determine whether a vertex w is *z-exceptional*, the algorithm has to explore at least one vertex in the subtree of the vertex w . Hence, the running time of the algorithm is bounded from below by the number of vertices w that have to be probed in order to find a *z-exceptional* vertex. But this quantity follows the geometric distribution with success probability $\mathbb{P}(E_w) \leq e^{-cz\sqrt{N/2}}$. Altogether, for any $z' < cz/\sqrt{2}$ and for N sufficiently large, this shows that the running time τ of any algorithm solving the simpler problem is at least $e^{z'\sqrt{N}}$ with probability $1 - \delta$. The statement readily follows.

We now make this argument formal. Recall that, by definition, an algorithm is a stochastic process $(v(n))_{n \geq 0}$ previsible with respect to the filtration $\tilde{\mathcal{F}}$, defined by

$$\tilde{\mathcal{F}}_k = \sigma(v(1), \dots, v(k); X_{v(1)}, \dots, X_{v(k)}; U_1, \dots, U_{k+1})$$

where $(U_k)_{k \geq 1}$ is a sequence of iid uniform random variables on $[0, 1]$, independent of the branching random walk \mathbf{X} . We now define a larger filtration \mathcal{G} . For this, define for any $v \in \mathbb{T}_N$ the following set of vertices:

$$\mathcal{V}_v = \begin{cases} w \in \mathbb{T}_N : |v \wedge w| \geq \lfloor N/2 \rfloor, & \text{if } |v| \geq \lfloor N/2 \rfloor \\ v, & \text{otherwise.} \end{cases}$$

Note that $v \in \mathcal{V}_v$ for every $v \in \mathbb{T}_N$. We then set

$$\mathcal{G}_k = \sigma(v(1), \dots, v(k); (X_w)_{w \in \mathcal{V}_{v(1)}}, \dots, (X_w)_{w \in \mathcal{V}_{v(k)}}; U_1, \dots, U_{k+1}).$$

Note that $\tilde{\mathcal{F}}_k \subset \mathcal{G}_k$ for all $k \geq 0$ — heuristically, \mathcal{G}_k adds to $\tilde{\mathcal{F}}_k$ the information about the values in the branching random walk of all vertices contained in $\mathcal{V}_{v(i)}$, $i = 1, \dots, k$. Note that trivially, the stochastic process $v(n)_{n \geq 0}$ is still previsible with respect to this larger filtration \mathcal{G} .

Now say that \mathcal{V}_v is *z-exceptional* if $|v| \geq \lfloor N/2 \rfloor$ and the ancestor of v at generation $\lfloor N/2 \rfloor$ is *z-exceptional* in the sense defined above — note that this definition does not depend on the choice of v . Define

$$\tau' = \inf\{k \geq 0 : \mathcal{V}_{v(k)} \text{ is } z\text{-exceptional}\},$$

and note that τ' is a stopping time with respect to the filtration \mathcal{G} . Now, by the equality of events

$$\{\tau' = k\} = \{\mathcal{V}_{v(k)} \text{ is } z\text{-exceptional}\} \cap \{\tau' > k - 1\},$$

and since $\{\tau' > k - 1\} \in \mathcal{G}_{k-1}$ and $v(k)$ is \mathcal{G}_{k-1} -measurable, we have

$$\mathbb{P}(\tau' = k | \mathcal{G}_{k-1}) = \sum_{v \in \mathbb{T}_N} \mathbb{P}(\mathcal{V}_v \text{ is } z\text{-exceptional} | \mathcal{G}_{k-1}) \mathbf{1}_{(v(k)=v, \tau' > k-1)}.$$

Now, for any $v \in \mathbb{T}_N$, if $|v| < \lfloor N/2 \rfloor$, or if $v \in \bigcup_{i=0}^{k-1} \mathcal{V}_{v(i)}$, the above probability is zero, because none of $\mathcal{V}_{v(i)}$, $i = 0, \dots, k - 1$ are z -exceptional on the event $\{\tau' > k - 1\}$. On the other hand, if $v \notin \bigcup_{i=0}^{k-1} \mathcal{V}_{v(i)}$, then, by the branching property, the above conditional probability is equal to the unconditioned probability that \mathcal{V}_v is z -exceptional, which is bounded by $e^{-cz\sqrt{N/2}}$ by (5.2). Hence, we get in total that

$$\mathbb{P}(\tau' = k | \tau' > k - 1) \leq e^{-cz\sqrt{N/2}},$$

and τ' is dominated from below by a geometric random variable with success probability $e^{-cz\sqrt{N/2}}$. The proof now continues as above. \square

6 Proof of Lemma 1.4

We first consider the case $\beta \in [0, \beta_c)$. Define

$$D_{\beta,N} = \frac{d}{d\beta} W_{\beta,N} = \sum_{|u|=n} (X_u - \varphi'(\beta)N) e^{\beta X_u - \varphi(\beta)N}.$$

We express the entropy by

$$\begin{aligned} H(\mu_{\beta,N}) &= \frac{1}{W_{\beta,N}} \sum_{|u|=N} (\varphi(\beta)N - \beta X_u) e^{\beta X_u - \varphi(\beta)N} + \log W_{\beta,N} \\ &= (\varphi(\beta) - \beta \varphi'(\beta))N - \beta \frac{D_{\beta,N}}{W_{\beta,N}} + \log W_{\beta,N}. \end{aligned}$$

By the assumption on β , we have $\varphi(\beta) - \beta \varphi'(\beta) > 0$. Furthermore, $W_{\beta,N}$ converges almost surely to a positive random variable as $N \rightarrow \infty$ by Fact 1.2 and $D_{\beta,N}$ converges almost surely as well as $N \rightarrow \infty$, see [12]. The first statement follows.

Now let $\beta > \beta_c$ and assume $\beta_c \in \mathcal{D}(\varphi)^\circ$. Define

$$\widetilde{W}_{\beta,N} = \sum_{|u|=N} e^{\beta(X_u - \varphi'(\beta_c)N - \frac{3}{2\beta_c} \log N)}.$$

We now write

$$\begin{aligned} H(\mu_{\beta,N}) &= \frac{1}{\widetilde{W}_{\beta,N}} \sum_{|u|=N} \left(-\beta \left(X_u - \varphi'(\beta_c)N - \frac{3}{2\beta_c} \log N \right) \right) e^{\beta(X_u - \varphi'(\beta_c)N - \frac{3}{2\beta_c} \log N)} + \log \widetilde{W}_{\beta,N}. \end{aligned}$$

Fix $\beta' \in (\beta_c, \beta)$. Then there exists $C > 0$, such that $x e^{-\beta x} \leq C e^{-\beta' x}$ for every $x \in \mathbb{R}$. We have

$$H(\mu_{\beta,N}) \leq C \frac{\widetilde{W}_{\beta',N}}{\widetilde{W}_{\beta,N}} + \log \widetilde{W}_{\beta,N}.$$

Now, $\widetilde{W}_{\beta,N}$ and $\widetilde{W}_{\beta',N}$ converge almost surely as $N \rightarrow \infty$ to positive random variables, see [25]. The second statement follows.

7 Open questions and further directions

In this section, we state a few questions and further directions related to the CREM and the branching random walks for future study.

1. It was conjectured in [1] that for the CREM, there exists a threshold $\beta_G \geq 0$ such that its Gibbs measure with parameter β can be efficiently approximated if $\beta < \beta_G$ and cannot if $\beta > \beta_G$. A conjectured explicit expression² of β_G appears in Item 1, Section 5 of [1]. Our results confirm the conjecture in the case where the CREM has correlation function $A(x) = x$, with the notion of *approximation* from Definition 1.5. Moreover, our results imply that $\beta_G = \beta_c = \sqrt{2 \log 2}$, the (*static*) *critical inverse temperature*. One can check that β_c equals the expression of β_G from [1]. Ongoing work of the authors is trying to generalize the result to the CREM with a general correlation function.
2. Back to the branching random walk, one might be interested in the near critical regime to understand how the transition happens near β_c . To do this, one can take a sequence $\beta(N) = \beta_c - N^{-\delta}$ for some $\delta > 0$. It might be interesting to study the time complexity of any algorithm approximating the Gibbs measure $\mu_{\beta(N), N}$ in the sense of Definition 1.5. One should expect a phase transition at $\delta = 1/2$, in line with a phase transition for the asymptotics of the partition function obtained by Alberts and Ortgièse [4]. See also the introduction of Pain [27]. This should be related to Pemantle's [28] study of optimization algorithms discussed in Section 1.2.

References

- [1] Louigi Addario-Berry and Pascal Maillard, *The algorithmic hardness threshold for continuous random energy models*, *Mathematical Statistics and Learning* **2** (2020), 77–101. MR4073148
- [2] Elie Aïdékon, *Convergence in law of the minimum of a branching random walk*, *Annals of Probability* **41** (2013), 1362–1426. MR3098680
- [3] Ahmed El Alaoui, Andrea Montanari, and Mark Sellke, *Sampling from the Sherrington-Kirkpatrick Gibbs measure via algorithmic stochastic localization*, arXiv:2203.05093 [cond-mat] (2022).
- [4] Tom Alberts and Marcel Ortgièse, *The near-critical scaling window for directed polymers on disordered trees*, *Electronic Journal of Probability* **18** (2013), 1–24. MR3035747
- [5] David Aldous, *Greedy Search on the Binary Tree with Random Edge-Weights*, *Combinatorics, Probability and Computing* **1** (1992), 281–293. MR1211318
- [6] Fabrizio Antenucci, Silvio Franz, Pierfrancesco Urbani, and Lenka Zdeborová, *Glassy Nature of the Hard Phase in Inference Problems*, *Physical Review X* **9** (2019), 011020.
- [7] Gérard Ben Arous and Aukosh Jagannath, *Spectral gap estimates in mean field spin glasses*, *Communications in Mathematical Physics* **361** (2018), 1–52. MR3825934
- [8] Roland Bauerschmidt and Thierry Bodineau, *Spectral gap critical exponent for Glauber dynamics of hierarchical spin models*, *Communications in Mathematical Physics* **373** (2020), 1167–1206. MR4061408
- [9] Jean Bérard and Jean-Baptiste Gouéré, *Survival Probability of the Branching Random Walk Killed Below a Linear Boundary*, *Electronic Journal of Probability* **16** (2011), 396–418. MR2774095
- [10] John D. Biggins, *Chernoff's Theorem in the Branching Random Walk*, *Journal of Applied Probability* **14** (1977), 630–636. MR0464415
- [11] John D. Biggins, *Martingale Convergence in the Branching Random Walk*, *Journal of Applied Probability* **14** (1977), 25–37. MR0433619

²There is a typo in the expression of β_G in the case of the CREM in Item 1, Section 5 of [1]. The definition of t_G should be replaced by the following one: $t_G = \sup\{t \in [0, 1] : A(s) = \hat{A}(s) \text{ for all } s \leq t\}$.

- [12] John D. Biggins, *Uniform Convergence of Martingales in the Branching Random Walk*, *Annals of Probability* **20** (1992), 137–151. MR1143415
- [13] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart, *Concentration Inequalities: A Nonasymptotic Theory of Independence*, Oxford University Press, February 2013. MR3185193
- [14] Anton Bovier and Irina Kurkova, *Derrida’s Generalized Random Energy models 2: Models with continuous hierarchies*, *Annales de l’Institut Henri Poincaré (B) Probability and Statistics* **40** (2004), 481–495. MR2070335
- [15] Xinxin Chen, Thomas Madaule, and Bastien Mallein, *On the trajectory of an individual chosen according to supercritical Gibbs measure in the branching random walk*, *Stochastic Processes and their Applications* **129** (2019), 3821–3858. MR3997663
- [16] Bernard Derrida and Herbert Spohn, *Polymers on disordered trees, spin glasses, and traveling waves*, *Journal of Statistical Physics* **51** (1988), 817–840. MR0971033
- [17] Ronen Eldan, Frederic Koehler, and Ofer Zeitouni, *A Spectral Condition for Spectral Gap: Fast Mixing in High-Temperature Ising Models*, arXiv:2007.08200 [math-ph] (2021). MR4408509
- [18] David Gamarnik and Aukosh Jagannath, *The overlap gap property and approximate message passing algorithms for p -spin models*, *The Annals of Probability* **49** (2021), 180–205. MR4203336
- [19] Nina Gantert, Yueyun Hu, and Zhan Shi, *Asymptotics for the survival probability in a killed branching random walk*, *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques* **47** (2011), 111–129. MR2779399
- [20] Reza Gheissari and Aukosh Jagannath, *On the spectral gap of spherical spin glass dynamics*, *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques* **55** (2019), 756–776. MR3949952
- [21] Jean-Pierre Kahane and Jacques Peyrière, *Sur certaines martingales de Benoit Mandelbrot*, *Advances in Mathematics* **22** (1976), 131–145. MR0431355
- [22] Richard M. Karp and Judea Pearl, *Searching for an optimal path in a tree with random costs*, *Artificial Intelligence* **21** (1983), 99–116. MR0716388
- [23] Quansheng Liu, *Asymptotic properties and absolute continuity of laws stable by random weighted mean*, *Stochastic Processes and their Applications* **95** (2001), 83–107. MR1847093
- [24] Russell Lyons, *A Simple Path to Biggins’ Martingale Convergence for Branching Random Walk*, *Classical and Modern Branching Processes* (Krishna B. Athreya and Peter Jagers, eds.), *The IMA Volumes in Mathematics and Its Applications*, Springer, New York, NY, 1997, pp. 217–221. MR1601749
- [25] Thomas Madaule, *Convergence in Law for the Branching Random Walk Seen from Its Tip*, *Journal of Theoretical Probability* **30** (2017), 27–63. MR3615081
- [26] Andrea Montanari, *Optimization of the Sherrington-Kirkpatrick Hamiltonian*, 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS), November 2019, pp. 1417–1433. MR4228234
- [27] Michel Pain, *The near-critical Gibbs measure of the branching random walk*, *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques* **54** (2018), 1622–1666. MR3825893
- [28] Robin Pemantle, *Search cost for a nearly optimal path in a binary tree*, *Annals of Applied Probability* **19** (2009), 1273–1291. MR2538070
- [29] Haskell P. Rosenthal, *On the subspaces of L^p ($p > 2$) spanned by sequences of independent random variables*, *Israel Journal of Mathematics* **8** (1970), 273–303. MR0271721
- [30] Eliran Subag, *Following the Ground States of Full-RSB Spherical Spin Glasses*, *Communications on Pure and Applied Mathematics* **74** (2021), 1021–1044. MR4230065
- [31] Hugo Touchette, *Equivalence and nonequivalence of ensembles: Thermodynamic, macrostate, and measure levels*, *Journal of Statistical Physics* **159** (2015), 987–1016. MR3345408
- [32] Ofer Zeitouni, *Branching random walks and Gaussian fields*, *Probability and statistical physics in St. Petersburg* **91** (2016), 437–471. MR3526836

Acknowledgments. We are grateful to an anonymous referee for several helpful suggestions improving the presentation.