Università degli studi Roma Tre

Dottorato di ricerca in Matematica

XXXVIII Ciclo

Nonlinear Exchange Dynamics for Spin Systems

Tesi di

Mario Morellini

Candidato:

Mario Morellini .....

Relatore:

Prof. Pietro Caputo .....

Coordinatore:

Prof. Alessandro Giuliani .....

Acknowledgements

I would like to begin by expressing my deep and sincere gratitude to my advisor, Pietro Caputo. To my mind, a teacher does not pass on a subject, but a way of doing it. From him I learned that there is no point in getting discouraged when facing problems larger than ourselves: small steps count too, and understanding comes by understanding. Being constructive and humble in the face of complexity is, in the end, a form of kindness.

"One has to understand whether the problem is conceptual or technical": this too is a lesson I carry with me. I would not claim to have learned to see things from above, but I count myself fortunate to have at least some idea of what the right way of doing research might be. Beyond this, I thank him for the entire weeks spent discussing, for his availability and his patience, and for accepting that there are bad stretches, simply waiting for them to pass.

It has been a privilege to work with Mehrad Abbaszadeh Minab, Zongchen Chen and Alistair Sinclair. I never felt the weight of their experience, only its benefit, and working with them taught me a great deal about the systems I was studying. The collaboration with them and with Pietro led to a paper I am proud of, and to time I have no doubt was well spent.

I would also like to thank Matteo Quattropani for his generosity with time and advice. What I am most grateful for is his deep and natural sense of solidarity towards other young mathematicians, which has the effect of making everyone, myself among them, feel part of the scientific community.

I am grateful to the Department of Mathematics and Physics of Roma Tre University, which provides an excellent environment for research and, at the same time, a serene and welcoming atmosphere that makes it a pleasure to spend one's working days there.

Contents

Acknowledgements i

Introduction 1

0.1 The models . . . 2
0.1.1 Ising Model . . . 2
0.1.2 Hardcore Model . . . 3
0.2 Exchange dynamics . . . 4
0.3 Main results . . . 6
0.3.1 Qualitative convergence. . . . 6
0.3.2 Quantitative entropy decay. . . . 7
0.3.3 Wasserstein contraction. . . . 9
0.4 Sampling Hardcore Model with nonlinear dynamics . . . 10
0.4.1 Nonlinear dynamics as Gibbs sampler . . . 13
0.4.2 Algorithms via particle systems . . . 15
0.5 Overview of the chapters . . . 18
0.6 Related work . . . 19

1 Nonlinear Ising Model 21

1.1 Introduction and main results . . . 21
1.1.1 The collision kernel . . . 21
1.1.2 Entropy decay . . . 24
1.1.3 Kac's program . . . 25
1.1.4 Wasserstein decay . . . 28
1.1.5 Plan of the chapter . . . 29
1.2 Structural properties of the nonlinear equation . . . 29
1.2.1 Preliminary facts . . . 29
1.2.2 Entropy and stationary states . . . 30
1.2.3 Conservation laws . . . 32
1.2.4 Proof of Theorem 1.5 . . . 33
1.2.5 Graphical representation . . . 35
1.2.6 Reduction to discrete time and the case $J = 0$ . . . 37
1.2.7 Proof of Lemma 1.18 . . . 39

1.2.8 Marked Partition Process . . . 39
1.2.9 Proof of Lemma 1.22 . . . 43
1.3 Proof of Theorem 1.11 . . . 44
1.3.1 Proof of Theorem 1.11 assuming Lemma 1.24 . . . 45
1.3.2 Proof of Lemma 1.24 . . . 46
1.4 Relative entropy decay for the particle system . . . 52
1.4.1 Mean field exchange dynamics . . . 52
1.4.2 Down-Up Walk for Conservative Particle Systems . . . 53
1.4.3 Proof of Theorem 1.27 . . . 55
1.4.4 Multicomponent mean field exchange dynamics . . . 61
1.4.5 Proof of Theorem 1.10 . . . 63
1.5 Completing Kac's program . . . 64
1.5.1 Nonlinear MLSI . . . 64
1.5.2 Local Central Limit Theorem and Kac Chaos . . . 65
1.5.3 Entropic and Fisher Chaos . . . 67
1.5.4 Proof of Theorem 1.7 . . . 70
1.5.5 Open problem: exponential ergodicity at all temperatures? . . . 71

2 Nonlinear Mean-Field Hardcore Model **72**
2.1 Introduction and Main results . . . 72
2.1.1 The collision kernel . . . 72
2.1.2 Nonlinear dynamics and convergence theorem . . . 73
2.1.3 Rates of convergence . . . 74
2.1.4 Kac's program . . . 75
2.1.5 Plan of the chapter . . . 78
2.2 Structural properties of nonlinear equation . . . 78
2.2.1 Entropy and stationary states . . . 79
2.2.2 Proof of Theorem 2.2 . . . 82
2.3 Wasserstein Contraction for low fugacity . . . 85
2.3.1 Proof of Lemma 2.15 . . . 87
2.4 Proof of Theorem 2.6 . . . 91
2.4.1 Reduction to Small Densities . . . 92
2.4.2 Martingale's Method in the Reduced Regime . . . 95
2.4.3 Proof of Theorem 2.21 . . . 97
2.4.4 Induced Chain . . . 101
2.4.5 Proof of Lemma 2.25 . . . 105
2.5 Completing Kac's Program . . . 110

3 Algorithmic Applications to Hardcore Model **114**
3.0.1 Plan of the chapter . . . 115
3.1 The particle system . . . 115

3.2 Algorithmic Applications . . . 123
3.2.1 Sampling with prescribed density . . . 124
3.2.2 Hardness of prescribed-density sampling . . . 133
3.2.3 Sampling with prescribed vertex marginals . . . 137

A Proof of Lemma 2.34 and Lemma 2.35 **158**

B **171**
B.1 The marginal polytope . . . 171
B.1.1 Gradient descent for prescribed marginals . . . 172
B.1.2 Bisection for prescribed density . . . 183
B.1.3 Missing proofs from Section 3.2.2 . . . 193

Introduction

A central problem in probability theory and statistical physics is to understand how large interacting systems relax to equilibrium. For linear Markov chains (Glauber dynamics, Kawasaki dynamics, and their relatives) this question has a rich theory, with spectral gaps, logarithmic Sobolev inequalities, and coupling methods providing sharp quantitative answers. These chains are also the standard tool for sampling, where the same estimates that describe relaxation to equilibrium bound the cost of the algorithm, and they have the further advantage of being simple to implement. This thesis addresses the same question for a class of *nonlinear* evolution equations. We consider Boltzmann-type evolution equations of the form

$$\frac{dp_t}{dt} = p_t \circ p_t - p_t, \tag{0.1}$$

where p_t is a probability distribution on a configuration space Ω, and the *collision product* $p_t \circ p_t$ is defined through a probability kernel $Q(\tau \mid \sigma, \sigma')$. The kernel takes as input two independent configurations $\sigma, \sigma' \in \Omega$, each sampled from p_t, and produces an output configuration $\tau \in \Omega$, with law

$$p_t \circ p_t(\tau) = \int p_t(d\sigma) \int p_t(d\sigma')\, Q(\tau \mid \sigma, \sigma'). \tag{0.2}$$

A classical example is Kac's caricature of a Maxwellian gas [Kac59], which captures essential features of the Boltzmann equation [Bol72] in a simplified setting; see also [Bas+22] for a discrete version of Kac's model. Such nonlinear equations admit a natural probabilistic representation in terms of random binary trees [Wil51; McK67]. A key structural property is the existence of *conservation laws*, associated to quantities that remain invariant under the collision product. Each conservation law gives rise to a family of stationary measures, which determine the possible limiting distributions of the dynamics.

For instance, in Kac's model, where $\Omega = \mathbb{R}$ is the set of one-dimensional velocities and collisions conserve the kinetic energy, for any initial distribution p_0 on $\mathbb{R}$ with mean zero and finite variance one has convergence to the centered Gaussian distribution with the same variance as p_0. Understanding the *rate of convergence to equilibrium* has been a longstanding question, and even in this simple setting it presents significant challenges [Kac56; Vil02; Vil03; Car+10; DGR09].

In this work we focus on a *combinatorial counterpart* of the Boltzmann equation, where Ω is the Boolean cube, and the analogue of the Maxwellian equilibrium is played by a high-

dimensional Gibbs measure, such as the Ising model on a graph. At infinite temperature, the model reduces to a version of the recombination equations commonly studied in population genetics [RSW92; RRS98; CP24]. When interactions are present, however, the system becomes substantially richer, and analyzing its convergence to stationarity raises new challenges. This thesis develops a quantitative theory of convergence to equilibrium for two important families of Gibbs measures: the *Ising model* and the *Hardcore model*.

0.1 The models

We consider two models from statistical physics: the Ising model and the Hardcore model. Both were introduced as toy models to describe mathematically physical phenomena known as phase transitions. Their mathematical study has since attracted interest across several fields of research, including probability theory [Mar97; Geo11; FV17], computer science [Wei06; Sly10; GŠV16], combinatorics [BW99], and social science [Kel85; CFL09].

Both models belong to the family of graphical *spin systems*, which model interactions through a graph. Given a graph $G = (V, E)$, the states of the system are called *configurations* $\sigma = (\sigma_i,\, i \in V)$, consisting of an assignment of a value σ_i from a list S, called *spin*, to each vertex of the graph, which in our case will take the form $S = \{-1, +1\}$ for the Ising model, where spins represent the magnetic moment of a particle, and $S = \{0, 1\}$ for the hardcore model, where spins represent occupation, with 1 indicating the presence of a particle and 0 its absence. A statistical weight is then assigned to each configuration, depending on the interaction between spins along the edges of the graph.

We now present each model with its precise mathematical definition. Throughout, we use the notation $[n] = \{1, \cdots, n\}$ to identify the set of vertex of the graph $V = [n]$ of the underlying graph.

0.1.1 Ising Model

The Ising model is one of the most studied models in statistical mechanics. It plays a key role in the rigorous understanding of phase transitions. We refer e.g. to [FV17; Dum23] for background and key facts.

The model is defined as follows. Let $\Omega = \{-1, 1\}^n$ denote the set of spin configurations. The Ising model is defined through the Gibbs distribution

$$\mu_{J,h}(\sigma) = \frac{1}{Z_{J,h}} \exp\Big(\tfrac{1}{2}\textstyle\sum_{i,j\in[n]} J_{i,j}\sigma_i\sigma_j + \sum_{i\in[n]} h_i\sigma_i\Big)\,, \qquad \sigma \in \Omega, \tag{0.3}$$

where $J = (J_{i,j})_{i,j\in[n]}$ is the interaction matrix, $h = (h_i)_{i\in[n]}$ is the external magnetic field, and $Z_{J,h}$ is the partition function ensuring normalisation, so that $\mu_{J,h}$ is a probability measure over Ω. We assume that J is a real symmetric matrix and h is a real vector. When $h = 0$ we simply write μ_J.

When $J = aA_G$, where $a \in \mathbb{R}$ and A_G is the adjacency matrix of an undirected graph G with n vertices, then $\mu_{J,h}$ represents the Ising model on the graph G with external field

h and interaction strength a. If $a > 0$ the model is called *ferromagnetic*, while it is called *antiferromagnetic* if $a < 0$.

A classical example is the Curie–Weiss model (or mean-field Ising model), where the interaction matrix is given by $J_{i,j} = \beta/n$ for all $i, j \in [n]$, with $\beta > 0$ denoting the inverse temperature. This corresponds to the Ising model on the complete graph with interaction strength $a = \beta/n$. Other examples include disordered systems such as the Sherrington–Kirkpatrick model, where $J_{i,j} = \frac{\beta}{\sqrt{n}} g_{ij}$, where $g_{ij} = g_{ji}$, $i < j$, are i.i.d. standard normal random variables, or the Ising model on a random d-regular graph with interaction strength $a \in \mathbb{R}$, where $d \geq 3$ is a fixed integer, in which case $J = aA_G$, where A_G is the adjacency matrix of a random d-regular graph with n vertices.

It is well known that if the interaction is sufficiently weak, the Gibbs measure exhibits rapid decay of correlations, and the system can be sampled efficiently through Markov chain algorithms such as Glauber dynamics; see [Mar97] for a comprehensive review in the case of the integer lattice. Weakness of the interaction can be formulated in several ways. A classical approach is to use a Dobrushin condition, see [Dob68] and [Hay06]. We also refer to [CLV23; Bla+22] and references therein for some recent extensions based on the concept of Spectral Independence introduced in [ALG21].

Here we study a nonlinear counterpart of the above classical results. That is, the linear evolution represented by a Markov chain such as Glauber or Kawasaki dynamics, in our case will be replaced by the quadratic evolution defined by a dynamical system such as (0.1).

We shall see that for a natural class of collision kernels, whose definition depends on J, one can obtain a general convergence result stating that the nonlinear process converges, for all initial distributions, to the target Gibbs measure (0.3) for a suitable choice of external fields h. Here an important role is played by the conservation laws of the dynamical system. Moreover, we shall obtain quantitative convergence statements under a suitable weakness hypothesis on the interaction J. Our results for the Ising model are obtained in collaboration with P.Caputo and have recently appeared in the paper [CM25].

0.1.2 Hardcore Model

In this model, a configuration assigns to each site an occupation variable $\sigma_i \in \{0, 1\}$, subject to the *hardcore constraint*: no two adjacent sites can be simultaneously occupied. In other words, a configuration is an independent set of the graph G.
Formally, given a graph $G = (V, E)$ with $|V| = n$, let Ω_G denote the set of *independent sets* of G, i.e. configurations $\sigma \in \{0, 1\}^n$ such that $\sigma_i \sigma_j = 0$ for all $\{i, j\} \in E$. The Gibbs measure with parameter $\lambda > 0$ is

$$\mu_{G,\lambda}(\sigma) = \frac{\lambda^{|\sigma|}}{Z_G(\lambda)} \mathbf{1}(\sigma \in \Omega_G), \qquad |\sigma| = \textstyle\sum_i \sigma_i, \tag{0.4}$$

where $Z_G(\lambda) = \sum_{\sigma \in \Omega_G} \lambda^{|\sigma|}$ is the *partition function*, which normalizes the measure. The

parameter λ plays the role of a fugacity: larger values of λ favor configurations with more occupied sites. In analogy with the weak interaction regime for the Ising model, one can find a low-fugacity regime, often called the uniqueness regime for the hardcore gas, such that the model exhibits rapid decay of correlations and natural randomized algorithms such as Glauber dynamics provide efficient sampling, see e.g. [ALG21; CLV23].

In this setting, the *uniqueness threshold* is given by

$$\lambda_c(\Delta) = \frac{(\Delta-1)^{\Delta-1}}{(\Delta-2)^{\Delta}}, \tag{0.5}$$

where Δ denotes the maximum degree of G.

This value can be understood as the critical fugacity of the infinite Δ-regular tree: for $\lambda \leq \lambda_c(\Delta)$ there exists a unique infinite-volume Gibbs measure, while for $\lambda > \lambda_c(\Delta)$ uniqueness fails, marking the presence of a phase transition [Kel85]. This threshold is also sharp from a computational perspective, see [Wei06] and [Sly10; SS12; GŠV16]. We refer to Section 0.4 for a full explained computational picture.

In our work, we shall develop a nonlinear counterpart to the above picture. Namely, we study convergence to stationarity of a nonlinear process of the form (0.1), for a collision kernel, whose definition depends on the graph G, that will guarantee the convergence to the Gibbs distribution (2.1) for a suitable choice or λ. As we will see, the collision conserves the initial density and thus, given any initial distribution p on Ω_G, we obtain convergence to the target $\mu_{G,\lambda}$, where λ is such that $\mu_{G,\lambda}$ and p have the same density α. This picture holds provided $\alpha < 1/(\Delta+1)$. Moreover, we shall derive quantitative convergence results in the uniqueness regime, that is for initial densities corresponding to a fugacity $\lambda < \lambda_c(\Delta)$. The conservation law also makes these dynamics natural candidates for sampling: since the density is preserved along the evolution, the process converges to the target measure without any knowledge of the corresponding fugacity, which is precisely the situation one faces when the model is parameterized by its density rather than by λ. Direct simulation of the nonlinear evolution is prohibitively expensive, but this can be bypassed by approximating it with a finite particle system, which yields polynomial-time samplers both at prescribed density and, using a different collision kernel, at prescribed marginal vector. We discuss this in Section 0.4. All our results for the hardcore model are obtained in collaboration with P. Caputo, Z. Chen, M. Abbaszadeh, and A. Sinclair, and is appeared recently in the paper [Min+26].

0.2 Exchange dynamics

To define equation (0.1) for these models, it is convenient to work with the following version of the kernel Q introduced in (0.2). Rather than specifying only the law of the output configuration τ, we prescribe the joint law of the pair (τ, τ') produced by the collision: this

is encoded in a kernel $\mathcal{Q}(\sigma, \sigma'; \tau, \tau')$, related to Q by

$$Q(\tau \mid \sigma, \sigma') = \sum_{\tau'} \mathcal{Q}(\sigma, \sigma'; \tau, \tau').$$

Working at the level of pairs is natural for exchange dynamics, where (τ, τ') is obtained from (σ, σ') by swapping spins, and makes the conservation laws of the dynamics transparent.

The precise definition of the kernel $\mathcal{Q}$ will be given in later sections. For the moment, let us only observe that $\mathcal{Q}$ consists of two main ingredients: a *mechanism*, which prescribes how the new pair (τ, τ') is generated from (σ, σ') and encodes the conserved quantities of the system, and a *rate*, which determines the likelihood of each such transition as a function of the current configurations. These nature of the kernel determines the family of stationary measures of the nonlinear process.

We focus on *exchange models*: when two configurations σ and σ' collide, they produce a new pair by exchanging one spin of σ with one spin of σ'. For the hardcore model we allow exchanges between σ_ℓ, σ'_k, where k, ℓ are pairs of sites chosen uniformly at random. In the case of the Ising model we have an additional layer of structure, in that we exchange σ_ℓ, σ'_k, where $\ell \in [n]$ is chosen uniformly at random, and k is chosen according to the distribution $\mathcal{K}(\ell, \cdot)$, where $\mathcal{K}$ is a symmetric stochastic matrix, the *transport kernel*.

Two transport kernels play a central role in this thesis. The *single-site model* corresponds to $\mathcal{K} = I$, so that only spins at the same site may be exchanged. The *mean-field model* corresponds to $\mathcal{K}(\ell, k) = 1/n$ for all ℓ, k, so that every pair of sites is exchanged with equal probability. Both are special cases of the *multi-component mean-field model*: given a partition $\mathcal{A}$ of the site set $[n]$, the kernel $\mathcal{K}_\mathcal{A}$ allows exchanges uniformly within each block of $\mathcal{A}$ and forbids exchanges across different blocks. See Subsection 1.1.1 for a detailed definition.

In the Ising model, the proposed exchange is accepted with a probability that depends on the Gibbs weights of the configurations involved in the swap, chosen so as to make the kernel $\mathcal{Q}$ reversible with respect to the product $\mu_J \otimes \mu_J$. In the hardcore model, the exchange is accepted whenever it produces two valid independent sets, and rejected otherwise. In both cases the collision kernel $\mathcal{Q}$ satisfies the *detailed balance condition* with respect to the product of two copies of the target Gibbs measure, which guarantees that every measure in the target family is a stationary solution of equation (0.1).

We refer to later sections for a detailed discussion. For the time being, we point out that in both models the total spin count is conserved by each collision. It follows that in the hardcore model the *density* $\alpha(p_t) = \frac{1}{n}\mathbb{E}_{\sigma \sim p_t}[|\sigma|]$ is a conserved quantity, and the nonlinear dynamics will converge to the unique stationary distribution of the form (2.1) with the same density as the initial distribution p_0. Similarly, in the Ising model with transport kernel $\mathcal{K}_\mathcal{A}$, the dynamics preserves the $\mathcal{A}$-*magnetization profile* $m_\mathcal{A}(p_t)$, defined as the vector of average magnetizations within each block of $\mathcal{A}$, and the nonlinear dynamics will converge to the unique stationary distribution of the form (0.3) with the same magnetization profile as the initial distribution p_0.

0.3 Main results

We now state the main results of the thesis. We present them for the two models separately, following the order in which the arguments are developed.

0.3.1 Qualitative convergence.

The first result establishes that, for the Ising model, the exchange dynamics always converges, regardless of the interaction strength. The notion of convergence we use throughout is the *total variation distance*, defined for two probability measures $p, q \in \mathcal{P}(\Omega)$ by

$$\|p - q\|_{TV} = \frac{1}{2} \sum_{\sigma \in \Omega} |p(\sigma) - q(\sigma)|.$$

Theorem 0.1 (Convergence — Ising model, informal; see Theorem 1.5)**.** *Let $J \in \mathbb{R}^{n \times n}$ be any symmetric interaction matrix and let $\mathcal{K}$ be any symmetric stochastic matrix with irreducible components forming a partition $\mathcal{A}$. For any initial state p_0 whose $\mathcal{A}$-magnetization profile is non-degenerate, the solution p_t of the nonlinear exchange dynamics (0.1) satisfies*

$$\|p_t - \mu_{J,h}\|_{TV} \longrightarrow 0 \quad \text{as } t \to \infty,$$

where h is the unique external field, constant on each block of $\mathcal{A}$, such that $\mu_{J,h}$ has the same $\mathcal{A}$-magnetization as p_0. Moreover, the $\mathcal{A}$-magnetization profile is constant along the evolution.

Theorem 0.2 (Convergence — hardcore model, informal; see Theorem 2.2)**.** *Let G be any graph with maximum degree $\Delta \geq 3$. For any initial state p_0 with density $\alpha(p_0) \in (0, \frac{1}{\Delta+1})$, the solution p_t of the mean-field hardcore exchange dynamics satisfies*

$$\|p_t - \mu_{G,\lambda}\|_{TV} \longrightarrow 0 \quad \text{as } t \to \infty,$$

where $\lambda = \lambda(p_0)$ is the unique fugacity such that $\mu_{G,\lambda}$ has density $\alpha(p_0)$. Moreover, the density is constant along the evolution.

We note that Theorem 0.2 holds for densities up to $1/(\Delta + 1)$. This is optimal since already with density $\alpha = 1/(\Delta + 1)$ one can have an example violating the convergence in Theorem 0.2. For example, take the complete graph with $\Delta + 1$ vertices and consider the probability measure which gives mass $1/(\Delta + 1)$ to each singleton. In this case one has $\alpha(p) = 1/(\Delta + 1)$ and the system is frozen at the initial state with $p_t = p$ for all $t \geq 0$.

The proofs of both theorems rest on a common strategy. The first ingredient is the *H-theorem*: the relative entropy is non-increasing along the dynamics for any fixed Gibbs measure μ, and the monotonicity is strict at a non-stationary state. The second ingredient is an *irreducibility* property: along any trajectory, the probability of every configuration eventually becomes and stays uniformly bounded away from zero. Together, these two

facts force convergence to the unique stationary state selected by the conserved quantity; see Section 1.2 for the Ising Case and Section 2.2 for the Hardcore Case.

This approach to the convergence of quadratic systems has its roots in the classical theory of the Boltzmann equation and the subsequent developments in the context of Mass Action Kinetics [Fei19]. In our context it was used by [RSW92] for non-interacting recombination models, and was subsequently developed in [CS24b] for the interacting single-site Ising dynamics. In both cases the argument to obtain the irreducibility relies on a fragmentation structure that is available only when exchanges are restricted to spins with the same label. When the transport kernel $\mathcal{K}$ allows exchanges between spins at different sites, a new difficulty arises: the fragmentation alone is no longer sufficient, and must be combined with a mixing step induced by $\mathcal{K}$ within each block of the partition $\mathcal{A}$. For the Ising model our analysis covers all symmetric transport kernels.

For the hardcore model the proof of irreducibility is rather different. A key feature of the irreducibility argument for the Ising model is that it can be reduced to the convergence of the same model in the non-interacting case $J = 0$, which is simple and tractable. This reduction is not available for the hardcore model: the hard exclusion constraint cannot be relaxed, and there is no simple non-interacting limit to compare the system with. The irreducibility argument therefore requires genuinely new combinatorial ideas, developed in [Min+26] for the first time; see Subsection 2.2.2.

0.3.2 Quantitative entropy decay.

The qualitative convergence results stated above give no information on the speed of convergence. A natural quantity to measure this speed is the *relative entropy* (or KL-divergence) of p_t with respect to the target Gibbs measure: for $\nu, \mu \in \mathcal{P}(\Omega)$,

$$H(\nu \mid \mu) = \begin{cases} \sum_{\sigma \in \Omega} \nu(\sigma) \log \dfrac{\nu(\sigma)}{\mu(\sigma)} & \text{if } \nu \ll \mu, \\ +\infty & \text{otherwise,} \end{cases} \tag{0.6}$$

where $\nu \ll \mu$ denotes absolute continuity of ν with respect to μ. By Pinsker's inequality,

$$\|p_t - \mu\|_{TV}^2 \leq \tfrac{1}{2} H(p_t \mid \mu),$$

the decay in relative entropy also implies decay of total variation distance to stationarity. A classical feature of nonlinear dynamics of Boltzmann type is that relative entropy is non-increasing along solutions; this is the content of the celebrated *H-theorem*, which has been at the heart of kinetic theory since Boltzmann's original work. The H-theorem identifies $H(p_t \mid \mu)$ as a Lyapunov functional, but does not by itself say at what rate this functional decays.

Exponential decay of the relative entropy $H(p_t \mid \mu)$ is equivalent to a functional inequality relating $H(p_t \mid \mu)$ to its time derivative, which gives rise to a nonlinear variant of the

so-called *modified logarithmic Sobolev inequality* (MLSI). Establishing such nonlinear MLSI inequality is notoriously difficult: in the Boltzmann setting, this is known as the Cercignani's conjecture, whose investigations led to impressive developments in the last few decades, see e.g. [Vil03].
In the linear Markov chain setting, by contrast, the corresponding family of inequalities, namely the standard log-Sobolev inequality and its modified version, is one of the most studied tools for quantifying convergence to equilibrium of Markov chains; see e.g. [DS96; Mar97; CP07; CDP09; Cap23] for an introduction. In recent years there has been substantial progress [BB19; Ana+21; CP21; Bla+22; HS23; CE22; CS24a; BBD24a].
Our approach is to bridge the nonlinear and linear settings via *Kac's program* [Kac56; Car+10; MM13]: we will show that a nonlinear MLSI for the dynamics p_t can be deduced from a (uniform-in-N) linear MLSI for an associated N-particle system. This reduces the nonlinear inequality to a linear one, at the price of having to deal with a high-dimensional setting.

The main quantitative contributions of this work are the following two theorems, one for each model.

Theorem 0.3 (Entropy decay — Ising model, informal; see Theorem 1.7)**.** *Let $J \in \mathbb{R}^{n\times n}$ be nonnegative definite with largest eigenvalue $\lambda(J) < \frac{1}{2}$, and let $\overline{J} = \max_i \sum_j |J_{ij}|$. For any partition $\mathcal{A}$ of $[n]$ and any initial state p_0, the solution p_t of the nonlinear multi-component mean-field Ising exchange dynamics satisfies*

$$H(p_t \mid \mu_{J,h}) \;\le\; H(p_0 \mid \mu_{J,h})\, e^{-rt/n}, \qquad t \ge 0,$$

where h is the unique external field compatible with the $\mathcal{A}$-magnetization of p_0, and the rate satisfies

$$r \;\ge\; \frac{1}{8}(1-2\lambda(J))^2\, e^{-16\overline{J}}.$$

For the next result we introduce the *critical density*

$$\alpha_c(\Delta) = \frac{\lambda_c(\Delta)}{1+(\Delta+1)\lambda_c(\Delta)}\,, \tag{0.7}$$

This corresponds to the occupation ratio of the complete graph with $\Delta+1$ vertices for the hardcore model with fugacity $\lambda_c(\Delta)$, see [DP23].

Theorem 0.4 (Entropy decay — hardcore model, informal; see Theorem 2.4)**.** *Let G be a graph with maximum degree $\Delta \ge 3$. For any $\gamma, \delta > 0$ and any initial state p_0 with density $\alpha(p_0) \in (\gamma, (1-\delta)\alpha_c(\Delta))$, the solution p_t of the mean-field hardcore exchange dynamics satisfies*

$$H(p_t \mid \mu_{G,\lambda}) \;\le\; e^{-rt/n}\, H(p_0 \mid \mu_{G,\lambda}),$$

for a rate $r = r(\gamma,\delta,\Delta) > 0$, where $\lambda = \lambda(p_0)$ is the unique fugacity matching the density of p_0.

The rate in Theorem 0.3 is of order $1/n$, which is optimal: even at zero interaction ($J = 0$), the relative entropy cannot decay faster than $e^{-ct/n}$ for some constant $c > 0$; see [CS24b]. It is also uniform over all choices of partition $\mathcal{A}$ and all values of the conserved magnetization. To our knowledge, Theorem 0.3 is the first instance of a high-dimensional interacting spin system for which Kac's program can be fully implemented quantitatively, yielding uniform exponential relative entropy decay independent of the conserved quantities. It is worthy of note that this situation is in sharp contrast with the classical Boltzmann-Kac model in kinetic theory; see [Kac56; Vil02; Car+10].

Both theorems are proved by implementing Kac's program: one lifts the nonlinear equation to a linear mean-field N-particle system, proves a uniform MLSI for it, and transfers the estimate back to the nonlinear equation as $N \to \infty$ via Entropic and Fisher Chaos, which in turn are based on the Local Central Limit Theorem. The two models require different constructions of the particle system and different methods to establish the MLSI, reflecting their different structures.

For the *Ising model*, the N-particle system is a reversible Markov chain on $\{-1, +1\}^{Nn}$ whose invariant measure is $\mu_J^{\otimes N}$ conditioned on a fixed $\mathcal{A}$-magnetization profile, the *multi-canonical Ising measure*. Inspired by recent work on Kawasaki dynamics [BBD24a], we derive the uniform MLSI for this system (Theorem 1.10) by using *stochastic localisation* [CE22; BBD24b]. This technique allows us to decompose the multi-canonical measure into a mixture of simpler measures, for each of which the the MLSI can be derived from log-concavity of the associated generating polynomials [CGM21; Ana+22] ; see Subsection 1.4.3 and Subsection 1.4.4. The quantitative transfer between components uses entropy factorisations in the spirit of the recent work [CCP25].

For the *hardcore model*, the N-particle system is a Kawasaki-type exchange dynamics on fixed-size independent sets of the disjoint union of N copies of G. Following the recent works [Jai+23; JPV26], we prove the MLSI for this exchange dynamics by first reducing the density to a low-enough value, and then by proving a Log-Sobolev inequality for the *Down–Up walk* on fixed-size independent sets; see Section 2.4.

0.3.3 Wasserstein contraction.

Alongside the entropy-based results, we establish exponential decay in the L^1-Wasserstein distance induced by the Hamming metric. Recall that the Hamming distance between two configurations $\sigma, \sigma' \in \Omega$ is $|D_{\sigma,\sigma'}| = \#\{i \in [n] : \sigma_i \neq \sigma'_i\}$; the corresponding Wasserstein distance on $\mathcal{P}(\Omega)$ is

$$W_1(p, q) = \inf_{\pi \in \Gamma(p,q)} \sum_{\sigma,\sigma' \in \Omega} \pi(\sigma, \sigma') \, |D_{\sigma,\sigma'}|, \tag{0.8}$$

where $\Gamma(p, q)$ is the set of couplings of p and q. This distance is methodologically significant for two reasons. First, unlike the total variation distance — which, as noted in [CLL25], need not be monotone along the nonlinear dynamics — W_1 is *monotone decreasing* along the nonlinear evolution, making it a natural Lyapunov functional. Second, the proof

is direct: it rests on an explicit coupling of two solutions starting from different initial conditions, without any auxiliary particle system or functional inequality. Kac's program, by contrast, requires a Local Central Limit Theorem, a uniform MLSI for the N-particle system, and Entropic Chaos estimates. The two approaches are thus complementary: Wasserstein contraction is simpler and self-contained; on the other hand, entropy decay requires more work while covering a larger temperature range in the Ising model, and a larger density range in the hardcore model.

Theorem 0.5 (Wasserstein contraction, informal; see Theorems 1.11 and 2.3)**.** *The following exponential decay in W_1 holds.*

- (Nonlinear multi-component mean-field Ising Model.) *If $\overline{J} = \max_i \sum_j |J_{ij}| < \frac{1}{10}$, then for any initial state p_0,*

$$W_1(p_t,\, \mu_{J,h}) \;\le\; W_1(p_0,\, \mu_{J,h})\, e^{-\varepsilon t/n}, \qquad \varepsilon = \tfrac{1}{2}(1 - 10\overline{J}),$$

  where h is the unique external field compatible with the $\mathcal{A}$-magnetization of p_0.

- (Nonlinear Hardcore model.) *If $\alpha(p_0) < \frac{1}{3(\Delta+1)}$, then*

$$W_1(p_t,\, \mu_{G,\lambda}) \;\le\; W_1(p_0,\, \mu_{G,\lambda})\, e^{-\varepsilon t/n}, \qquad \varepsilon = 1 - 3(\Delta+1)\alpha(p_0),$$

  where $\lambda = \lambda(p_0)$ is the unique fugacity matching the density of p_0.

We remark that, in both cases, one has

$$\|p_t - \mu\|_{TV} \le n\, e^{-\varepsilon t/n},$$

since $\|p - q\|_{TV} \le W_1(p, q) \le n$ for all $p, q \in \mathcal{P}(\Omega)$. The rate ε/n has the correct n-dependance: even at zero interaction, the bound decays no faster than $e^{-ct/n}$ for some constant $c > 0$; see [CS24b; RRS98]. For the Ising model, Wasserstein contraction requires $\overline{J} < 1/10$, which is more restrictive than the $\lambda(J) < 1/2$ of Theorem 0.3, so entropy decay covers a strictly larger regime at the cost of a deeper analysis. Similarly, for the hardcore model, the Wasserstein condition on the density is more restrictive than the density range in Theorem 0.4.

0.4 Sampling Hardcore Model with nonlinear dynamics

Sampling from discrete high-dimensional distributions is a fundamental problem in computer science. Many important applications, such as statistical estimation, causal inference, and experimental simulation, rely on massive samples drawn independently and approximately from a complex target distribution. Developing fast samplers for general distributions is thus a crucial task. In this thesis we focus on the *hard-core model*, and in particular on the sampling problem associated with nonlinear dynamics.

One of the most popular methods for generating samples approximately from a target distribution is the Markov Chain Monte Carlo (MCMC) method. This involves simulating a suitable Markov chain for a sufficiently long period and outputting the final state. Among the various MCMC methods, *Glauber dynamics*, also known as the *Gibbs sampler*, is a standard approach widely used in many applications. At each step, this dynamics picks a random coordinate and updates it according to the conditional marginal distribution with all other coordinates fixed.

The Glauber dynamics for the hard-core model has a simple description. If the current state is $\sigma \in \Omega_G$, we pick a vertex $v \in V$ uniformly at random and let $\sigma' := \sigma \cup \{v\}$. If $\sigma' \notin \Omega_G$, then we set the new state $\tau = \sigma$; otherwise, we set $\tau = \sigma'$ with probability $\frac{\lambda}{1+\lambda}$, and $\tau = \sigma' \setminus \{v\}$ with the remaining probability $\frac{1}{1+\lambda}$. It is easy to verify from elementary Markov chain theory that this process is ergodic (i.e., irreducible and aperiodic) and converges to the Gibbs distribution $\mu_{G,\lambda}$ as the number of updates tends to infinity. Its utility as an algorithm for sampling from $\mu_{G,\lambda}$ thus depends on its rate of convergence.

Let Δ be the maximum degree of the graph. Recent research has determined the exact threshold $\lambda_c(\Delta)$ as in (0.5) for the phase transition between fast and slow convergence of the Glauber dynamics. Specifically, the Glauber dynamics converges rapidly whenever $\lambda \leq \lambda_c(\Delta)$ [ALG21; Che+22; CE22; Che+25], and can be exponentially slow when $\lambda > \lambda_c(\Delta)$ [MWW09].[1] Here, the speed of convergence is characterized by the *mixing time*, which is the time until the distribution of the chain is close to $\mu_{G,\lambda}$ in total variation distance. Thus, Glauber dynamics provides an efficient sampling algorithm for the hard-core model at all fugacities $\lambda \leq \lambda_c(\Delta)$.

While fugacity is the classical parameterization of the hard-core model in statistical physics, many modern applications actually call for alternative parameterizations in terms of the density or occupancy probabilities of vertices in the independent set. The *density* of the hard-core model, defined as

$$\alpha = \alpha(\mu_{G,\lambda}) := \frac{1}{n}\mathbb{E}_{\mu_{G,\lambda}}[|\sigma|],$$

describes the average (normalized) size of an independent set drawn from the hard-core model.[2] The map $\lambda \mapsto \alpha(\mu_{G,\lambda})$ is strictly increasing and gives a one-to-one correspondence between the fugacity $\lambda \in (0, \infty)$ and the density $\alpha \in (0, i_{\max}/n)$, where $i_{\max}$ is the size of a largest independent set; see Appendix B.1.2. Thus, we can alternatively parameterize the hard-core model using α and represent its Gibbs distribution as μ_G^{α}. This approach is applicable when we want to sample at a specific density but do not have access to the corresponding fugacity.

We can extend the density to a finer parameterization in terms of the vector of marginal occupation probabilities of the vertices. To introduce this, we first consider a generalized

[1] This computational phase transition actually coincides exactly with a physical phase transition on the Δ-regular tree, defined by the onset of long-range correlations between vertices in the independent set; see, e.g., [Wei06].

[2] Throughout, we write $n = |V|$ for the number of vertices in G.

version of the hard-core model in which each vertex has a distinct fugacity. Given a vector of positive fugacities $\boldsymbol{\lambda} = (\lambda_v)_{v \in V} \in (0, \infty)^V$, we define the Gibbs distribution as

$$\mu_{G,\boldsymbol{\lambda}}(\sigma) = \frac{1}{Z_G(\boldsymbol{\lambda})} \prod_{v:\sigma_v=1} \lambda_v, \qquad \text{where} \quad Z_G(\boldsymbol{\lambda}) := \sum_{\sigma \in \Omega_G} \prod_{v:\sigma_v=1} \lambda_v.$$

If $\lambda_v = \lambda$ for all $v \in V$, we recover the hard-core model $\mu_{G,\lambda}$ with uniform fugacity λ. We now define the marginal vector $\boldsymbol{m} = (m_v)_{v \in V} \in (0, 1)^V$ as

$$m_v = \Pr_{\sigma \sim \mu_{G,\lambda}}(\sigma_v = 1).$$

Analogously to the bijection between fugacity λ and density α mentioned above, there is also a one-to-one correspondence between the fugacity vector $\boldsymbol{\lambda} \in (0, \infty)^V$ and the marginal vector $\boldsymbol{m} \in \mathcal{P}_{\mathrm{IS}}^\circ$; see Appendix B.1.1. Here, $\mathcal{P}_{\mathrm{IS}}$ denotes the independent set polytope (i.e., the convex hull of vectors in Ω_G), and $\mathcal{P}_{\mathrm{IS}}^\circ$ its interior. Therefore, the hard-core model can alternatively be parameterized using the marginal vector $\boldsymbol{m}$, with associated Gibbs distribution denoted by $\mu_G^{\boldsymbol{m}}$.

In this thesis, we study the sampling problem when we are given either the density α or the marginal vector $\boldsymbol{m}$, while the corresponding fugacity λ or, respectively, fugacity vector $\boldsymbol{\lambda}$ are unknown to us. The Gibbs distribution μ_G^α or $\mu_G^{\boldsymbol{m}}$ is known as the *maximum entropy distribution*: it is the unique distribution supported on Ω_G that achieves maximum entropy, when either the density or the marginal vector is fixed [Jay57a; Jay57b]. Maximum entropy distributions have been extensively studied for their applications in exponential families, machine learning, and combinatorial optimization: see, e.g., [Kap89; WJ08; OSS11; SV14b; Asa+17; SV19; CKV20; MSP22] for representative examples.

As a concrete toy example, consider a social network in which edges between two people represent mutual friendships. We want to select a committee from a distribution that is "as random as possible" (to prevent predictability or manipulation), subject to the constraints that no two selected individuals are friends and that each individual v is chosen with a given probability m_v (e.g, one important scenario is to set $m_v = c$ uniformly across individuals, thereby ensuring fairness of the selection). To achieve this, we should sample from the maximum entropy distribution over independent sets with marginal vector $\boldsymbol{m}$, which corresponds to the Gibbs distribution $\mu_G^{\boldsymbol{m}}$.

It is important to note that, without access to the fugacity parameter(s) λ or $\boldsymbol{\lambda}$, we cannot implement the Glauber dynamics for the hard-core model specified by a density α or a marginal vector $\boldsymbol{m}$. Indeed, the classical approach to this problem is to first estimate the fugacity parameter(s) given the density or marginals, and then use Glauber dynamics. This approach is known as solving the "inverse problem", and is well studied; see [NZB17] for a survey and [SV14b; BGS14; Mon15] for related results. Solving the inverse problem is in general a complex learning task relying on methods such as bisection or gradient descent, and involves tuning multiple parameters and the use of sampling or inference as subroutines. In the appendix, we sketch these approaches in the context of the hard-core

model and compare them with our new algorithms.

This discussion motivates the central question we aim to address in this thesis, namely:

- *Is there an efficient and easy-to-implement stochastic process that samples directly from the hard-core model with a given density or marginal vector?*

0.4.1 Nonlinear dynamics as Gibbs sampler

To address the algorithmic question above we consider a discrete time version of (0.1) specialized to independent sets. Namely, one iteration of the discrete-time dynamics is described by the quadratic map

$$T : p \mapsto p \circ p := \sum_{\sigma,\sigma' \in \Omega_G} p(\sigma)p(\sigma')\, Q(\cdot \mid \sigma, \sigma'), \tag{0.9}$$

where Q is based on exchange dynamics described in Section 0.2. For our sampling problem, we consider two different types of exchange collision kernels: mean field, suited for the density parameterization, and the second, single site is suited to the marginal vector parameterization. Recall that the mean-field kernel allows exchanges between any pair of vertices, while the single-site kernel only allows exchanges at matching vertices. We call the associated nonlinear dynamics, defined as in (0.9), the mean-field and the single-site dynamics respectively, and denote their one-step maps by T_{mf} and T_{ss}.We refer to Figure 1 for an illustration.

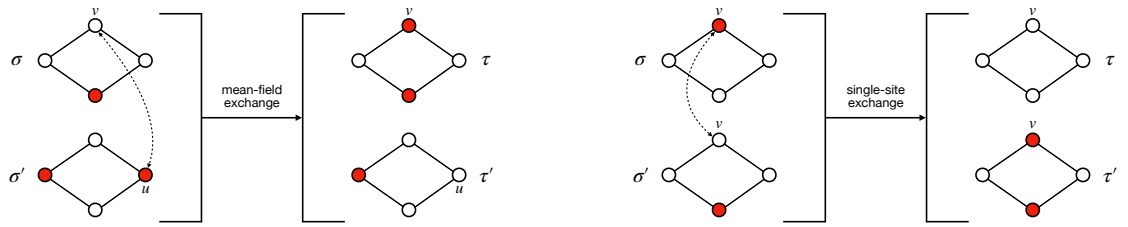


Figure 1: An illustration of the mean-field (left) and single-site (right) exchange kernels.

We can now explain why nonlinear dynamics are well suited to the sampling problem we want to solve. First, updates in nonlinear dynamics do not require knowledge of the fugacities. Second, as follows directly from the definitions of the corresponding collision kernels, the mean-field dynamics *preserves the density*, i.e. $\alpha(T_{\mathrm{mf}}(p)) = \alpha(p)$ for all distributions p over Ω_G, while the single-site dynamics *preserves the marginal vector*, i.e. $\boldsymbol{m}(T_{\mathrm{ss}}(p)) = \boldsymbol{m}(p)$; see Lemma 2.11 for the mean-field case. Moreover, both dynamics converge to a Gibbs distribution, a non-trivial fact stated in Proposition 0.6 below. Consequently, when initialized with the desired density or marginal vector, they converge to the corresponding maximum entropy distribution.

Proposition 0.6 ([Min+26]). *Consider the mean-field and single-site nonlinear dynamics on independent sets in a graph $G = (V, E)$ of maximum degree Δ.*

(1) If p is a distribution on Ω_G with density $\alpha = \alpha(p)$ satisfying $0 < \alpha < \frac{1}{\Delta+1}$, then

$$\lim_{t\to\infty} T^t_{\mathrm{mf}}(p) = \mu^{\alpha}_G.$$

(2) If p is a distribution on Ω_G with marginal vector $\boldsymbol{m} = \boldsymbol{m}(p)$ satisfying $0 < m_v < \frac{1}{\Delta+1}$ for all $v \in V$, then

$$\lim_{t\to\infty} T^t_{\mathrm{ss}}(p) = \mu^{\boldsymbol{m}}_G.$$

We omit the proof, which amounts to establishing an irreducibility property for the discrete-time process: for all large enough times, every state is reached with positive probability, uniformly over the initial distributions. For the mean-field kernel, this property is proved in the continuous-time setting in Section 2.2.2, and the argument adapts easily to discrete time; the single-site kernel is handled similarly. Full details can be found in [Min+26, Propositions 3.5 and 3.7], for the mean-field and the single-site case respectively.

While nonlinear dynamics in principle provide a simple and elegant solution to our sampling problem, a number of substantial technical challenges remain if we are to turn this idea into an efficient algorithm. First, in contrast to linear Markov chains, which have a well-established theoretical foundation and extensive literature, research on nonlinear dynamics is very limited. For instance, we have a comprehensive understanding of the convergence of Markov chains, under simple conditions, at an exponential rate. However, for nonlinear Markov chains, there are no general conditions that ensure convergence, at *any* rate, to a unique stationary distribution.[3] Additionally, many techniques, such as coupling and functional inequalities, that have been successfully applied to the analysis of the speed of convergence, or mixing time, of Markov chains do not have immediate analogs in the nonlinear case.

Furthermore, while nonlinear dynamics may be straightforward to describe, their efficient implementation is not trivial. As is clear from (0.9), each iteration of the dynamics requires two independent samples from the current state. Consequently, to obtain a sample after t iterations, 2^t independent samples from the initial distribution are needed, which is computationally prohibitive. In contrast, a linear Markov chain can be simulated by following the trajectory of a single sample.

[3] Notable exceptions for general nonlinear Markov chains are [And+11] and [But14], but the conditions for convergence that these works require are far too restrictive to apply in our setting. A famous open problem known as the "Global Attractor Conjecture" [Fei19] postulates the convergence of any nonlinear dynamics under fairly mild conditions.

0.4.2 Algorithms via particle systems

The direct implementation of nonlinear dynamics can be prohibitive because generating a single sample from $T(p)$ requires two independent samples from p. Consequently, producing an output from $T^t(p)$ after t iterations necessitates 2^t independent samples from the initial distribution p.

Applying Lemma 2.15 iteratively, used in the proof of Theorem 0.5, and comparing the Wasserstein and total variation distances, one deduces that for every distribution $p \in \mathcal{P}(\Omega_G)$ with $\alpha(p) < \frac{1}{3(\Delta+1)}$,

$$\|T^t_{\mathrm{mf}}(p) - \mu^{\alpha}_G\|_{\mathrm{TV}} \leq n\, e^{-\varepsilon t/n},$$

where $\varepsilon := \varepsilon(\alpha)$ is as prescribed in the lemma. This means that our mixing time is $t = O(n \log n)$ and the overall runtime becomes exponential. We claim that similar result of Lemma 2.15 can be established for the single site case, leading to the same negative conclusion for T_{ss}, see [Min+26, Theorem 4.3 and Lemma 4.4].

We can overcome this obstacle using the idea of *particle systems*, a concept from kinetic theory dating back to Mark Kac [Kac56]. The idea is to approximate the nonlinear evolution via a finite population of N particles $\sigma^{(i)}$, $i = 1, \ldots, N$, each represented by an independent set $\sigma^{(i)} \in \Omega_G$, which undergo pairwise collisions

$$(\sigma^{(i)}, \sigma^{(j)}) \mapsto (\tau^{(i)}, \tau^{(j)})$$

where the pair $(\tau^{(i)}, \tau^{(j)})$ is obtained from the pair $(\sigma^{(i)}, \sigma^{(j)})$ through the exchange of two bits $\sigma^{(i)}_v$ and $\sigma^{(j)}_u$, according to a choice of sites u, v that matches the behavior of the nonlinear collision kernel. This evolution is actually a (linear) Markov chain on the set Ω^N_G; crucially, it is efficiently implementable (provided N is not too large). Moreover, one expects that if the number of particles N is sufficiently large, then the particle system should approximate the evolution of the nonlinear dynamics. In particular, since the nonlinear dynamics converges to the desired Gibbs distribution, then after sufficiently many steps the law of a typical particle should be close to this target measure.

This suggests the following general algorithmic template:

(1) Initialize N particles so that the relevant conserved quantity (density α or marginal vector $\boldsymbol{m}$) has the desired value.

(2) Run the corresponding particle-system chain for t steps, where t is large enough.

(3) Output one particle.

The analysis of this algorithm involves two main steps. First, we show that, for N sufficiently large, the one-particle marginal of the stationary distribution of the particle system is close to the desired hard-core measure. Second, we bound the mixing time of the particle-system chain, which determines how large t should be. Combining these two

steps shows that the output particle is approximately distributed according to the target measure in polynomial time.

For the mean-field dynamics, we carry out this program for densities up to the critical density $\alpha_c(\Delta)$ and obtain a new sampling algorithm for the hard-core Gibbs distribution parameterized by the density.

Theorem 0.7 (informal version of Theorem 3.9)**.** *There is a sampling algorithm (Algorithm 4) such that, for any graph G on n vertices with maximum degree $\Delta \geq 3$, any density $\delta \leq \alpha \leq (1-\delta)\alpha_c(\Delta)$ where $\delta \in (0,1)$, and any approximation error $\varepsilon \in (0,1)$, the algorithm runs in time*[4] $\widetilde{O}(\frac{n}{\varepsilon})$ *and outputs an independent set from a distribution μ_{alg} on Ω_G satisfying*

$$\|\mu_{\mathrm{alg}} - \mu_G^\alpha\|_{\mathrm{TV}} \leq \varepsilon.$$

We note that the critical threshold $\alpha_c(\Delta)$ in Theorem 0.7 is the best possible, beyond which the sampling problem becomes computationally hard. Specifically, in Theorem 3.11, we show that there is no polynomial-time algorithm to sample from μ_G^α for any fixed $\alpha \in (\alpha_c(\Delta), \frac{1}{\Delta+1})$, assuming NP $\neq$ RP. This result follows from an adaptation of the approach used in [DP23] for proving the hardness of sampling from $\mu_{G,\lambda}(\cdot \mid |\sigma| = \alpha n)$ under the same regime (that is, uniformly random independent sets of size αn).

The proof of Theorem 0.7 follows the two-step strategy mentioned above. For the first step, let $\mu_{N,\alpha}^{\mathrm{mf}}$ be the stationary distribution of the mean-field particle-system chain, where $M = \lfloor Nn\alpha \rfloor$ is the prescribed total occupation over all particles. This distribution is uniform over all N-particle configurations with total occupation M. Equivalently, it can be viewed as N independent samples from μ_G^α conditioned on the event that the total number of occupied vertices over all the particles is exactly M. We prove that this conditioning does not significantly change the law of a typical particle. The intuition is that M is chosen close to the expected total occupation under the product measure $(\mu_G^\alpha)^{\otimes N}$, and fixing one particle σ only changes the required total occupation of the remaining particles by $|\sigma|$, which is at most n. When N is large, this change is small compared to the fluctuations of the total occupation over all particles. We make this intuition precise using local central limit theorem estimates, which show that the one-particle marginal of $\mu_{N,\alpha}^{\mathrm{mf}}$ is within variation distance ε of μ_G^α when $N = \widetilde{O}(\frac{1}{\varepsilon})$. For the second step, we use entropy contraction for the mean-field particle system to show that the particle-system chain converges to $\mu_{N,\alpha}^{\mathrm{mf}}$ within time $\widetilde{O}(Nn)$. This step relies crucially on a recent result from [Jai+23; JPV26], which allows us to obtain a tight entropy production bound for the particle system at all subcritical densities. Therefore, outputting one particle from the resulting configuration gives a sample close to μ_G^α.

We go on to provide an analogous analysis for the particle system simulation of the single-site dynamics.

[4] Hereafter, we use $\widetilde{O}(\cdot)$ to suppress dependencies on $(\log n, \log(1/\varepsilon))$.

Theorem 0.8 (informal version of Theorem 3.18)**.** *There is a sampling algorithm (Algorithm 6) such that, for any graph G on n vertices with maximum degree $\Delta \geq 0$, any marginal vector $\boldsymbol{m} \in (0,1)^V$ such that $\delta \leq m_v \leq \frac{1-\delta}{2(\Delta+1)}$ for all $v \in V$ where $\delta \in (0,1)$, and any approximation error $\varepsilon \in (0,1)$, the algorithm runs in time $\widetilde{O}(\frac{n^3}{\varepsilon^4})$ and outputs an independent set from a distribution μ_{alg} on Ω_G satisfying*

$$\|\mu_{\mathrm{alg}} - \mu_G^{\boldsymbol{m}}\|_{\mathrm{TV}} \leq \varepsilon.$$

Remark 0.9. In fact, our sampling algorithm works as long as all marginals are less than $\frac{1}{c_{\mathrm{col}}(\Delta+1)}$, where c_{col} is the smallest constant such that the Glauber dynamics for sampling q-colorings on any graph of maximum degree Δ (not necessarily bounded) mixes in nearly linear time whenever $q \geq c_{\mathrm{col}}\Delta$. Under the widely believed conjecture that Glauber dynamics has optimal mixing when $q \geq \Delta + 2$, we would achieve the marginal bound $\frac{1-\delta}{\Delta+1}$ for any $\delta \in (0,1)$, matching the constraint in Theorem 0.6. This would not contradict the hardness derived from [DP23] as discussed immediately following Theorem 0.7, since that hardness corresponds to constraints on the global density rather than individual marginals.

The proof of Theorem 0.8 is again based on the two steps as before, but the ingredients are different from the density case in Theorem 0.7. Let $\vec{r} = (r_v)_{v \in V}$ be an integer vector with $r_v = Nm_v$, and let $\mu^{\mathrm{ss}}_{N,\vec{r}}$ be the stationary distribution of the single-site particle-system chain. This distribution is uniform over all N-particle configurations in which exactly r_v particles contain vertex v, for every $v \in V$. Equivalently, it can be viewed as N independent samples from $\mu_G^{\boldsymbol{m}}$ conditioned on having these exact vertex counts. For the first step, we show that this conditioning does not significantly change the law of a typical particle. By a standard relative entropy argument, it is enough to show that the exact-count event has sufficiently large probability under $(\mu_G^{\boldsymbol{m}})^{\otimes N}$. To lower bound this probability, we first use concentration inequalities and a union bound to show that the vertex counts are close to their expected values with high probability, and then compare the probabilities of nearby count classes. This shows that, for $N = \widetilde{O}(\frac{n^2}{\varepsilon^4})$, the one-particle marginal of $\mu^{\mathrm{ss}}_{N,\vec{r}}$ is close enough to $\mu_G^{\boldsymbol{m}}$ with ε total variation error. For the second step, we prove rapid mixing of the single-site particle-system chain by comparing it with Glauber dynamics for proper colorings of an auxiliary graph constructed from G and $\vec{r}$. This gives mixing after $\widetilde{O}(Nn)$ many steps. Combining these two steps shows that, after running the chain for its mixing time and outputting one particle, the output distribution is close to $\mu_G^{\boldsymbol{m}}$.

We remark that, while our algorithms based on particle systems are crucially inspired by the nonlinear dynamics, their analysis does not depend on the convergence rate of those nonlinear dynamics, nor do they imply such a rate.

We conclude the introduction by comparing our algorithms based on nonlinear dynamics with more classical approaches that involve solving the inverse problem (i.e., learning the relevant fugacities given the density or marginal vector). The formal statements and proofs are included in Appendices B.1.1 and B.1.2, and we summarize them here.

To sample from μ_G^α with a specified density $\alpha \leq (1-\delta)\alpha_c(\Delta)$, we can begin by utilizing the bisection method (as outlined in Algorithm 8) to estimate the corresponding fugacity λ of the hard-core model. Then, we can generate samples using efficient sampling algorithms, such as Glauber dynamics, based on this value of λ. The running time of such an algorithm is $\widetilde{O}(\frac{n}{\varepsilon^2})$, whereas our algorithm achieves a faster running time $\widetilde{O}(\frac{n}{\varepsilon})$.

To sample from the distribution $\mu_G^{\boldsymbol{m}}$ with a specified marginal vector $\boldsymbol{m}$, we can utilize the gradient descent method (see Algorithm 7) to determine the fugacity vector $\boldsymbol{\lambda}$. Once we have (an estimate of) $\boldsymbol{\lambda}$, we can sample from $\mu_{G,\boldsymbol{\lambda}}$ using Glauber dynamics. The running time for this algorithm is $\widetilde{O}(\frac{n^2}{\varepsilon^2})$, which is more efficient than what we claimed in Theorem 0.8 for our algorithm. It is thus interesting to see whether our analysis can be improved to match this quadratic running time. One possibility is to apply strong local central limit theorems as we did in the mean-field case; however, this becomes technically much more challenging in the single-site case and the bounds we have been able to obtain so far are not sufficient to establish the desired result.

The gradient descent algorithm works when all marginal probabilities are less than $\alpha_c(\Delta)$. We note that our algorithm is potentially effective over the wider range of marginals up to $\frac{1}{\Delta+1}$, assuming the widely accepted mixing conjecture on colorings (see Theorem 0.9). This raises the interesting question of whether the sampler based on gradient descent could also work efficiently in this regime.

In summary, we see that in the prescribed-density case, our nonlinear dynamics based algorithm has a better performance guarantee than the traditional learning method. In the prescribed-marginal case, however, the performance guarantee for our algorithm is slightly worse than that of the gradient descent approach. Nevertheless, in both cases, our algorithms have the advantage of being very simple to implement, involving only the direct simulation of a particle system of suitable size with no additional computation. Moreover, they represent a potentially powerful new computational paradigm whose full power is far from fully understood.

0.5 Overview of the chapters

Chapter 1: Nonlinear Ising Model. This chapter develops all our results for the Ising model. With the exception of the Wasserstein contraction (Theorem 0.5), the results of this chapter appeared in the preprint [CM25]. The chapter is organised as follows. Section 1.2 establishes the structural properties of the nonlinear dynamics: detailed balance, conservation of magnetization, characterisation of stationary states, and proof of Theorem 0.1. Section 1.3 contains the proof of Wasserstein contraction (Theorem 0.5 for Ising Model), based on an explicit coupling. Section 1.4 contains the proof of the uniform MLSI for the particle system, via stochastic localisation and entropy factorisation. Section 1.5 completes Kac's program by transferring the MLSI to the nonlinear equation via Entropic and Fisher Chaos, thereby proving Theorem 0.3.

Chapter 2: Nonlinear Hardcore Model. This chapter studies the nonlinear hard-core

dynamics with the mean-field kernel and establishes its convergence to equilibrium. Part of the material comes from joint work with Abbaszadeh, Caputo, Chen, and Sinclair [Min+26]. Section 2.2 establishes structural properties of equation (0.1) and proves Theorem 0.2, including the irreducibility argument specific to the hard-constraint setting. Section 2.3 contains the proof of Wasserstein contraction (Theorem 0.5 fo the Hardcore model). The uniform MLSI for the hardcore particle system is proved in Section 2.4. Section 2.5 then completes Kac's program for the hardcore model (Theorem 0.4), establishing propagation of chaos by means of a local central limit theorem.

Chapter 3: Algorithmic Applications to Hardcore Model. We presents the algorithmic consequences of the nonlinear dynamics for the hard-core model. The material is joint work with Abbaszadeh, Caputo, Chen, and Sinclair [Min+26]. Section 3.1 introduces the mean-field and single-site particle systems in discrete time and establishes their mixing properties. Section 3.2.1 turns these into an algorithm for sampling μ_G^α at a prescribed density α as in Theorem 0.7, and Section 3.2.2 proves a matching hardness result, showing that the regime $\alpha < \alpha_c(\Delta)$ in which the algorithm operates is optimal, see Theorem 3.11 for a detailed statement. Finally, Section 3.2.3 treats the prescribed marginals case, where the single-site particle system yields a polynomial-time sampler for $\mu_G^{\boldsymbol{m}}$, as presented in Theorem 0.8.

0.6 Related work

Equation (0.1) appears across several fields. In kinetic theory it is a central object of study [Vil02]; in chemistry it appears as mass-action kinetics [Fei19]; in biology it models genetic recombination [RSW92; RRS98]; in computer science it underlies genetic algorithms [Gol89; Mit98] and random sampling [SV14a]. The framework of nonlinear Markov chains [McK66] provides a convenient unifying language.

The spin-system version of equation (0.1) was introduced as *Reversible Quadratic Systems* in [CS18], where relative entropy decay was analysed for recombination models at infinite temperature. Quantitative mixing and cutoff results for non-interacting models were subsequently obtained in [CP24; CLL25]. The first quantitative results for an interacting quadratic system — exponential total variation decay at high temperature for a single-site model — appear in [CS24b]. The present work extends those results in three directions simultaneously: arbitrary transport kernels, arbitrary magnetization profiles, and exponential decay in relative entropy rather than total variation.

Proving exponential entropy decay for kinetic equations is notoriously difficult [Vil03; Car+10; MM13; Bon+18; EFS20]. Kac's program and the associated concept of Propagation of Chaos [Kac56] are central tools; we refer to [MM13] for a detailed analysis in the Boltzmann setting and to [CD22] for a recent survey. In the linear setting, MLSI inequalities for conservative spin systems (Kawasaki dynamics) under spatial mixing conditions have been studied in [Yau96; CMR02], but they do not provide uniformity in the number of particles N in the mean-field setting that is relevant to our problem. Our

approach, inspired by recent developments in Kawasaki dynamics [BBD24a], achieves this uniformity via stochastic localisation [CE22; BBD24b], entropy factorisation [CCP25], and the log-concavity of generating polynomials of strongly Rayleigh measures [CGM21; Ana+22].

For the hardcore model, the Kac program naturally leads to the analysis of an exchange dynamics for fixed-size independent sets, which is reversible with respect to the canonical ensemble. Canonical ensembles for the hardcore model were previously studied by Davies and Perkins [DP23]. The log-Sobolev inequality for the Down–Up walk that underpins our MLSI is due to [Jai+23; JPV26], which is the first instance of optimal mixing for sub-critical fixed-size independent sets. To our knowledge, the present work is the first to study a nonlinear evolution equation of the form (0.1) whose stationary measures are Gibbs measures of the hardcore model; the hard exclusion constraint makes this setting genuinely new compared to all previously studied instances of nonlinear spin dynamics.

The exponential decay in terms of Wasserstein distance, stated Theorem 0.5, rests on coupling techniques that are novel and differ from those used for linear Markov chains. For instance, considering the nonlinear hardcore model case, the well-known path coupling argument for Markov chains [BD97] cannot be applied directly to nonlinear dynamics. This is because creating a path between two independent sets either requires the inclusion of non-independent sets along the path, or alters the density or marginal vector of the intermediate independent sets, which poses considerable challenges for analysis. Additionally, the traditional coupling approach requires showing that the expected one-step contraction in Hamming distance occurs for two variables X_0 and Y_0 generated from any pair of coupled initial distributions p and q. Thus, it is sufficient to consider *fixed* initializations, such that $X_0 = x_0$ and $Y_0 = y_0$ (in other words, p and q are both delta measures). However, for our analysis, both p and q must have the specified density or marginal vector, and our coupling analysis critically depends on this fact.

Moreover, these coupling techniques yield a simpler route to convergence results, and at the same time cover a broader parameter regime. For the single-site dynamics for the Ising model as in Theorem 0.5, for instance, they give convergence whenever $\max_i \sum_j |J_{ij}| < \frac{1}{10}$, strictly improving on the range obtained in [CS24b] with a more elaborate recursive construction.

Chapter 1

Nonlinear Ising Model

This chapter is entirely devoted to the analysis of the nonlinear exchange dynamics for the Ising model. The content presented here is based on the article [CM25].
For the sake of completeness, we briefly recall the mathematical definition of the model. Let $\Omega = \{-1, +1\}^n$ denote the spin configuration space. The Ising model is defined through the Gibbs measure

$$\mu_{J,h}(\sigma) = \frac{1}{Z_{J,h}} \exp\left(\tfrac{1}{2} \sum_{i,j \in [n]} J_{i,j}\, \sigma_i \sigma_j + \sum_{i \in [n]} h_i\, \sigma_i \right), \qquad \sigma \in \Omega,$$

where $J \in \mathbb{R}^{n \times n}$ is a symmetric matrix, referred to as the *interaction matrix*, $h \in \mathbb{R}^n$ is the vector of *external fields*, and $Z_{J,h}$ is the normalizing partition function.
The results discussed in this chapter were already presented informally in the introduction of this thesis. Here, we expose them in full detail, providing the complete statements of the theorems together with their proofs.

1.1 Introduction and main results

1.1.1 The collision kernel

A collision between two spin configurations $(\sigma, \sigma') \in \Omega \times \Omega$ consists in the exchange of the ℓ-th spin of σ with the k-th spin of σ', that is

$$(\sigma, \sigma') \mapsto \left(S_\ell(\sigma, \sigma'_k),\ S_k(\sigma', \sigma_\ell) \right), \tag{1.1}$$

where $S_\ell : \Omega \times \{-1, 1\} \to \Omega$ is the configuration defined as

$$S_\ell(\sigma, a)_i = \begin{cases} \sigma_i & \text{if } i \neq \ell, \\ a & \text{if } i = \ell. \end{cases}$$

The sites (ℓ, k) at which the exchange occurs are chosen by picking ℓ uniformly at random in $[n]$ and then by picking k according to a prescribed *symmetric stochastic* $n \times n$ *matrix*,

denoted $\mathcal{K}(\ell, k)$, which we refer to as the transport kernel. In order to guarantee the reversibility of our system with respect to the Ising measure, the proposed exchange (1.1) is accepted with probability

$$p_J(\ell, k \,|\, \sigma, \sigma') = \frac{\mu_J(S_\ell(\sigma, \sigma'_k))\, \mu_J(S_k(\sigma', \sigma_\ell))}{\mu_J(\sigma)\, \mu_J(\sigma') + \mu_J(S_\ell(\sigma, \sigma'_k))\, \mu_J(S_k(\sigma', \sigma_\ell))}, \tag{1.2}$$

and is otherwise rejected. Summarizing, the pair $(\sigma, \sigma') \in \Omega \times \Omega$ produces the new pair $(\tau, \tau') \in \Omega \times \Omega$ according to the probability kernel on $\Omega \times \Omega$ given by

$$\mathcal{Q}_J^{\mathcal{K}}(\sigma, \sigma'; \tau, \tau') = \frac{1}{n} \sum_{\ell,k=1}^{n} \mathcal{K}(\ell, k)\; Q_J^{\ell,k}(\sigma, \sigma'; \tau, \tau'), \tag{1.3}$$

where, for each pair of sites $\ell, k \in [n]$, $Q_J^{\ell,k}$ is the probability kernel on $\Omega \times \Omega$ defined as

$$\begin{aligned} Q_J^{\ell,k}(\sigma, \sigma'; \tau, \tau') = p_J(\ell, k \,|\, \sigma, \sigma') \mathbf{1}(\tau = S_\ell(\sigma, \sigma'_k), \tau' = S_k(\sigma', \sigma_\ell)) + \\ + \; (1 - p_J(\ell, k \,|\, \sigma, \sigma')) \mathbf{1}(\tau = \sigma, \tau' = \sigma'), \end{aligned} \tag{1.4}$$

and we use the notation $\mathbf{1}(E)$ for the indicator function of the event E. If $\mathcal{A}$ denotes the set of equivalence classes corresponding to the irreducible components of the kernel $\mathcal{K}$, it follows that $\mathcal{Q}_J^{\mathcal{K}}$ satisfies the detailed balance condition with respect to $\mu_{J,h} \otimes \mu_{J,h}$, for any choice of external fields $h \in \mathbb{R}^n$ such that h is constant over every subset $A \in \mathcal{A}$:

$$\mu_{J,h}(\sigma)\mu_{J,h}(\sigma')\mathcal{Q}_J^{\mathcal{K}}(\sigma, \sigma'; \tau, \tau') = \mu_{J,h}(\tau)\mu_{J,h}(\tau')\mathcal{Q}_J^{\mathcal{K}}(\tau, \tau'; \sigma, \sigma'),$$

for arbitrary $(\sigma, \sigma'), (\tau, \tau') \in \Omega \times \Omega$ and $h \in \mathbb{R}^n$ such that $h_i = h_j$ for all $i, j \in A$, for all $A \in \mathcal{A}$; see Lemma 1.13.

In conclusion, for any given interaction $J \in \mathbb{R}^{n \times n}$, and any transport kernel $\mathcal{K}$, we have constructed a probability kernel $\mathcal{Q}_J^{\mathcal{K}}$ on $\Omega \times \Omega$, which satisfies the detailed balance condition with respect to $\mu_{J,h} \otimes \mu_{J,h}$, independently of the choice of h, provided h is constant over the irreducible components of $\mathcal{K}$.

Nonlinear dynamics

Let $\mathcal{P}(\Omega)$ denote the set of probability measures on Ω. Given the collision kernel $\mathcal{Q}_J^{\mathcal{K}}$, the nonlinear evolution on $\mathcal{P}(\Omega)$ is defined by the dynamical system

$$\frac{dp_t}{dt} = p_t \circ p_t - p_t, \qquad p_0 = p, \tag{1.5}$$

for some initial value $p \in \mathcal{P}(\Omega)$, where $p \circ q \in \mathcal{P}(\Omega)$, for $p, q \in \mathcal{P}(\Omega)$, denotes the commutative collision product defined as

$$p \circ q\,(\tau) = \frac{1}{2} \sum_{\sigma, \sigma', \tau' \in \Omega} \big(p(\sigma)q(\sigma') + p(\sigma')q(\sigma)\big)\, \mathcal{Q}_J^{\mathcal{K}}(\sigma, \sigma'; \tau, \tau'). \tag{1.6}$$

Note that this has the form (0.2) with

$$Q(\tau \mid \sigma, \sigma') = \textstyle\sum_{\tau' \in \Omega} \mathcal{Q}_J^{\mathcal{K}}(\sigma, \sigma'; \tau, \tau') \,.$$

For any $J \in \mathbb{R}^{n \times n}$ and any stochastic matrix $\mathcal{K}$, the Cauchy problem (1.5) has a unique solution $p_t \in \mathcal{P}(\Omega)$, $t \geq 0$, for any initial state $p \in \mathcal{P}(\Omega)$. This can be seen by standard existence and uniqueness results, see e.g. [BBS15] for related statements. Moreover, a representation of the solution based on random binary trees can be obtained, see e.g. [CCG00; CP24; CLL25] for similar constructions. Here, we are interested in the structure of the stationary states of (1.5), that is the set of $p \in \mathcal{P}(\Omega)$ such that $p_t \equiv p$, and in the analysis of the convergence to stationarity. As we will see, the set $\mathcal{A}$ of irreducible components of the kernel $\mathcal{K}$ determines the conservation laws of the dynamical system (1.5) and the set of stationary states is fully determined by J and $\mathcal{A}$. The following are some examples of kernels $\mathcal{K}$ to keep in mind.

Example 1.1 (Single-Site Dynamics)**.** If $\mathcal{K}(i,j) = \mathbf{1}(i = j)$ for all $i, j \in [n]$, exchanges are allowed only between spins at the same site. With the notation (1.4), the collision kernel is given by

$$\mathcal{Q}_J^{\mathcal{K}}(\sigma, \sigma'; \tau, \tau') = \frac{1}{n} \sum_{\ell=1}^{n} Q_J^{\ell,\ell}(\sigma, \sigma'; \tau, \tau') \,.$$

In this case the partition associated to $\mathcal{K}$ is made of singletons: $\mathcal{A} = (\{i\}, i \in [n])$. The system (1.5) is the continuous-time counterpart of the nonlinear Single-Site Dynamics introduced in [CS24b].

Example 1.2 (Mean-Field Exchange Dynamics)**.** If $\mathcal{K}(i,j) = 1/n$ for all $i, j \in [n]$, the proposed exchanges are uniformly chosen between all pairs of sites. The collision kernel is given by

$$\mathcal{Q}_J^{\mathcal{K}}(\sigma, \sigma'; \tau, \tau') = \frac{1}{n^2} \sum_{\ell,k=1}^{n} Q_J^{\ell,k}(\sigma, \sigma'; \tau, \tau') \,.$$

In this case the partition $\mathcal{A}$ consists of a single component $\mathcal{A} = \{[n]\}$, and the process (1.5) is a nonlinear analogue of a mean-field Kawasaki dynamics.

The two examples discussed above are both covered by the general class of examples defined as follows.

Definition 1.3 (Multi-component Mean-Field Exchange Dynamics)**.** Fix an arbitrary partition $\mathcal{A}$ of $[n]$, and define the matrix $\mathcal{K} = \mathcal{K}_{\mathcal{A}}$ by

$$\mathcal{K}_{\mathcal{A}}(\ell, k) = \begin{cases} \frac{1}{|A|}, & \text{if } \ell, k \in A \text{ for some } A \in \mathcal{A}, \\ 0, & \text{otherwise,} \end{cases} \tag{1.7}$$

In this case, the nonlinear dynamics associated to the kernel $\mathcal{Q}_J^{\mathcal{K}}$ in (1.3) is referred to as the Nonlinear Multi-component Mean-Field Exchange Dynamics.

Definition 1.4. Given a partition $\mathcal{A}$ of $[n]$, and $p \in \mathcal{P}(\Omega)$, we define the $\mathcal{A}$-*magnetization profile* as the vector

$$m_{\mathcal{A}}(p) = (m(p, A)\,,\ A \in \mathcal{A})\,, \qquad m(p, A) = \frac{1}{|A|} \sum_{\ell \in A} p[\sigma_\ell]\,,$$

where $p[\sigma_\ell]$ is the expected value of the spin at site ℓ according to p. The vector $m_{\mathcal{A}}(p)$ is in $[-1, 1]^{\mathcal{A}}$, and we say that p is $\mathcal{A}$-*regular* if $m_{\mathcal{A}}(p) \in (-1, 1)^{\mathcal{A}}$, that is if $|m(p, A)| \neq 1$ for all $A \in \mathcal{A}$. We write $\mathcal{P}_{\mathcal{A}}(\Omega)$ for the set of $p \in \mathcal{P}(\Omega)$ that are $\mathcal{A}$-regular.

We note that in Example 1.1, specifying the vector $m_{\mathcal{A}}(p) \in [-1, 1]^n$ is equivalent to giving all the single-site marginals of p, while in Example 1.2, $m_{\mathcal{A}}(p) \in [-1, 1]$ is just the global mean magnetization of p.

The fist result is the following general convergence theorem, which summarizes the main properties of our nonlinear dynamical system. We use the notation $\Gamma_{\mathcal{A}} \subset \mathbb{R}^n$ for the set of external fields that are constant over each element of the partition $\mathcal{A}$:

$$\Gamma_{\mathcal{A}} = \{h \in \mathbb{R}^n :\ h_i = h_j,\ \forall i, j \in A\,,\ \forall A \in \mathcal{A}\}.$$

It can be shown that for each $p \in \mathcal{P}_{\mathcal{A}}(\Omega)$ there is a unique vector $h \in \Gamma_{\mathcal{A}}$ depending only on J and $m_{\mathcal{A}}(p)$, such that $m_{\mathcal{A}}(p) = m_{\mathcal{A}}(\mu_{J,h})$; see Lemma 1.14.

Theorem 1.5. *Fix a symmetric matrix $J \in \mathbb{R}^{n\times n}$ and a symmetric stochastic $n \times n$ matrix $\mathcal{K}$. Let $\mathcal{A}$ denote the partition of $[n]$ induced by the irreducible components of $\mathcal{K}$. A probability $p \in \mathcal{P}_{\mathcal{A}}(\Omega)$ is a stationary state if and only if $p = \mu_{J,h}$ for some $h \in \Gamma_{\mathcal{A}}$. Moreover, the expected values $m(p_t, A)$, $A \in \mathcal{A}$, are independent of $t \geq 0$. Finally, for any initial state $p \in \mathcal{P}_{\mathcal{A}}(\Omega)$, one has the convergence*

$$\|p_t - \mu_{J,h}\|_{\mathrm{TV}} \ \longrightarrow\ 0\,, \qquad t \to \infty\,,$$

where $\|\cdot\|_{\mathrm{TV}}$ denotes the total variation distance and h is the unique vector $h \in \Gamma_{\mathcal{A}}$ such that $m_{\mathcal{A}}(p) = m_{\mathcal{A}}(\mu_{J,h})$.

Remark 1.6. We note that the requirement that p be $\mathcal{A}$-regular in Theorem 1.5 is not essential, as it can be removed if one allows the external fields h to take the value $\pm\infty$.

Theorem 1.5 establishes convergence to stationarity, but it does not provide information on the speed of convergence. The next results address this question under a high–temperature assumption on the interaction matrix J.

1.1.2 Entropy decay

We restrict to the case of Multi-component Mean Field Dynamics, that is we take $\mathcal{K} = \mathcal{K}_{\mathcal{A}}$ as in Definition 1.3 with an arbitrary choice of the partition $\mathcal{A}$.
Recall the definition of relative entropy given in (0.6).

Our main quantitative result is formulated as follows.

Theorem 1.7. *Let $J \in \mathbb{R}^{n\times n}$ be a nonnegative definite matrix with largest eigenvalue $\lambda(J) < \frac{1}{2}$, and let $\overline{J} = \max_{i\in[n]} \sum_{j\in[n]} |J_{i,j}|$. Then, for any partition $\mathcal{A}$ of $[n]$, the nonlinear system (1.5) defined by J and $\mathcal{K}$, with $\mathcal{K} = \mathcal{K}_{\mathcal{A}}$ given in (1.7), satisfies*

$$H(p_t \,|\, \mu_{J,h}) \le H(p \,|\, \mu_{J,h})\, e^{-r\,t}\,, \qquad t \ge 0, \tag{1.8}$$

for any $p \in \mathcal{P}_{\mathcal{A}}(\Omega)$, where h is the unique vector $h \in \Gamma_{\mathcal{A}}$ such that $m_{\mathcal{A}}(p) = m_{\mathcal{A}}(\mu_{J,h})$, and the positive constant α satisfies

$$r \ge \tfrac{1}{8n}(1 - 2\lambda(J))^2\, e^{-16\overline{J}}\,.$$

Remark 1.8. We note that the estimate (1.8) implies exponential decay in total variation as well. Indeed, by Pinsker's inequality, $\|p_t - \mu_{J,h}\|_{\mathrm{TV}}^2 \le \frac{1}{2} H(p_t \,|\, \mu_{J,h})$. Using $H(p \,|\, \mu_{J,h}) \le \max_{\sigma\in\Omega} \log(1/\mu_{J,h}(\sigma))$, which can be bounded by Cn, where $C = \lambda(J) + 2\overline{h} + \log(2)$ and $\overline{h} = \max_{i\in[n]} |h_i|$, it follows that, in the setting of Theorem 1.7, for all $p \in \mathcal{P}_{\mathcal{A}}(\Omega)$, one has

$$\|p_t - \mu_{J,h}\|_{\mathrm{TV}} \le \sqrt{\frac{C\,n}{2}}\, e^{-r\,t/2}\,. \tag{1.9}$$

In particular, Theorem 1.7 is a quantitative version of Theorem 1.5 under the additional assumption on the spectral radius of J, which can be regarded as a high temperature requirement. It is worthy of note that the rate α is bounded below independently of the choice of partition $\mathcal{A}$. In the single-site case of Example 1.1 this reproduces essentially the same rate that was obtained, under a more restrictive high temperature constraint, for the discrete-time version of the single-site model in [CS24b].

By comparison with the non-interacting case $J = 0$, we note that the rate α has the optimal dependence with respect to the dimension n, up to a constant factor. For instance, in the single-site case, one can use a coupon collector argument to show that if $J = 0$ and the initial distribution p is all $+1$ with probability $1/2$ and all -1 with probability $1/2$, then the left hand side in (1.9) is at least $c\,e^{-\frac{t}{c\,n}}$, $t \ge 0$, for some absolute constant $c > 0$; see [CS24b].

The estimate in Theorem 1.7 is equivalent to a certain nonlinear logarithmic Sobolev inequality, which provides a lower bound on the relative entropy dissipation in terms of the relative entropy; see Section 1.5.1 below. Rather than establishing such a bound directly, to prove Theorem 1.7 we follow the strategy introduced by Mark Kac in his seminal 1956 paper [Kac56], that is we approximate the nonlinear evolution of one particle by a linear mean-field type Markov process involving a large number of particles.

1.1.3 Kac's program

We consider a system of $N \in \mathbb{N}$ *particles* $\{\sigma(i), i = 1, \dots, N\}$, where each $\sigma(i) = (\sigma_1(i), \dots, \sigma_n(i))$ is a spin configuration in $\Omega = \{-1, +1\}^n$. In words, the process is described as follows. Independently, each ordered pair of particles (i, j), including $i = j$, is

updated at the arrival times of a Poisson process with rate $1/N$. When an update occurs at particles (i, j), a site ℓ is picked uniformly at random, a site k is picked according to the transport kernel $\mathcal{K}(\ell, k)$, and then with probability $\hat{p}_J(\ell, k \mid \sigma(i), \sigma(j))$ the ℓ-th spin of the i-th particle is exchanged with the k-th spin of the j-th particle, while with probability $1 - \hat{p}_J(\ell, k \mid \sigma(i), \sigma(j))$ the particle configuration is left unchanged. Using the notation (1.1), the exchange can be written as

$$(\sigma(i), \sigma(j)) \mapsto (S_\ell(\sigma(i), \sigma(j)_k)\,,\, S_k(\sigma(j), \sigma(i)_\ell),$$

with the understanding that when $i = j$ this is replaced by $\sigma(i) \mapsto \sigma(i)^{\ell,k}$, where $\sigma(i)^{\ell,k}$ denotes the configuration $\sigma(i)$ after the exchange $(\sigma_\ell(i), \sigma_k(i)) \mapsto (\sigma_k(i), \sigma_\ell(i))$. Here the probability $\hat{p}_J(\ell, k \mid \sigma(i), \sigma(j))$ is defined by

$$\hat{p}_J(\ell, k \mid \sigma(i), \sigma(j)) = \begin{cases} p_J(\ell, k \mid \sigma(i), \sigma(j)) & i \neq j \\ \frac{\mu_J(\sigma(i)^{\ell k})}{\mu_J(\sigma(i)) + \mu_J(\sigma(i)^{\ell k})} & i = j \end{cases} \tag{1.10}$$

where we use the notation (1.2). This defines a continuous-time Markov chain with state space $\Omega^N = \{-1, +1\}^{Nn}$, which is described by the infinitesimal generator

$$\mathcal{L}_N F(\sigma) = \frac{1}{Nn} \sum_{i,j=1}^{N} \sum_{\ell,k \in [n]} \mathcal{K}(\ell, k)\, \hat{p}_J(\ell, k \mid \sigma(i), \sigma(j))\, \nabla^{i,\ell;j,k} F(\sigma), \tag{1.11}$$

where $\sigma \in \Omega^N$, $F : \Omega^N \to \mathbb{R}$, and the gradient operator $\nabla^{i,\ell;j,k}$ is defined by

$$\nabla^{i,\ell;j,k} F(\sigma) := F(\sigma^{i,\ell;j,k}) - F(\sigma), \quad \text{with } \sigma^{i,\ell;j,k}(u) := \begin{cases} \sigma(u) & u \neq i, j, \\ S_\ell(\sigma(i), \sigma(j)_k) & u = i, \\ S_k(\sigma(j), \sigma(i)_\ell) & u = j, \end{cases} \tag{1.12}$$

with the convention that, when $i = j$, then $\sigma^{i,\ell;i,k}(u) = \sigma(u)$ if $u \neq i$, and $\sigma^{i,\ell;i,k}(i) = \sigma(i)^{\ell,k}$.

We may picture the system as a spin exchange dynamics in a $N \times n$ grid, with exchanges occurring between sites (i, ℓ) and (j, k). By construction, the Markov chain is reversible with respect to the product measure $\mu_J^{\otimes N}$. Since only pairs (ℓ, k) belonging to sites within the same equivalence class are exchanged, to turn the Markov chain into an irreducible process we need to impose a constraint of fixed magnetization in each block $B_{N,A} := \{(i, \ell)\,,\ i \in [N], \ell \in A\}$, $A \in \mathcal{A}$. To this end, we need the following definitions.

Definition 1.9. Given a partition $\mathcal{A}$ of $[n]$, we write $D_{\mathcal{A},N}$ for the set of all vectors $\alpha = (\alpha(A),\ A \in \mathcal{A}) \in [0, 1]^{\mathcal{A}}$ such that $N|A|\alpha(A)$ is a nonnegative integer. Given $\mathcal{A}$, $D_{\mathcal{A},N}$ is called the set of $\mathcal{A}$-*density profiles*. Given $\alpha \in D_{\mathcal{A},N}$, we define

$$\Omega_{N,\alpha} = \left\{\sigma \in \Omega^N :\ \textstyle\sum_{i \in [N]} \sum_{\ell \in A} \sigma_\ell(i) = (2\alpha(A) - 1)N|A|,\ \ \forall A \in \mathcal{A}\right\}. \tag{1.13}$$

It follows that the Markov chain with generator (1.11) is irreducible in the state space

$\Omega_{N,\alpha}$, for every choice of $\alpha \in D_{\mathcal{A},N}$. Moreover, the probability measure

$$\mu_{N,J,\alpha} := \mu_J^{\otimes N}(\cdot \mid \Omega_{N,\alpha}) \tag{1.14}$$

obtained as tensor product of μ_J restricted to $\Omega_{N,\alpha}$ is the associated unique invariant distribution. We refer to $\mu_{N,J,\alpha}$ as the multi-canonical Ising measure.

Our main result concerning the particle system is the following modified log-Sobolev inequality (MLSI), with bounds that are uniform in N and in the density profile $\alpha \in D_{\mathcal{A},N}$. To introduce the relevant quantities, observe that the Dirichlet form associated to the reversible pair $(\mathcal{L}_N, \mu_{N,J,\alpha})$ is given by

$$\mathcal{E}_{N,J,\alpha}(F,G) := \frac{1}{2Nn} \sum_{i,j=1}^{N} \sum_{\ell,k=1}^{n} \mathcal{K}(\ell,k) \mu_{N,J,\alpha} \left[\hat{p}_J(\ell,k \mid \sigma(i),\sigma(j)) \nabla^{i,\ell;j,k} F \, \nabla^{i,\ell;j,k} G \right], \tag{1.15}$$

for functions $F, G : \Omega_{N,\alpha} \to \mathbb{R}$. Recall that the *entropy* of a function $F \geq 0$ with respect to a probability measure μ is given by $\operatorname{Ent}_\mu F := \mu[F \log F] - \mu[F] \log \mu[F]$.

Theorem 1.10 (MLSI for the particle system)**.** *Let $J \in \mathbb{R}^{n \times n}$ be a nonnegative definite matrix with largest eigenvalue $\lambda(J) < 1/2$, and let $\overline{J} = \max_{i \in [n]} \sum_{j \in [n]} |J_{i,j}|$. Then, for any partition $\mathcal{A}$ of $[n]$, for any $N \geq 2$, and $\alpha \in \mathcal{D}_{\mathcal{A},N}$, the Markov chain with Dirichlet form (1.15) defined by J and $\mathcal{K} = \mathcal{K}_{\mathcal{A}}$ given in (1.7), satisfies*

$$\mathcal{E}_{N,J,\alpha}(F, \log F) \geq r \operatorname{Ent}_{\mu_{N,J,\alpha}} F,$$

for all $F \geq 0$, with constant α such that

$$r \geq \tfrac{1}{4n}(1 - 2\lambda(J))^2 \, e^{-16\overline{J}}. \tag{1.16}$$

It is well known that a MLSI of the form (1.16) is equivalent to an exponential decay statement for the relative entropy, with decay rate given by the same constant α. That is, for any initial distribution ν on $\Omega_{N,\alpha}$, the distribution $\nu_t := \nu e^{t\mathcal{L}_N}$ of the particle system at time t satisfies

$$H_N(\nu_t \mid \mu_{N,J,\alpha}) \leq H_N(\nu \mid \mu_{N,J,\alpha}) \, e^{-r\,t}, \qquad t \geq 0, \tag{1.17}$$

where H_N denotes the relative entropy for distributions on Ω^N. The key feature of the bound in Theorem 1.10 is the independence on N and $\alpha \in \mathcal{D}_{\mathcal{A},N}$. This, together with so-called Entropic Chaos and Fisher Chaos results for our system, will allow us to transfer the result (1.17) on the linear particle system to the desired entropy decay estimate (1.8) for the nonlinear system, as stated in Theorem 1.7; see Section 1.5.

1.1.4 Wasserstein decay

The results presented in this subsection concern the multi-mean-field case described in Definition 1.3, namely the choice of the transportation kernel $\mathcal{K}_{\mathcal{A}}$ associated with a given partition $\mathcal{A}$ of the set $[n]$. Recall the definition of Wasserstein metric (0.8).

Theorem 1.11. *Fix a partition $\mathcal{A}$ of the set $[n]$. Let $J \in \mathbb{R}^{n\times n}$ be a matrix such that $\overline{J} = \max_{i\in[n]} \sum_{j\in[n]} |J_{i,j}| < \frac{1}{10}$. Then the nonlinear system (1.5) defined by J and $\mathcal{K}_{\mathcal{A}}$ satisfies*

$$W_1(p_t, \mu_{J,h}) \leq W_1(p, \mu_{J,h}) e^{-\varepsilon t/n}$$

for any initial state $p \in \mathcal{P}(\Omega)$, where $h \in \mathbb{R}^n$ is the unique external field in $\Gamma_{\mathcal{A}}$ such that $m(\mu_{J,h}, \mathcal{A}) = m(p, \mathcal{A})$, and $\varepsilon = \frac{1}{2}(1 - 10\overline{J})$.

Remark 1.12. Theorem 1.11 implies exponential convergence to equilibrium in total variation for the multi-mean-field dynamics. More precisely, let p_t be the solution of the nonlinear system (1.5) defined by the interaction matrix J and the kernel $\mathcal{K}_{\mathcal{A}}$ associated with a given partition $\mathcal{A}$ of the set $[n]$. Then

$$\|p_t - \mu_{J,h(p)}\|_{TV} \leq n e^{-\varepsilon t/n},$$

where $\varepsilon = \frac{1}{2}(1 - 10\overline{J})$, under the restriction $\overline{J} < \frac{1}{10}$, and $h(p) \in \mathbb{R}^n$ is the unique external field in $\Gamma_{\mathcal{A}}$ such that $m(\mu_{J,h(p)}, \mathcal{A}) = m(p, \mathcal{A})$.

This follows from the coupling characterization of the total variation distance: for all $p, q \in \mathcal{P}(\Omega)$,

$$\|p - q\|_{TV} = \inf_{\pi\in\Gamma(p,q)} \mathbb{P}_{\sigma,\sigma'\sim\pi}(\sigma \neq \sigma'),$$

which is a standard identity (see, for instance, [LP17]). Denote $\mu := \mu_{J,h(p)}$. Using Theorem 1.11, we obtain

$$\begin{aligned}
\|p_t - \mu\|_{TV} &= \inf_{\pi\in\Gamma(p_t,\mu)} \sum_{\sigma,\sigma'\in\Omega} \pi(\sigma,\sigma')\mathbf{1}(\sigma \neq \sigma') \\
&\leq \inf_{\pi\in\Gamma(p_t,\mu)} \sum_{\sigma,\sigma'\in\Omega} \pi(\sigma,\sigma')|D_{\sigma,\sigma'}| \\
&= W_1(p_t, \mu) \\
&\leq e^{-\varepsilon t/n} W_1(p, \mu) \leq n e^{-\varepsilon t/n}.
\end{aligned}$$

Here we used the elementary bounds $\mathbf{1}(\sigma \neq \sigma') \leq |D_{\sigma,\sigma'}| \leq n$.

The result of Theorem 1.11 does not depend on the choice of the partition of the set $[n]$. Moreover, restricting to the case of single-site dynamics, i.e., $\mathcal{A} = \{\{i\}, i \in [n]\}$, yields the correct order of decay, as already noted in Remark 1.8, and improves the condition on J obtained in [CS24b].

However, the corresponding high-temperature assumption is more restrictive than the one required in Theorem 1.7, since $\lambda(J) \leq \overline{J} < 1/10$ here, whereas Theorem 1.7 only requires $\lambda(J) < 1/2$.

Therefore, while this approach yields the correct exponential rate in terms of the dimension n, it does not cover a wide range of interactions. We stress that the strategy used to prove Theorem 1.7 is considerably more involved: it relies on Kac's program together with a uniform modified logarithmic Sobolev inequality for the associated particle system (see Theorem 1.10). By contrast, the Wasserstein contraction argument follows from a direct construction of a suitable coupling.

1.1.5 Plan of the chapter

General structural properties of the dynamical system are discussed in Section 1.2, which also contains the proof of Theorem 1.5. The proof of Theorem 1.11 is in Section **??**. The analysis of the associated Kac particle system, together with the proof of Theorem 1.10, is presented in Section 1.4. Finally, in Section 1.5 we show how to transfer these results from the Kac particle system to the nonlinear equation, thereby completing the proof of Theorem 1.7. We conclude with some open problems about establishing exponential decay to equilibrium without the high temperature assumption.

1.2 Structural properties of the nonlinear equation

The statements in this section hold for any integer $n \in \mathbb{N}$, for any fixed choice of the symmetric matrix $J \in \mathbb{R}^{n\times n}$ and the symmetric stochastic $n \times n$ matrix $\mathcal{K}$.

1.2.1 Preliminary facts

Recall the notation $\mathcal{A}$ for the partition of $[n]$ induced by the irreducible components of $\mathcal{K}$ and the set $\Gamma_{\mathcal{A}}$ of external fields associated to $\mathcal{A}$.

Lemma 1.13. *The kernel $\mathcal{Q}_J^{\mathcal{K}}$ defined by (1.3) satisfies*

$$\mu_{J,h}(\sigma)\mu_{J,h}(\sigma')Q_J^{\mathcal{K}}(\sigma,\sigma';\tau,\tau') = \mu_{J,h}(\tau)\mu_{J,h}(\tau')Q_J^{\mathcal{K}}(\tau,\tau';\sigma,\sigma'), \tag{1.18}$$

for all $h \in \Gamma_{\mathcal{A}}$, and $\sigma,\sigma',\tau,\tau' \in \Omega$.

Proof. Let $p_{J,h}(\ell,k|\sigma,\sigma')$ denote the quantity defined in (1.2) with the difference that μ_J is replaced by $\mu_{J,h}$, with $h \in \mathbb{R}^n$. By the definition of $\Gamma_{\mathcal{A}}$, it follows that, whenever ℓ,k belong to the same irreducible class,

$$p_J(\ell,k|\sigma,\sigma') = p_{J,h}(\ell,k|\sigma,\sigma')\,, \qquad h \in \Gamma_{\mathcal{A}}.$$

In particular, if $(\tau,\tau') = (S_\ell(\sigma,\sigma'_k),\ S_k(\sigma',\sigma_\ell))$, and $h \in \Gamma_{\mathcal{A}}$, then

$$\mu_{J,h}(\sigma)\mu_{J,h}(\sigma')p_J(\ell,k|\sigma,\sigma') = \mu_{J,h}(\tau)\mu_{J,h}(\tau')p_J(\ell,k|\tau,\tau').$$

Since only pairs (ℓ,k) such that $\mathcal{K}(\ell,k) > 0$ enter the definition (1.3) of $\mathcal{Q}_J^{\mathcal{K}}$, this shows that the kernel $\mathcal{Q}_J^{\mathcal{K}}$ satisfies the claimed detailed balance (1.18). □

Recall the Definition 1.4 of the magnetization profile $m_{\mathcal{A}}(p)$.

Lemma 1.14. *Fix a partition $\mathcal{A}$ of $[n]$. For any vector $v = (v_A,\, A \in \mathcal{A})$, with $v_A \in (-1,1)$, for all $A \in \mathcal{A}$, there exists a unique choice of $h \in \Gamma_{\mathcal{A}}$ such that $m_{\mathcal{A}}(\mu_{J,h}) = v$.*

Proof. We adapt the argument from Lemma 1 in [BE19]. We need to show that the map

$$h \mapsto m_{\mathcal{A}}(\mu_{J,h})$$

is a bijection between $\Gamma_{\mathcal{A}}$ and $(-1,1)^{\mathcal{A}}$. Using the change of variables $\theta_A := \sum_{i\in A} h_i$, $A \in \mathcal{A}$, this is equivalent to showing that the map

$$\theta \in \mathbb{R}^{\mathcal{A}} \quad \longmapsto \quad f(\theta) := \nabla\varphi(\theta) \in (-1,1)^{\mathcal{A}}$$

is a bijection, where

$$\varphi(\theta) = \log\left(\sum_{\sigma\in\Omega} e^{\sum_{A\in\mathcal{A}} \theta_A s_A(\sigma)} \mu_J(\sigma)\right), \qquad s_A(\sigma) := \frac{1}{|A|}\sum_{i\in A}\sigma_i\,.$$

The Jacobian $\nabla f = \nabla^2\varphi$ coincides with the covariance matrix of the random vector $(s_A)_{A\in\mathcal{A}}$ when σ is distributed according to $\mu_{J,h}$, which is easily seen to be strictly positive definite for all $\theta \in \mathbb{R}^{\mathcal{A}}$. Therefore, φ is strictly convex and $f = \nabla\varphi$ is injective. To show that it is also surjective, it is sufficient to show that the image of f, $\mathrm{Im}(f)$, satisfies

$$\partial\mathrm{Im}(f) \subset \partial(-1,1)^{\mathcal{A}}.$$

Indeed, suppose by contradiction that there exists $y \in (-1,1)^{\mathcal{A}} \setminus \mathrm{Im}(f)$, and consider the line segment $[f(0), y]$. Since $\mathrm{Im}(f)$ is open, this segment must intersect $\partial\mathrm{Im}(f)$ at some point z. If $\partial\mathrm{Im}(f) \subset \partial(-1,1)^{\mathcal{A}}$, then $z \in \partial(-1,1)^{\mathcal{A}}$, which contradicts the fact that $(-1,1)^{\mathcal{A}}$ is convex.

To prove $\partial\mathrm{Im}(f) \subset \partial(-1,1)^{\mathcal{A}}$, take $w \in \partial\mathrm{Im}(f)$. Then there exists a sequence $(\theta_k) \subset \mathbb{R}^{\mathcal{A}}$ such that $\lim_{k\to\infty} f(\theta_k) = w$. By continuity of f, the sequence (θ_k) must be unbounded, since otherwise one could extract a converging subsequence $\theta'_k \to \theta \in \mathbb{R}^{\mathcal{A}}$ which would imply that $w = f(\theta) \in \mathrm{Im}(f)$. Therefore, there exists a subsequence (θ'_k) and some $A \in \mathcal{A}$ such that $(\theta'_k)_A \to +\infty$ or $(\theta'_k)_A \to -\infty$. It follows that

$$w_A = \lim_{j\to\infty} f(\theta'_k)_A \in \{-1, 1\},$$

and therefore $w \in \partial(-1,1)^{\mathcal{A}}$. □

1.2.2 Entropy and stationary states

A distribution $p \in \mathcal{P}(\Omega)$ is called *stationary* if $p \circ p = p$, where the collision product "$\circ$" is defined in (1.6). Equivalently, p is stationary iff $p_t = p$, for all $t \geq 0$. It follows from the

reversibility in Lemma 1.13 that $\mu_{J,h} \circ \mu_{J,h} = \mu_{J,h}$, for any $h \in \Gamma_{\mathcal{A}}$, and thus all measures $\mu_{J,h}$, $h \in \Gamma_{\mathcal{A}}$, are stationary. To characterize all stationary states we use the relative entropy functional.

The following statement is a version of Boltzmann's H-theorem from kinetic theory.

Lemma 1.15. *Let p_t denote the solution of (1.5) with initial datum $p \in \mathcal{P}(\Omega)$, and let $\mu = \mu_{J,h}$ for some arbitrary, fixed $h \in \Gamma_{\mathcal{A}}$. Then,*

$$\frac{d}{dt} H(p_t \mid \mu) = -\mathcal{D}_\mu(f_t, f_t), \qquad t \geq 0, \tag{1.19}$$

where $f_t := p_t/\mu$, and the entropy dissipation functional $\mathcal{D}_\mu$ is defined by

$$\mathcal{D}_\mu(f,f) = \frac{1}{4} \sum_{\sigma,\sigma',\tau,\tau'} \mu(\tau)\mu(\tau') \mathcal{Q}_J^{\mathcal{K}}(\tau,\tau';\sigma,\sigma') \left(f(\tau)f(\tau') - f(\sigma)f(\sigma')\right) \log\left(\frac{f(\tau)f(\tau')}{f(\sigma)f(\sigma')}\right).$$

In particular, $H(p_t \mid \mu)$ is non-increasing in $t \geq 0$. Moreover, $p \in \mathcal{P}(\Omega)$ is stationary iff the density $f = p/\mu$ satisfies $f(\tau)f(\tau') = f(\sigma)f(\sigma')$ whenever $\mathcal{Q}_J^{\mathcal{K}}(\tau,\tau';\sigma,\sigma') > 0$.

Proof. For any $p \in \mathcal{P}(\Omega)$, setting $f = p/\mu$, and using Lemma 1.13, for any $\tau \in \Omega$,

$$p \circ p\,(\tau) - p(\tau) = \sum_{\sigma,\sigma',\tau'} \mu(\sigma)\mu(\sigma') \mathcal{Q}_J^{\mathcal{K}}(\sigma,\sigma';\tau,\tau') \left(f(\sigma)f(\sigma') - f(\tau)f(\tau')\right). \tag{1.20}$$

Therefore, differentiating the relative entropy,

$$\begin{aligned} \frac{d}{dt} H(p_t \mid \mu) &= \sum_{\tau} \left[\frac{d}{dt}\, p_t(\tau)\right] \log \frac{p_t(\tau)}{\mu(\tau)} \\ &= -\sum_{\sigma,\sigma',\tau,\tau'} \mu(\tau)\mu(\tau') \mathcal{Q}_J^{\mathcal{K}}(\tau,\tau';\sigma,\sigma') \left(f_t(\tau)f_t(\tau') - f_t(\sigma)f_t(\sigma')\right) \log f_t(\tau). \end{aligned}$$

Symmetrizing first with respect to $\tau \leftrightarrow \tau'$ and then with respect to $(\tau,\tau') \leftrightarrow (\sigma,\sigma')$, Lemma 1.13 shows that the above expression can be rewritten as (2.9). Note that $\mathcal{D}_\mu(f) \geq 0$, for any function $f : \Omega \to \mathbb{R}_+$, since every summand in its definition is nonnegative. Now, if p is stationary, then the density $f = p/\mu$ satisfies $\mathcal{D}_\mu(f) = 0$, and therefore, by positivity of μ, $f(\tau)f(\tau') = f(\sigma)f(\sigma')$ whenever $\mathcal{Q}_J^{\mathcal{K}}(\tau,\tau';\sigma,\sigma') > 0$. Conversely, if $f(\tau)f(\tau') = f(\sigma)f(\sigma')$ whenever $\mathcal{Q}_J^{\mathcal{K}}(\tau,\tau';\sigma,\sigma') > 0$, then (1.20) shows that $p \circ p = p$. □

We turn to a characterization of the stationary states. Let $\mathcal{P}_+(\Omega)$ denote the set of positive distributions, that is the set of $p \in \mathcal{P}(\Omega)$ such that $p(\sigma) > 0$ for all $\sigma \in \Omega$. Clearly, $\mu_{J,h} \in \mathcal{P}_+(\Omega)$ for any $h \in \mathbb{R}^n$.

Lemma 1.16 (Stationary measures)**.** *A probability measure $p \in \mathcal{P}_+(\Omega)$ is stationary if and only if $p = \mu_{J,h}$ for some $h \in \Gamma_{\mathcal{A}}$.*

Proof. We already know that $\mu_{J,h}$ is stationary, for any $h \in \Gamma_{\mathcal{A}}$. It remains to show that these are the only positive stationary distributions. Fix $\mu = \mu_{J,h}$ for some $h \in \Gamma_{\mathcal{A}}$, and let

$p \in \mathcal{P}_+(\Omega)$ be a stationary measure. From Lemma 1.15, $f = p/\mu$ satisfies

$$f(\sigma)f(\sigma') = f(S_\ell(\sigma, \sigma'_k))f(S_k(\sigma', \sigma_\ell)), \tag{1.21}$$

for all $\ell, k \in [n]$ such that $\mathcal{K}(\ell, k) > 0$, for all $\sigma, \sigma' \in \Omega$. For any ℓ, fix some $k = k(\ell) \in [n]$ such that $\mathcal{K}(\ell, k) > 0$. For any $\sigma \in \Omega$ and any $j \in [n]$, let $\sigma[j] \in \Omega$ denote the configuration equal to $+1$ on all sites $\{1, \dots, j\}$ and equal to σ on $\{j+1, \dots, n\}$. Setting $\sigma[0] = \sigma$, we have:

$$f(\sigma) = f(+) \prod_{j=1}^{n} \frac{f(\sigma[j-1])}{f(\sigma[j])},$$

where $\sigma[n] = +$ is the all $+$ configuration. Define $\varphi_j(a) := \frac{f(+^{k(j),a})}{f(+)}$, where $+^{k,a}$ denotes the configuration equal to a at site k and $+1$ elsewhere. Taking $\sigma = \sigma[j-1]$, $\sigma' = +$, from (1.21) with $j = \ell$ and $k = k(j)$,

$$\frac{f(\sigma[j-1])}{f(\sigma[j])} = \frac{f(+^{k(j),\sigma_j})}{f(+)} = \varphi_j(\sigma_j).$$

Since φ_j is a positive function on $\{-1, +1\}$, we can write $\varphi_j(a) = c_j e^{h_j a}$, $a = \pm 1$, for some constants $c_j > 0, h_j \in \mathbb{R}$. It follows that

$$f(\sigma) = C \exp\left(\textstyle\sum_{j\in[n]} h_j \sigma_j\right),$$

for some $h \in \mathbb{R}^n$, where $C > 0$ is a constant. It remains to prove that $h \in \Gamma_{\mathcal{A}}$. To see this, note that by (1.21), whenever $\mathcal{K}(\ell, k) > 0$ one has

$$(h_\ell - h_k)(\sigma'_k - \sigma_\ell) = 0\,,$$

for all σ, σ'. It follows that $h_\ell = h_k$ if $\mathcal{K}(\ell, k) > 0$. Therefore $h \in \Gamma_{\mathcal{A}}$. □

1.2.3 Conservation laws

Recall the Definition 1.4 of the magnetization profile $m_{\mathcal{A}}(p) = (m(p, A), A \in \mathcal{A})$.

Lemma 1.17. *For all $p, q \in \mathcal{P}(\Omega)$,*

$$m_{\mathcal{A}}(p \circ q) = \tfrac{1}{2}\, m_{\mathcal{A}}(p) + \tfrac{1}{2}\, m_{\mathcal{A}}(q).$$

In particular, if p_t is the solution of (1.5) with initial datum p, then $\frac{d}{dt} m_{\mathcal{A}}(p_t) = 0$

Proof. By definition,

$$m(p \circ q, A) = \frac{1}{|A|} \sum_{i \in A} p \circ q\,[\tau_i],$$

where

$$p \circ q\,[\tau_i] = \frac{1}{2} \sum_{\sigma,\sigma'} p(\sigma) q(\sigma') \sum_{\tau,\tau'} \left(\mathcal{Q}^{\mathcal{K}}_J(\sigma, \sigma'; \tau, \tau') + \mathcal{Q}^{\mathcal{K}}_J(\sigma', \sigma; \tau, \tau')\right) \tau_i.$$

Moreover,

$$\sum_{\tau,\tau'} \mathcal{Q}_J^{\mathcal{K}}(\sigma,\sigma';\tau,\tau')\,\tau_i = \sigma_i + \frac{1}{n}\sum_{k=1}^{n} \mathcal{K}(i,k) p_J(i,k \mid \sigma,\sigma')(\sigma'_k - \sigma_i).$$

Therefore,

$$\frac{1}{2}\sum_{\tau,\tau'} \Big(\mathcal{Q}_J^{\mathcal{K}}(\sigma,\sigma';\tau,\tau') + \mathcal{Q}_J^{\mathcal{K}}(\sigma',\sigma;\tau,\tau')\Big)\,\tau_i = \frac{1}{2}(\sigma_i + \sigma'_i) + \frac{1}{2n}\sum_{k=1}^{n} R_{i,k}(\sigma,\sigma')$$

where

$$R_{i,k}(\sigma,\sigma') = \mathcal{K}(i,k)\left[p_J(i,k \mid \sigma,\sigma')(\sigma'_k - \sigma_i) - p_J(k,i \mid \sigma,\sigma')(\sigma'_i - \sigma_k)\right],$$

and we have used the symmetries $p_J(i,k \mid \sigma,\sigma') = p_J(k,i \mid \sigma',\sigma)$. Since $\mathcal{K}(i,k) = \mathcal{K}(k,i)$, $R_{i,k}(\sigma,\sigma')$ is antisymmetric in the exchange $i \leftrightarrow k$. Therefore, for any fixed $A \subset [n]$,

$$\sum_{i,k\in A} R_{i,k}(\sigma,\sigma') = 0.$$

In particular, using that $\mathcal{K}(i,k) = 0$ when i,k do not belong to the same irreducible component, this proves that for any fixed $A \in \mathcal{A}$,

$$\frac{1}{|A|}\sum_{i\in A} p \circ q\,[\tau_i] = \frac{1}{2}m(p,A) + \frac{1}{2}m(q,A)\,.$$

□

1.2.4 Proof of Theorem 1.5

The key to the proof of Theorem 1.5 is the following statement about positivity of solutions, that we refer to as the irreducibility lemma. Recall the notation $\mathcal{P}_{\mathcal{A}}(\Omega)$ for the set of $\mathcal{A}$-regular distributions.

Lemma 1.18 (Irreducibility)**.** *For any $p \in \mathcal{P}_{\mathcal{A}}(\Omega)$, there exists $t_0 > 0$ and $\varepsilon_0 > 0$ such that for all $t \geq t_0$ the solution p_t of (1.5) satisfies*

$$p_t(\tau) \geq \varepsilon_0, \qquad \forall \tau \in \Omega. \tag{1.22}$$

The proof requires several steps and is given in the next subsections. Here we show how to prove Theorem 1.5 once we have Lemma 1.18. If $p \in \mathcal{P}_{\mathcal{A}}(\Omega)$, then Lemma 1.18 ensures that $p_t \in \mathcal{P}_+(\Omega)$ eventually. Therefore, if p is $\mathcal{A}$-regular and stationary, then it must be in $\mathcal{P}_+(\Omega)$, and therefore by Lemma 1.16, $p = \mu_{J,h}$ for some $h \in \Gamma_{\mathcal{A}}$. Next, we show that if p is $\mathcal{A}$-regular, then p_t converges as $t \to \infty$ to a stationary state ν. If this is the case, then by the above discussion $\nu = \mu_{J,h}$ for some $h \in \Gamma_{\mathcal{A}}$, and, using Lemma 1.14 and Lemma 1.17, h is the unique element of $\Gamma_{\mathcal{A}}$ such that $m_{\mathcal{A}}(p) = m_{\mathcal{A}}(\mu_{J,h})$. Thus, the

proof of Theorem 1.5 is complete once we establish the following statement.

Lemma 1.19. *If $p \in \mathcal{P}_{\mathcal{A}}(\Omega)$, then there exists a stationary state ν such that $p_t \to \nu$, as $t \to \infty$.*

Proof. Fix $\mu = \mu_{J,h}$, for some $h \in \Gamma_{\mathcal{A}}$, and $p \in \mathcal{P}_{\mathcal{A}}(\Omega)$, and let $S_t(p) = p_t$ denote the solution of (1.5). By compactness, there exists a diverging sequence of times $(t_k)_{k\in\mathbb{N}}$ and $\nu \in \mathcal{P}(\Omega)$ such that

$$S_{t_k}(p) \to \nu \quad \text{as} \quad k \to \infty.$$

By Lemma 1.15, it follows that

$$\lim_{t\to\infty} H(S_t(p) \mid \mu) = \inf_{t\geq 0} H(S_t(p) \mid \mu) = H(\nu \mid \mu). \tag{1.23}$$

Note that this holds for any $\mu = \mu_{J,h}$, $h \in \Gamma_{\mathcal{A}}$. If ν is stationary, then by Lemma 1.18 and Lemma 1.16 we know that $\nu = \mu_{J,h}$, for some $h \in \Gamma_{\mathcal{A}}$. Therefore, in (1.23) we can take $\mu = \nu$, which then implies that $H(S_t(p) \mid \nu) \to 0$, $t \to \infty$. In turn, e.g. by Pinsker's inequality, this implies the desired convergence $S_t(p) \to \nu$. Thus, it suffices to prove that ν is stationary.

Assume, for contradiction, that ν is not stationary. Then, by Lemma 1.15 one has that $f = \nu/\mu$ satisfies $\mathcal{D}_\mu(f) > 0$. Therefore, there exist $u > 0$ and $\delta > 0$ such that

$$H(S_u(\nu) \mid \mu) \leq H(\nu \mid \mu) - \delta u.$$

On the other hand, by the convergence $S_{t_k}(p) \to \nu$, and the continuity of $S_t(\cdot)$ and $H(\cdot \mid \mu)$, there exists $k \in \mathbb{N}$ such that $H(S_u(\nu) \mid \mu) \geq H(S_u(S_{t_k}(p)) \mid \mu) - \delta u/2$. Since $S_u(S_{t_k}(p)) = S_{u+t_k}(p)$, this would then imply

$$H(S_{u+t_k}(p) \mid \mu) \leq H(\nu \mid \mu) - \delta u/2,$$

which contradicts (1.23). This ends the proof. □

Thus, we have proved Theorem 1.5 assuming the validity of Lemma 1.18. We now turn to the proof of Lemma 1.18. The strategy is based on ideas introduced in [CS24b] to prove convergence for the discrete time version of the single site model in Example 1.1, where the irreducibility was deduced from the analysis of a fragmentation process and a coupon collector argument. However, the dynamics considered here differs from that in [CS24b] in two respects: first, it evolves in continuous time, and second, exchanges are allowed between any pair of sites. The passage to continuous time entails a natural modification of the original argument, which is based on a representation of the solution in terms of random binary trees. More substantially, allowing exchanges across arbitrary pairs of sites requires a new analysis, since fragmentation and coupon-collecting arguments alone are no longer sufficient; one must also account for the "mixing" within each class $A \in \mathcal{A}$ in

order to conclude. The proof is broken down into several steps. We start with a graphical representation needed to take care of the continuous time nature of the process.

1.2.5 Graphical representation

We introduce some notation to handle binary rooted trees. The infinite rooted binary tree is represented as the set $\mathbb{T} := \{\emptyset\} \cup_{k\geq 1} \{0,1\}^k$, where $\emptyset$ denotes the root, and an element $x \in \{0,1\}^k$ represents a node in the k-th generation. For $x \in \mathbb{T}$ we let $|x|$ denote the length of the sequence x, interpreted as the depth of the node x. We equip $\mathbb{T}$ with the partial order $\prec$ defined by saying that $x \prec y$, if $|x| < |y|$ and $y = (x, z)$ for some $z \in \{0,1\}^{|y|-|x|}$, where (x, z) denotes the concatenation of the two sequences. A finite rooted binary tree is a finite subset γ of $\mathbb{T}$ which satisfies the following:

1. For any $x \in \gamma$ and any $y \in \mathbb{T}$, if $y \prec x$ then $y \in \gamma$.

2. For any $x \in \gamma$ either $\{(x,0),(x,1)\} \subset \gamma$ or $\{(x,0),(x,1)\} \cap \gamma = \emptyset$.

The set of maximal elements in γ (for the order $\prec$) that is, its leaves, is denoted by

$$L(\gamma) := \{x \in \gamma \ : \ \nexists y \in \gamma, \ x \prec y\}.$$

Given a finite rooted binary tree γ, and $p \in \mathcal{P}(\Omega)$, we define the distribution $T_\gamma(p)$ recursively as follows. Assign to each node u in γ a probability distribution $\nu_u \in \mathcal{P}(\Omega)$, in such a way that each leaf node $u \in L(\gamma)$ is assigned the distribution p, and each internal node is assigned the distribution resulting from the binary collision $\mu_1 \circ \mu_2$, where the measures μ_1, μ_2 are assigned to its two children. The distribution $T_\gamma(p)$ is defined as the distribution at the root of the tree obtained in this way. See Figure 1.1 for an illustration.

More generally, given a collection of measures $(p_1, p_2, \ldots, p_L) \in \mathcal{P}(\Omega)^L$ and a binary tree γ with L leaves $L(\gamma) = \{x_1, \ldots, x_L\}$, we define $T_\gamma(p_1, \ldots, p_L)$ by assigning p_i to the leaf x_i, and then recursively applying the binary collision rule at internal nodes up to the root.

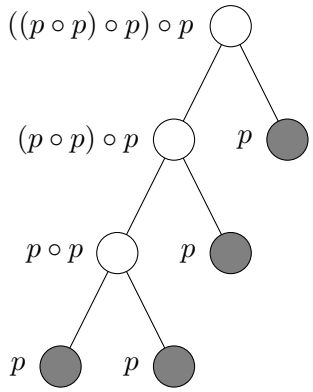


Figure 1.1: An example of a binary tree γ with four leaves, and the corresponding measure at the root: $T_\gamma(p) = (((p \circ p) \circ p) \circ p)$.

We now define the binary branching process, namely a stochastic process $(\gamma_t)_{t\geq 0}$ taking values in the space of finite rooted binary trees. The process evolves as follows: start with $\gamma_0 = \{\emptyset\}$, the single-node tree (the root). Each leaf of the tree is equipped with an exponential clock of mean one. When a clock associated with a leaf x rings at time t, the leaf splits, and two children $(x, 0)$ and $(x, 1)$ are added to the tree, making x an internal node. Equivalently, γ_t can be interpreted as a first-passage percolation process on the infinite binary tree $\mathbb{T}$. Let $(\mathcal{E}_x)_{x\in\mathbb{T}}$ be a collection of i.i.d. exponential random variables with mean one. Define the passage time to a vertex $x \in \mathbb{T}$ as

$$\tau_x := \sum_{y\prec x} \mathcal{E}_y, \tag{1.24}$$

where the sum runs over all ancestors y of x (excluding x itself). Then, the random tree at time t is given by

$$\gamma_t := \{x \in \mathbb{T} \ : \ \tau_x \leq t\}\,. \tag{1.25}$$

We denote by $\mathbf{P}$ the law of the random process $(\gamma_t)_{t\geq 0}$, and by $\mathbf{E}$ the associated expectation. The following representation of the continuous time equation can be derived as is [CLL25, Lemma 3.1] or, alternatively, as in [CP24, Section 3.2]. For completeness, we give a proof.

Lemma 1.20 ([CLL25], Lemma 3.1)**.** *Let (γ_t) be the branching process defined above. Then, for any initial measure $p \in \mathcal{P}(\Omega)$, the solution p_t to the nonlinear equation (1.5) is given by*

$$p_t = \mathbf{E}\left[T_{\gamma_t}(p)\right].$$

Proof. We set $\rho_t := \mathbf{E}\left[T_{\gamma_t}(p)\right]$ clearly we have $\rho_0 = p$ so we only need to show that ρ_t satisfies (1.5). Setting $\nu_t(\gamma) := \mathbf{P}\left[\gamma_t = \gamma\right]$ we have

$$\rho_t = \sum_{\gamma} \nu_t(\gamma) T_\gamma(p) = \nu_t(\{\emptyset\})\mu + \sum_{\gamma\neq\{\emptyset\}} \nu_t(\gamma) T_\gamma(p)\,. \tag{1.26}$$

By definition of γ_t one has (recall the notation (1.25))

$$\nu_t(\{\emptyset\}) = e^{-t} \quad \text{and} \quad \partial_t \nu_t(\gamma) = \mathbf{1}_{\{\gamma\neq\{\emptyset\}\}} \nu_t(\gamma^0)\nu_t(\gamma^1) - \nu_t(\gamma)\,. \tag{1.27}$$

The second identity in (1.27) can be checked as follows. If $\gamma = \{\emptyset\}$, then $\nu_t(\gamma) = e^{-t}$ and the identity follows. If $\gamma \neq \{\emptyset\}$ then the first branching occurs at some time $s \in [0, t]$ and we have

$$\nu_t(\gamma) = \int_0^t e^{-(t-s)} \nu_s(\gamma^0)\nu_s(\gamma^1) ds\,,$$

and differentiating this expression yields the asserted identity.

By construction, if $\gamma \neq \{\emptyset\}$, one has that $T_\gamma(p) = T_{\gamma^0}(p) \circ T_{\gamma^1}(p)$. Thus differentiating

(1.26) and using (1.27) we obtain

$$\begin{aligned}\partial\rho_t &= -e^{-t}p + \sum_{\gamma\neq\{\emptyset\}} \Big[\nu_t(\gamma^0)\nu_t(\gamma^1)T_{\gamma^0}(p)\circ T_{\gamma^1}(p) - \nu_t(\gamma)T_\gamma(p)\Big] \\ &= -e^{-t}p + \sum_{\gamma',\gamma''} \nu_t(\gamma')\nu_t(\gamma'')T_{\gamma'}(p)\circ T_{\gamma''}(p) - \sum_{\gamma\neq\{\emptyset\}} \nu_t(\gamma)T_{\gamma'}(p) \\ &= \rho_t\circ\rho_t - \rho_t\,,\end{aligned}$$

where the sum in γ', γ'' runs over all finite binary rooted trees. This concludes the proof. □

1.2.6 Reduction to discrete time and the case $J = 0$

The discrete time version of the equation (1.5) is the dynamical system $\Phi_u, u \in \mathbb{N}$, defined by

$$\Phi_u(p) = \Phi_{u-1}(p)\circ\Phi_{u-1}(p)\,, \quad u\in\mathbb{N}, \quad \Phi_0(p) = p\,.$$

Therefore, adopting the notation introduced above, the probability measure $\Phi_u(p)$ can be written as $T_{\Gamma_u}(p)$, where Γ_u denotes the rooted regular binary tree with depth u; see Figure 1.2 for an illustration.

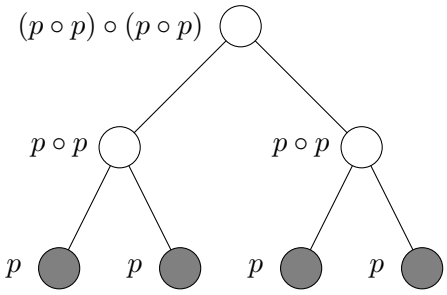


Figure 1.2: $\Phi_2(p) = (p\circ p)\circ(p\circ p)$ is the distribution at the root of the binary tree Γ_2 with depth 2.

We are going to compare the continuous time process with the discrete time approximation, and then reduce to the simpler case $J = 0$. As above, we extend the notation to include possibly distinct measures assigned to the leaves of the tree. Thus, we write $\Phi_u(p_1,\dots,p_{2^u}) = T_{\Gamma_u}(p_1,\dots,p_{2^u})$, for the distribution at the root of the regular rooted tree Γ_u when the 2^u leaves are assigned the distributions $p_1,\dots,p_{2^u}\in\mathcal{P}(\Omega)$. Similarly, we write $\Phi_u^0(p_1,\dots,p_{2^u})$ for the corresponding quantity when collisions are computed with the $J = 0$ kernel $\mathcal{Q}_0^{\mathcal{K}}$ instead of $\mathcal{Q}_J^{\mathcal{K}}$.

Lemma 1.21. *Fix a symmetric, stochastic $n\times n$ matrix $\mathcal{K}$, and let $\mathcal{A}$ denote its irreducible components. Fix $m\in(-1,1)^{\mathcal{A}}$. Suppose there exist a positive constant $\varepsilon := \varepsilon(m,\mathcal{K})$ and a positive integer $u := u(m,\mathcal{K})$ such that*

$$\Phi_u^0(p_1,\dots,p_{2^u})(\tau)\geq\varepsilon, \quad \forall\tau\in\Omega, \tag{1.28}$$

for every choice of $(p_1,\dots,p_{2^u})$ such that $m_{\mathcal{A}}(p_i) = m$. Then, there exist constants $t_0 > 0$

and $\varepsilon_0 > 0$ depending only on $m, \mathcal{K}, J$, such that for all $p \in \mathcal{P}(\Omega)$ with $m_{\mathcal{A}}(p) = m$, the solution p_t of (1.5) satisfies

$$p_t(\tau) \geq \varepsilon_0, \quad \forall \tau \in \Omega\,,\ \forall\, t \geq t_0\,. \tag{1.29}$$

Proof. As a first step, we show that assumption (1.28) extends to the case of arbitrary interaction J. Specifically, there exists a constant $\varepsilon_1 := \varepsilon_1(J, \mathcal{K}, m) > 0$ such that

$$\Phi_u(p_1, \dots, p_{2^u})(\tau) \geq \varepsilon_1, \tag{1.30}$$

for all $\tau \in \Omega$, and for every choice of $(p_1, \dots, p_{2^u})$ such that $m_{\mathcal{A}}(p_i) = m$. Indeed, observe that

$$\mathcal{Q}_J^{\mathcal{K}} = \delta_J \mathcal{Q}_0^{\mathcal{K}} + (1 - \delta_J)\overline{\mathcal{Q}}, \qquad \delta_J := \tfrac{1}{2}\, e^{-4\bar{J}} \tag{1.31}$$

where $\overline{\mathcal{Q}}$ is a probability kernel on $\Omega \times \Omega$ that does not require explicit specification, and $\bar{J} := \max_i \sum_j |J_{i,j}|$. To prove (1.31) it suffices to use the definition (1.3) and note that

$$p_J(\ell, k \mid \sigma, \sigma') \geq \frac{1}{1 + e^{4\bar{J}}}\,,$$

for all $\sigma, \sigma' \in \Omega$ and $\ell, k, \in [n]$. From the identity (1.31) it follows that, with probability δ_J the collision is performed according to the kernel $\mathcal{Q}_0^{\mathcal{K}}$. Since in the regular binary tree T_{Γ_u} there are $N = 2^u - 1$ collisions, one can write, for any $u \in \mathbb{N}$, and any choice of $p_1, \dots, p_{2^u} \in \mathcal{P}(\Omega)$,

$$\Phi_u(p_1, \dots, p_{2^u}) = \delta_J^N \Phi_u^0(p_1, \dots, p_{2^u}) + (1 - \delta_J^N)\Psi(u, p_1, \dots, p_{2^u})\,,$$

where $\Psi(u, p_1, \dots, p_{2^u}) \in \mathcal{P}(\Omega)$ is some unspecified distribution. Therefore, the assumption (1.28) implies (1.30) with the same u and with $\varepsilon_1 = \varepsilon \delta_J^N$.

Now, fix $p \in \mathcal{P}(\Omega)$ such that $m_{\mathcal{A}}(p) = m$. Using the representation formula for p_t from Lemma 1.20, and considering the regular tree Γ_u, one has the lower bound

$$p_t(\tau) = \mathbf{E}[T_{\gamma_t}(p)(\tau)] \geq \mathbf{E}[T_{\gamma_t}(p)(\tau);\, \Gamma_u \subset \gamma_t]. \tag{1.32}$$

On the event $\Gamma_u \subset \gamma_t$, there exists a (random) vector $(\mu_1, \dots, \mu_{2^u})$, where each $\mu_i \in \mathcal{P}(\Omega)$ is the result of recursive collisions along subtrees rooted at the leaves of Γ_u. Due to the conservation law in Lemma 1.17, each distribution μ_i satisfies $m_{\mathcal{A}}(\mu_i) = m_{\mathcal{A}}(p) = m$, and therefore the inequality (1.30) applies. It follows that, whenever $\Gamma_u \subset \gamma_t$,

$$T_{\gamma_t}(p)(\tau) = \Phi_u(\mu_1, \dots, \mu_{2^u})(\tau) \geq \varepsilon_1. \tag{1.33}$$

In conclusion, for all $\tau \in \Omega$,

$$p_t(\tau) \geq \varepsilon_1 \mathbf{P}(\Gamma_u \subset \gamma_t). \tag{1.34}$$

To estimate the latter probability, note that $\Gamma_u \subset \gamma_t$, iff for each leaf $x \in L(\Gamma_u)$, its

associated arrival time satisfies $\tau_x \le t$, where τ_x is distributed as a Gamma$(u,1)$ random variable; see (1.24). By a union bound, and using a standard tail estimate, for $t > u$,

$$\mathbf{P}(\Gamma_u \subset \gamma_t) \ge 1 - 2^u e^{-t}(et/u)^u .$$

Choosing $t_0 = t_0(u) > 0$ such that the above probability is at least $1/2$, one concludes the desired bound (1.29) with $\varepsilon_0 = \varepsilon_1/2$. □

1.2.7 Proof of Lemma 1.18

According to Lemma 1.21, it suffices to establish that the assumption (1.28) is satisfied. Note that in this case one actually obtains a slightly stronger statement, since the bound (1.22) is uniform over all $p \in \mathcal{P}(\Omega)$ with the given magnetization vector $m_{\mathcal{A}}(p) = m \in (-1,1)^{\mathcal{A}}$.

We now focus on establishing the validity of the assumption (1.28). The set of stationary measures consists of non-homogeneous products of Bernoulli measures $\pi = \pi^\alpha$, where the parameters $\alpha = (\alpha_i) \in (0,1)^n$ are constant within each component of the partition $\mathcal{A}$, that is

$$\pi(\sigma) = \textstyle\prod_{i=1}^n \pi_i(\sigma_i),$$

with $\pi_i(\sigma) = \mathrm{Be}(\alpha_i)$, and $\alpha_i = \alpha_j$ for all $i, j \in A$, for all $A \in \mathcal{A}$; see Lemma 1.16. Moreover, for any $m \in (-1,1)^{\mathcal{A}}$ there is a unique vector α with $\alpha_i = \alpha_j$ for all $i, j \in A$, for all $A \in \mathcal{A}$, such that the associated measure $\pi = \pi^\alpha$ satisfies $m_{\mathcal{A}}(\pi) = m$; see Lemma 1.14.

Lemma 1.22. *Fix $m \in (-1,1)^{\mathcal{A}}$, and let π denote the associated product of Bernoulli measures as above with $m_{\mathcal{A}}(\pi) = m$. Then, for every $\varepsilon > 0$, there exists $u := u(\varepsilon, \mathcal{K}, m)$ such that*

$$\|\Phi_u^0(p_1, \dots, p_{2^u}) - \pi\|_{\mathrm{TV}} \le \varepsilon, \tag{1.35}$$

for all choices of $p_1, \dots, p_{2^u} \in \mathcal{P}(\Omega)$ with $m_{\mathcal{A}}(p_i) = m$.

Before proving the lemma we show that the assumption (1.28) is satisfied if the lemma holds. Indeed, since $m \in (-1,1)^{\mathcal{A}}$ one has $\pi_* := \min_{\sigma\in\Omega} \pi(\sigma) > 0$. Thus, by Lemma 1.22 with $\varepsilon = \pi_*/2$, and by definition of the total variation distance,

$$\Phi_u^0(p_1, \dots, p_{2^u})(\tau) \ge \pi(\tau) - \frac{\pi_*}{2} \ge \varepsilon, \quad \forall \tau \in \Omega.$$

This ends the proof of Lemma 1.18 assuming the validity of Lemma 1.22.

1.2.8 Marked Partition Process

We prove Lemma 1.22 by deriving an explicit representation of the distribution $\Phi_u^0(p_1, \dots, p_{2^u})$ via a suitable discrete time fragmentation process, augmented with randomly walking labels, which we call the *Marked Partition Process* (MPP). Let $X = \{0\} \cup [n]$ denote the space of marks. At each time $u \in \mathbb{N}$, the marked partition process takes a value $\mathcal{C}^u$ of the form $\mathcal{C}^u = (C_1^u, \dots, C_{2^u}^u)$, with each fragment C_i^u given by $C_i^u = (A_i^u, x_i^u)$, where $A_i^u \subset [n]$

is called the *spatial part*, and $x_i^u \in X$ is called the *mark* of the fragment C_i^u. The spatial parts are such that $\{A_i^u,\, i = 1, \dots, 2^u\}$ form a partition of $[n]$, for each $u \in \mathbb{N}$, that is $A_i^u \cap A_j^u = \emptyset$ for $i \neq j$ and $\cup_i A_i^u = [n]$ (the sets A_i^u can be empty). To define the evolution of the process we introduce some more notation. For any $x \in [n]$ we let $X^x(u)$, for $u \in \mathbb{N}$ denote the Markov chain with initial state x and transition matrix $\mathcal{K}$. Similarly, we let $Y^x(u)$, for $u \in \mathbb{N}$, be the Markov chain with initial state $x \in [n]$, and transition matrix given by the lazy kernel

$$\bar{\mathcal{K}} = \frac{1}{n}\mathcal{K} + \frac{n-1}{n} I,$$

where I is the identity $n \times n$ matrix. With this notation, the MPP $\mathcal{C}^0, \mathcal{C}^1, \dots$ is defined as follows. At time 0, we start with the single fragment $\mathcal{C}^0 = (C_1^0)$, where $C_1^0 = ([n], 0)$. Recursively, the process is defined by the following splitting and marking rules. If $\mathcal{C}^u = (C_1^u, \dots, C_{2^u}^u)$, then, independently for every $i = 1, \dots, 2^u$, sample a uniform random variable $U \in [n]$ and a uniform random variable $B \in \{1, 2, 3, 4\}$, and set $\mathcal{C}^{u+1} = (C_1^{u+1}, \dots, C_{2^{u+1}}^{u+1})$, according to the rules

1. if $B = 1$, then $(C_{2i-1}^{u+1}, C_{2i}^{u+1}) = (C_i^u, (\emptyset, 0))$,
2. if $B = 2$, then $(C_{2i-1}^{u+1}, C_{2i}^{u+1}) = ((\emptyset, 0), C_i^u)$,
3. if $B = 3$, $C_i^u = (A, 0)$ and $U \in A$, then $(C_{2i-1}^{u+1}, C_{2i}^{u+1}) = ((A \setminus \{U\}, 0), (\{U\}, X^U(1))$,
4. if $B = 3$, $C_i^u = (A, 0)$ and $U \notin A$, then $(C_{2i-1}^{u+1}, C_{2i}^{u+1}) = ((A, 0), (\emptyset, 0))$,
5. if $B = 3$, $C_i^u = (\{i\}, x)$, $x \in [n]$, then $(C_{2i-1}^{u+1}, C_{2i}^{u+1}) = ((\{i\}, Y^x(1)), (\emptyset, 0))$,
6. if $B = 4$, repeat the $B = 3$ rules and then invert the order of $(C_{2i-1}^{u+1}, C_{2i}^{u+1})$.

We remark that, since we start with $\mathcal{C}^0 = (C_1^0)$, where $C_1^0 = ([n], 0)$, the above rules are such that a nonzero mark can only be assigned to a spatial part consisting of a single site. Thus, the above rules yield a well defined evolution. Note that, by construction, the spatial parts are such that $(A_i^u,\, i = 1, \dots, 2^u)$ forms a partition of $[n]$, at any time u. We remark that the partition (A_i^u) is monotone in the sense that it can only become finer with time, until it consists only of empty sets or single sites with nonzero marks. One can define the *fragmentation time*

$$H = \min\left\{u \in \mathbb{N} : \text{all nonempty sets have nonzero marks in } \mathcal{C}^u\right\}.$$

For each time u, there is at most one spatial part A_i^u which can have cardinality larger than 1, and, with probability $1/2$ it loses one element every time the uniform random variable $U \in [n]$ belongs to it. Therefore, the erosion of the largest spatial part can be described by a simple coupon collecting process, where at each step, with probability $1/2$, a coupon is chosen uniformly at random, while with probability $1/2$ no coupon is collected. The fragmentation time H coincides with the first time where n distinct coupons have

been collected, and thus a union bound gives

$$\mathbb{P}(H \geq u) \leq n \left(1 - \tfrac{1}{2n}\right)^u \leq n\, e^{-u/2n}\,. \tag{1.36}$$

We note that there is a natural representation of the Marked Partition Process in terms of a regular binary tree. Namely, we can assign to each node of the rooted regular binary tree with depth $u \in \mathbb{N}$, $\Gamma_u = \{\emptyset\} \cup_{k=1}^u \{0,1\}^k$, a marked set in such a way that the root is assigned the set $C_1^0 = ([n], 0)$, and if the nodes with depth k are assigned, say from left to right, the sets $(C_1^k, \dots, C_{2^k}^k)$, then the nodes at the next generation are assigned, respectively, from left to right, the sets $(C_1^{k+1}, \dots, C_{2^{k+1}}^{k+1})$, with the pair $(C_{2i-1}^{k+1}, C_{2i}^{k+1})$, being assigned to the two children of the node corresponding to C_i^k. In this way, starting from the root, the marks along any path is always 0 until the first time where it takes some value $x \in [n]$, after which it evolves, along the unique branch with nonempty fragments, with marks given the lazy Markov chain $\bar{\mathcal{K}}$. In particular, any marked set of the form $(\{i\}, x)$, $x \neq 0$, must be such that i and x belong to the same irreducible component $A \in \mathcal{A}$.

We now describe how the above defined process encodes the measure $\Phi_u^0(p_1, \dots, p_{2^u})$. For any $u \in \mathbb{N}$, and any realization of $\mathcal{C}^u$, for any vector $\vec{p}(u) = (p_1, \dots, p_{2^u})$, define the measure $\Psi(\vec{p}(u), \mathcal{C}^u) \in \mathcal{P}(\Omega)$ by the following rule. Given $\mathcal{C}^u = (C_1^u, \dots, C_{2^u}^u)$, there is at most one marked set C_i^u of the form $(A, 0)$, with $\emptyset \neq A \subset [n]$ and mark 0, and all other marked sets are either empty or of the form $(\{j\}, x_j)$ for some $j \in [n]$ and some nonzero mark $x_j \in [n]$. For any $\tau \in \Omega$, $p \in \mathcal{P}(\Omega)$, $i = 1, \dots, 2^u$, we define

$$\psi(p, C_i^u)(\tau) = \begin{cases} p_A(\tau_A) & \text{if } C_i^u = (A, 0) \\ p_x(\tau_j) & \text{if } C_i^u = (\{j\}, x) \end{cases}$$

In this formula, we use the notation p_A for the marginal at A of the measure p and p_x for the marginal at $\{x\}$ of the measure p, and as usual τ_A denotes the spin configuration $(\tau_x,\, x \in A)$, with τ_x denoting the spin at x. Moreover, we adopt the convention that, when $A = \emptyset$, then we set $p_A(\tau_A) = 1$. Finally, we set

$$\Psi(\vec{p}(u), \mathcal{C}^u)(\tau) = \prod_{i=1}^{2^u} \psi(p_i, C_i^u)(\tau)\,.$$

Note that $\Psi(\vec{p}(u), \mathcal{C}^u) \in \mathcal{P}(\Omega)$, for any u, any choice of the distribution vector $\vec{p}(u) = (p_1, \dots, p_{2^u})$, and for any realization of the MPP $(\mathcal{C}^u, u = 0, 1, \dots)$. In words, $\Psi(\vec{p}(u), \mathcal{C}^u)$ is obtained by taking products of the marginals over the largest spatial part $(A, 0)$ in the marked partition $\mathcal{C}^u$, and the marginals over the single sites identified by the nonzero marks, with the caveat that the marginal of $\Psi(\vec{p}(u), \mathcal{C}^u)$ on the single sites appearing with a nonzero mark is given by the single site marginal identified by the nonzero mark.

Lemma 1.23. *For any $u \in \mathbb{N}$, $\tau \in \Omega$, any $p_1, \dots, p_{2^u} \in \mathcal{P}(\Omega)$,*

$$\Phi_u^0(p_1, \dots, p_{2^u}) = \mathbb{E}\left[\Psi(\vec{p}(u), \mathcal{C}^u)(\tau)\right]\,, \tag{1.37}$$

where the expectation is over the MPP.

Proof. We use induction over u. Note that it is true at $u = 0$, since $\mathcal{C}^u$ is given by the single marked set $C_1^0 = ([n], 0)$ and $\Phi_u^0(p) = p$. Let us check that for all $p, q \in \mathcal{P}(\Omega)$,

$$\Phi_1^0(p,q)(\tau) = p \circ q\,(\tau) = \mathbb{E}[\psi(p, C_1^1)(\tau)\psi(q, C_2^1)(\tau)] = \mathbb{E}\left[\Psi(p, q, \mathcal{C}^1)(\tau)\right], \tag{1.38}$$

which is the identity (1.37) for $u = 1$. To see this, note that by definition of the collision kernel $\mathcal{Q}_0^{\mathcal{A}}$, in the case $J = 0$ one has

$$\begin{aligned}(p \circ q)(\tau) = \frac{1}{4}\,p(\tau) + \frac{1}{4}\,q(\tau) + \\ + \frac{1}{4n}\sum_{\ell,k\in[n]} \mathcal{K}(\ell,k)\left[p_{[n]\setminus\{\ell\}}(\tau_{[n]\setminus\{\ell\}})\,q_k(\tau_\ell) + p_k(\tau_\ell)\,q_{[n]\setminus\{\ell\}}(\tau_{[n]\setminus\{\ell\}})\right],\end{aligned} \tag{1.39}$$

where p_A denotes the marginal of p over a subset $A \subset [n]$, and we use p_k for the marginal of p over a single site $\{k\}$. Points (1) and (2) in the definition of the MPP take care of the terms $\frac{1}{4}\,p$ and $\frac{1}{4}\,q$ in (1.39). Moreover, points (3)-(4)-(5) provide the term

$$\frac{1}{4}\,p_{[n]\setminus\{U\}}(\tau_{[n]\setminus\{U\}})\,q_{X^U(1)}(\tau_U)$$

in (1.39). Finally, point (6) provides the term

$$\frac{1}{4}\,p_{X^U(1)}(\tau_U)\,q_{[n]\setminus\{U\}}(\tau_{[n]\setminus\{U\}})$$

in (1.39). Taking the expectation over $U, X^U(1)$ yields (1.38).

Reasoning in the same way, one obtains, for any realization of C_i^u, for all τ,

$$\psi(p \circ q, C_i^u)\,(\tau) = \mathbb{E}\left[\psi(p, C_{2i-1}^{u+1})(\tau)\psi(q, C_{2i}^{u+1})(\tau) \mid C_i^u\right], \tag{1.40}$$

where $\mathbb{E}[\cdot \mid C_i^u]$ denotes the expectation over one step of the MPP $C_i^u \mapsto (\mathcal{C}_{2i-1}^{u+1}, \mathcal{C}_{2i}^{u+1})$, conditionally on C_i^u.

Assuming the validity of (1.37) at u, the formula for $u+1$ can be derived by using (1.38) applied to the collisions $p_{2i-1} \circ p_{2i}$, for every $i = 1, \dots, 2^u$. Indeed, setting $\vec{q}(u) = (q_1, \dots, q_{2^u})$ with $q_i := p_{2i-1} \circ p_{2i}$, by repeated application of (1.40), for all $\tau \in \Omega$,

$$\begin{aligned}&\mathbb{E}\left[\Psi(\vec{p}(u+1), \mathcal{C}^{u+1})(\tau) \mid \mathcal{C}^s,\, s \le u\right] \\ &= \prod_{i=1}^{2^u} \mathbb{E}\left[\psi(p_{2i-1}, C_{2i-1}^{u+1})(\tau)\psi(p_{2i}, C_{2i}^{u+1})(\tau) \mid \mathcal{C}^s,\, s \le u\right] = \Psi(\vec{q}(u), \mathcal{C}^u)(\tau)\,.\end{aligned}$$

Taking the expectation, and using the inductive assumption, one finds

$$\mathbb{E}\left[\Psi(\vec{p}(u+1), \mathcal{C}^{u+1})\right] = \Phi_u^0(q_1, \dots, q_{2^u})\,.$$

Since $\Phi_u^0(q_1, \dots, q_{2^u}) = \Phi_{u+1}^0(p_1, \dots, p_{2^{u+1}})$, this proves the desired identity (1.37) for all

u and all $\vec{p}(u)$. □

1.2.9 Proof of Lemma 1.22

Once the representation in terms of the MPP is established we turn to the proof of Lemma 1.22. To prove (1.35), we observe that for any $u_0 < u$, using Lemma 1.23, Jensen's inequality and the fact that $\{H \le u_0\}$ is measurable with respect to $\mathcal{C}^{\le u_0} := (\mathcal{C}^s, s \le u_0)$, one has

$$\begin{aligned}\|\Phi_u^0(p_1,\dots,p_{2^u}) - \pi\|_{\mathrm{TV}} \le \mathbb{P}(H > u_0) + \\ + \mathbb{E}\left[\left\|\mathbb{E}\left[\Psi(p_1,\dots,p_{2^u},\mathcal{C}^u) \mid \mathcal{C}^{\le u_0}\right] - \pi\right\|_{\mathrm{TV}}; H \le u_0\right]. \end{aligned} \tag{1.41}$$

We take $u_0 = u_0(\varepsilon, n)$ such that $n\, e^{-u_0/2n} \le \varepsilon/2$, so that by (1.36) the first term above is at most $\varepsilon/2$. To estimate the second term, the key observation is that conditionally on $\mathcal{C}^{\le u_0}$, on the event $H \le u_0$, one has the non-homogeneous product of Bernoulli measures

$$\mathbb{E}\left[\Psi(p_1,\dots,p_{2^u},\mathcal{C}^u) \mid \mathcal{C}^{\le u_0}\right] = \bigotimes_{j=1}^n \mathrm{Be}(\xi_j)\,,$$

where, for each $j \in [n]$, the parameter $\xi_j = \xi_j(p_1,\dots,p_{2^u},\mathcal{C}^{\le u_0})$ is obtained as follows. Let $C^u_{i_*(j)} = (\{j\}, x_*(j))$ denote the unique marked set with spatial part given by $\{j\}$, and let $W_j = p_{i_*(j)}(\sigma_{x_*(j)} = +1)$ denote the probability of the event $\sigma_{x_*(j)} = +1$, when σ is distributed according to $p_{i_*(j)}$. Then

$$\xi_j = \mathbb{E}\left[W_j \mid \mathcal{C}^{\le u_0}\right].$$

Since the total variation distance between two Bernoulli variables is the distance between the parameters, one has

$$\left\|\mathbb{E}\left[\Psi(p_1,\dots,p_{2^u},\mathcal{C}^u) \mid \mathcal{C}^{\le u_0}\right] - \pi\right\|_{\mathrm{TV}} \le n \max_j |\xi_j - \alpha_j|\,.$$

An important feature is that the random site $x_*(j)$, conditionally on $\mathcal{C}^{\le u_0}$, on the event $H \le u_0$, has the distribution of $Z^j(u - u_*(j))$, where $u_*(j) \le u_0$ denotes the first time the site j receives a nonzero mark in the MPP, and $Z^j(k)$ stands for the distribution after k steps of the Markov chain with transition $\widehat{\mathcal{K}} := \frac{1}{2}I + \frac{1}{2}\bar{\mathcal{K}}$, with initial state j. On the other hand we note that the parameter α_j, by our assumption $m_{\mathcal{A}}(p_i) = m$ for all p_i, satisfies $\alpha_j = \mathbb{E}[p_i(\sigma_{y_*(j)} = +1)]$, for any $i = 1,\dots,2^u$, where $y_*(j)$ is distributed uniformly in the irreducible component $A(j) \in \mathcal{A}$ containing j. Call $u(j)$ the uniform distribution over $A(j)$. As a consequence, we may estimate the difference $|\xi_j - \alpha_j|$ by using the optimal coupling between the chain with transition $\widehat{\mathcal{K}}$ and the uniform distribution over $A(j)$. Since $u_*(j) \le u_0$, we can estimate

$$\max_j |\xi_j - \alpha_j| \le \max_j \|\widehat{\mathcal{K}}^{u-u_0}(j,\cdot) - u(j)\|_{\mathrm{TV}}\,.$$

By elementary mixing time theory, for any $\varepsilon > 0$ there exists an integer $u > u_0$ such that the above is at most $\varepsilon/2$. Together with (1.41) this concludes the proof.

1.3 Proof of Theorem 1.11

In this section we prove Theorem 1.11. We work with the multi-component mean-field nonlinear dynamics of Definition 1.3, associated to a partition $\mathcal{A}$ of $[n]$ and an interaction matrix J satisfying the high-temperature condition $\overline{J} = \max_{i\in[n]} \sum_{j\neq i} |J_{i,j}| < 1/10$. We show that, in the multi-component mean-field setting, the solution p_t of (1.5) satisfies

$$W_1(p_t, \mu_{J,h}) \leq e^{-\varepsilon t/n}\, W_1(p_0, \mu_{J,h}) \tag{1.42}$$

for all initial states $p_0 \in \mathcal{P}(\Omega)$, with rate $\varepsilon(\overline{J}) = \frac{1}{2}(1 - 10\overline{J})$.

For the reader's convenience, we recall the collision kernel for the multi-component mean-field dynamics, associated to the transport kernel $\mathcal{K}_{\mathcal{A}}(\ell, k) = \frac{1}{|A(\ell)|}\mathbf{1}(k \in A(\ell))$, where $A(\ell)$ denotes the unique block of $\mathcal{A}$ containing ℓ:

$$\begin{aligned}\mathcal{Q}_J^{\mathcal{K}_{\mathcal{A}}}(\sigma,\sigma';\tau,\tau') = \frac{1}{n}\sum_{\ell,k\in[n]} \mathcal{K}_{\mathcal{A}}(\ell,k) \times \Big[& p_J(\ell,k \mid \sigma,\sigma')\, \mathbf{1}(\tau = S_\ell(\sigma,\sigma'_k),\, \tau' = S_k(\sigma',\sigma_\ell)) \\ & + (1 - p_J(\ell,k \mid \sigma,\sigma'))\, \mathbf{1}(\tau = \sigma,\, \tau' = \sigma')\Big].\end{aligned} \tag{1.43}$$

where $p_J(\ell, k \mid \sigma, \sigma') \in (0,1)$ is the exchange probability defined in (1.2). Throughout this section, T denotes the collision operator associated to the kernel (1.43): for $p, q \in \mathcal{P}(\Omega)$ and $\tau \in \Omega$,

$$T(p,q)(\tau) = \frac{1}{2}\sum_{\sigma,\sigma',\tau'} \big(p(\sigma)q(\sigma') + p(\sigma')q(\sigma)\big)\, \mathcal{Q}_J^{\mathcal{K}_{\mathcal{A}}}(\sigma,\sigma';\tau,\tau'). \tag{1.44}$$

We use the simplified notation for (1.44) when $p = q$ writing $T(p) = T(p,p)$.

The proof of Theorem 1.11 rests on two ingredients. The first is Lemma 1.24 below, which establishes a contraction property for the operator $T(p,q)$ with the respect of Wasserstein distance. The second is the graphical representation of the solution to (1.5) described in Subsection 1.2.5, together with the conservation of the $\mathcal{A}$-magnetization profile $m(\cdot, \mathcal{A})$ stated in Proposition 1.17. We state the lemma first and then show how it implies Theorem 1.11, postponing the proof of the lemma to the end of the section.

Lemma 1.24. *Let $p(i), q(i) \in \mathcal{P}(\Omega)$ for $i \in \{1,2\}$ be probability measures satisfying $m(p(i), \mathcal{A}) = m(q(i), \mathcal{A})$ for $i \in \{1,2\}$. Then*

$$\begin{aligned}& W_1(T(p(1),q(1)),\, T(p(2),q(2))) \\ & \leq \frac{1}{4}\Big(1 - \frac{\varepsilon(\overline{J})}{n}\Big)\Big[W_1(p(1),p(2)) + W_1(p(1),q(2)) + W_1(q(1),p(2)) + W_1(q(1),q(2))\Big],\end{aligned} \tag{1.45}$$

where $\varepsilon(\overline{J}) = \frac{1}{2} - 5\overline{J}$.

1.3.1 Proof of Theorem 1.11 assuming Lemma 1.24

We refer to Subsection 1.2.5 for the graphical construction used in the proof.
Given an initial state $p \in \mathcal{P}(\Omega)$, the solution p_t, $t \geq 0$, admits the representation

$$p_t = \mathbf{E}[T_{\gamma_t}(p)],$$

where γ_t is a random binary tree describing the collision history up to time t, and $\mathbf{E}$ denotes expectation with respect to the law of γ_t. Here, $T_\gamma(p)$ denotes the extension of definition of collision operator (1.44) to an arbitrary binary tree γ and applied to the distribution p, as described in Subsection 1.2.5.
As seen in the proof of Lemma 1.20, the law $\mathbf{P}$ of γ_t satisfies

$$\mathbf{P}(\gamma_t = \gamma) = \mathbf{1}(\gamma = \emptyset)e^{-t} + \mathbf{1}(\gamma \neq \emptyset) \int_0^t e^{-(t-s)} \mathbf{P}(\gamma_s = \gamma^0)\mathbf{P}(\gamma_s = \gamma^1)\, ds, \tag{1.46}$$

where γ^0 and γ^1 denote the left and right subtrees of γ, respectively, and $\emptyset$ denotes the trivial tree consisting only of the root.
Let $\mu = \mu_{J,h(p)}$, where $h(p)$ is the unique external field $h(p) \in \mathbb{R}^n$ such that $m(\mu_{J,h(p)}) = m(p)$, as stated in Lemma 1.14. To prove the inequality (1.42), and conclude the proof of Theorem 1.11, it is sufficient to estimate the function $f(t) := \mathbf{E}[W_1(T_{\gamma_t}(p), \mu)]$, and prove that

$$f(t) \leq e^{-\frac{\varepsilon t}{n}} f(0), \tag{1.47}$$

with $\varepsilon = \varepsilon(\overline{J}) = \frac{1}{2}(1 - 10\overline{J})$. Indeed, by convexity of the Wasserstein distance and assuming (1.47), one has

$$W_1(p_t, \mu) \leq \mathbf{E}\left[W_1(T_{\gamma_t}(p), \mu)\right] \leq e^{-\varepsilon t/n} W_1(p, \mu)$$

which proves (1.42). Now, we turn to the proof of inequality (1.47).
We start by expanding $f(t)$ with respect of the law $\mathbf{P}$ of γ_t and we obtain

$$f(t) = \mathbf{P}(\gamma_t = \emptyset) W_1(p, \mu) + \sum_{\gamma \neq \emptyset} \mathbf{P}(\gamma_t = \gamma)\, W_1(T_\gamma(p), T_\gamma(\mu)),$$

where, in the in the second term on the right-hand side, we used that μ is stationary, that is, $\mu = T_\gamma(\mu)$. We apply to each term of $W_1(T_\gamma(p), T_\gamma(\mu))$ the inequality (1.45) stated in Lemma 1.24 and in order to do this we use the recursive structure of binary trees. For $\gamma \neq \emptyset$, we write $\gamma = (\gamma^0, \gamma^1)$, where γ^0 and γ^1 are the left and right sub-tree of γ. The recursive definition of the collision operator gives

$$T_\gamma(p) = T_{\gamma^0}(p) \circ T_{\gamma^1}(p). \tag{1.48}$$

By Lemma 1.24 combined with the recursive rule (1.48) we have

$$\begin{aligned} W_1(T_\gamma(p), T_\gamma(\mu)) &= W_1\Big(T_{\gamma^0}(p) \circ T_{\gamma^1}(p), T_{\gamma^0}(\mu) \circ T_{\gamma^1}(\mu)\Big) \\ &\le \frac{1-\varepsilon/n}{4}\Big[W_1(T_{\gamma^0}(p), T_{\gamma^0}(\mu)) + W_1(T_{\gamma^1}(p), T_{\gamma^0}(\mu)) \\ &\quad + W_1(T_{\gamma^0}(p), T_{\gamma^1}(\mu)) + W_1(T_{\gamma^1}(p), T_{\gamma^1}(\mu))\Big] \\ &= \frac{1-\varepsilon/n}{2}\Big[W_1(T_{\gamma^0}(p), T_{\gamma^0}(\mu)) + W_1(T_{\gamma^1}(p), T_{\gamma^1}(\mu))\Big]. \end{aligned} \tag{1.49}$$

Here we used that μ is stationary, that is $\mu = T_{\gamma^0}(\mu) = T_{\gamma^1}(\mu)$. Combining (1.46) and (1.49), we obtain

$$\begin{aligned} &f(t) \le e^{-t} f(0) + \\ &+ \frac{1-\varepsilon/n}{2} \sum_{\gamma^0,\gamma^1} \int_0^t e^{-(t-s)} \mathbf{P}(\gamma_s = \gamma^0)\mathbf{P}(\gamma_s = \gamma^1)\Big[W_1(T_{\gamma^0}(p), T_{\gamma^0}(\mu)) + W_1(T_{\gamma^1}(p), T_{\gamma^1}(\mu))\Big] ds. \end{aligned}$$

Separating the sum over γ^0 and γ^1 and exchanging the order of summation and integration yields

$$f(t) \le e^{-t} f(0) + (1-\varepsilon/n) \int_0^t e^{-(t-s)} f(s)\, ds.$$

Applying Grönwall's lemma to $e^t f(t)$ gives (1.47), which concludes the proof. □

1.3.2 Proof of Lemma 1.24

We begin with the following proposition which gives an alternative definition of the probability defined in (1.2), that we repeat for the reader's convenience

$$p_J(\ell, k \mid \sigma, \sigma') = \frac{\mu_J\left(S_\ell(\sigma, \sigma'_k)\right) \mu_J\left(S_k(\sigma', \sigma_\ell)\right)}{\mu_J\left(S_\ell(\sigma, \sigma'_k)\right) \mu_J\left(S_k(\sigma', \sigma_\ell)\right) + \mu_J(\sigma)\, \mu_J(\sigma')}, \quad \ell, k \in [n], \quad \sigma, \sigma' \in \Omega \tag{1.50}$$

Proposition 1.25. *Fix $\ell, k \in [n]$ and spin configurations $\sigma, \sigma' \in \Omega$. Then,*

$$p_J(\ell, k \mid \sigma, \sigma') = \frac{1}{1 + e^{E_{\ell,k}(\sigma,\sigma')}}, \quad E_{\ell,k}(\sigma, \sigma') = (\sigma_\ell - \sigma'_k)\Big[(J\sigma)_\ell - (J\sigma')_k\Big] \tag{1.51}$$

In particular, for every pair of configurations $(\sigma^{(1)}, \sigma'^{(1)})$ and $(\sigma^{(2)}, \sigma'^{(2)})$ we have

$$|p_J(\ell, k \mid \sigma^{(1)}, \sigma'^{(1)}) - p_J(\ell, k \mid \sigma^{(2)}, \sigma'^{(2)})| \le \frac{1}{4} |E_{\ell,k}(\sigma^{(1)}, \sigma'^{(1)}) - E_{\ell,k}(\sigma^{(2)}, \sigma'^{(2)})|, \tag{1.52}$$

and for all $\sigma, \sigma' \in \Omega$ and $\ell, k \in [n]$,

$$\frac{1}{2} - \bar{J} \;\le\; p_J(\ell, k \mid \sigma, \sigma') \;\le\; \frac{1}{2} + \bar{J}\,. \tag{1.53}$$

Proof. Let $\tau = S_\ell(\sigma, \sigma'_k)$, and $\tau' = S_k(\sigma', \sigma_\ell)$ denote the configurations obtained by

exchanging the ℓ-th spin of σ with the k-th spin of σ'. The definition (1.50) can be rewritten easily as in (1.51) with

$$E_{\ell,k}(\sigma,\sigma') = \frac{1}{2}\Big(\langle\sigma, J\sigma\rangle - \langle\tau, J\tau\rangle + \langle\sigma', J\sigma'\rangle - \langle\tau', J\tau'\rangle\Big). \tag{1.54}$$

A direct computation yields

$$\langle\sigma, J\sigma\rangle - \langle\tau, J\tau\rangle = \langle\sigma-\tau, J(\sigma-\tau)\rangle + 2\langle\sigma-\tau, J\tau\rangle. \tag{1.55}$$

Since $\sigma - \tau = (\sigma_\ell - \sigma'_k)\delta_\ell$, and $\sigma' - \tau' = (\sigma'_k - \sigma_\ell)\delta_k$, where δ_i denotes the Kronecker vector at $i \in [n]$, applying the identity (1.55) to expression of $E_{\ell,k}(\sigma,\sigma')$ in (1.54) we obtain

$$\begin{aligned} E_{\ell,k}(\sigma,\sigma') &= \frac{1}{2}(\sigma_\ell - \sigma'_k)^2[J_{\ell,\ell} + J_{k,k}] + (\sigma_\ell - \sigma'_k)[(J\tau)_\ell - (J\tau')_k)] \\ &= (\sigma_\ell - \sigma_k)[(J\sigma)_\ell - (J\sigma')_k]. \end{aligned}$$

that proves the identity (1.51).

Finally, inequality (1.52) follows from the fact that the function $f(x) = \frac{1}{1+e^x}$ is 1/4-Lipschitz. The estimates (1.53) are an immediate consequence. □

We now begin the proof of Lemma 1.24.
Note that we can decompose the collision operator T defined in (1.44) as

$$T(p,q) = \frac{1}{2}\widehat{T}(p,q) + \frac{1}{2}\widehat{T}(q,p)$$

where $\widehat{T}(p,q)$ can be written as

$$\begin{aligned} \widehat{T}(p,q)(\tau) &= \frac{1}{n}\sum_{\sigma,\sigma'} p(\sigma)q(\sigma') \qquad (1.56) \\ \times &\sum_{\ell,k\in[n]} \mathcal{K}_{\mathcal{A}}(\ell,k)\big[p_J(\ell,k \mid \sigma,\sigma')\mathbf{1}(\tau = S_\ell(\sigma,\sigma'_k)) + (1 - p_J(\ell,k \mid \sigma,\sigma'))\mathbf{1}(\tau=\sigma)\big]. \end{aligned}$$

The proof of the Lemma consists in proving

$$W_1(\widehat{T}(p(1),q(1)),\widehat{T}(p(2),q(2))) \le \Big(1 - \frac{1}{n}\Big(\frac{1}{2} - 2\overline{J}\Big)\Big)W_1(p(1),p(2)) + \frac{3\overline{J}}{n}W_1(q(1),q(2)). \tag{1.57}$$

By convexity of the W_1-distance, it is straightforward to check that (1.57) implies the statement of the Lemma. The remainder of the proof is devoted to establishing (1.57).

Algorithm 1 defines a coupling between $\widehat{T}(p(1),q(1))$ and $\widehat{T}(p(2),q(2))$. Here we use the notation π^p_{opt} (and analogously π^q_{opt}) for the optimal coupling realizing the minimum in the definition of the Wasserstein distance, see (0.8): if $(\sigma^{p(1)},\sigma^{p(2)}) \sim \pi^p_{\text{opt}}$, then $W_1(p(1),p(2)) = \mathbb{E}[|D_{\sigma(1),\sigma(2)}|]$.

To see that Algorithm 1 indeed defines a valid coupling, we examine the role of the uniform random variable $U \in [0,1]$ in the definition of $\widehat{T}(p,q)$ and show that, conditionally

Algorithm 1 Multi-mean-field coupling on partition $\mathcal{A}$ between $\widehat{T}(p(1), q(1))$ and $\widehat{T}(p(2), q(2))$

Require: A partition $\mathcal{A}$ of $[n]$. Distributions $p(i), q(i) \in \mathcal{P}(\Omega)$ for $i \in \{1, 2\}$; optimal couplings $\pi^p_{\text{opt}} \in \Gamma(p(1), p(2))$ and $\pi^q_{\text{opt}} \in \Gamma(q(1), q(2))$ with respect to the expected Hamming distance; the parameter $\overline{J}$. Given $\ell \in [n]$, $A(\ell)$ denotes the unique block $A \in \mathcal{A}$ such that $\ell \in A$.

Ensure: Coupled outputs $\tau^1 \sim \widehat{T}(p(1), q(1))$ and $\tau^2 \sim \widehat{T}(p(2), q(2))$

Sample $(\sigma^{p(1)}, \sigma^{p(2)}) \sim \pi^p_{\text{opt}}$.
Sample U uniformly in the interval $[0, 1]$.
Sample $\ell \in [n]$ uniformly at random.
Sample $k \in A(\ell)$ uniformly at random.
if $U \leq \frac{1}{2} - \overline{J}$ **then**
    Sample $\sigma^{q(1)} \sim q(1)$.
    Set $\sigma^{q(2)} = \sigma^{q(1)}$.
else
    Sample $(\sigma^{q(1)}, \sigma^{q(2)}) \sim \pi^q_{\text{opt}}$.
end if
if $p_J(\ell, k \mid \sigma^{p(1)}, \sigma^{q(1)}) \geq U$ **then**
    Set $\tau^1, \hat{\tau}^1 = S_\ell(\sigma^{p(1)}, \sigma^{q(1)}_k), S_k(\sigma^{q(1)}, \sigma^{p(1)}_\ell)$ and output τ^1.
else
    $\tau^1, \hat{\tau}^1 = \sigma^{p(1)}, \sigma^{q(1)}$ and output τ^1.
end if
if $p_J(\ell, k \mid \sigma^{p(2)}, \sigma^{q(2)}) \geq U$ **then**
    Set $\tau^2, \hat{\tau}^2 = S_\ell(\sigma^{p(2)}, \sigma^{q(2)}_k), S_k(\sigma^{q(2)}, \sigma^{p(2)}_\ell)$ and output τ^2.
else
    $\tau^2, \hat{\tau}^2 = \sigma^{p(2)}, \sigma^{q(2)}$ and output τ^2.
end if
return (τ^1, τ^2).

on $\{U \leq \frac{1}{2} - \overline{J}\}$, replacing the distribution q with any other distribution $q' \in \mathcal{P}(\Omega)$ leaves $\widehat{T}$ unchanged, provided q and q' have the same magnetization, i.e. $m(q, \mathcal{A}) = m(q', \mathcal{A})$. We start by rewriting expression (1.56) in terms of the independent random variable U, uniformly distributed on $[0, 1]$. For general $p, q \in \mathcal{P}(\Omega)$, we can define

$$\widehat{T}(p, q \mid U)(\tau) = \frac{1}{n} \sum_{\ell, k \in [n]} \mathcal{K}_{\mathcal{A}}(\ell, k) \sum_{\sigma, \sigma' \in \Omega} p(\sigma) q(\sigma') \times \\ \times \Big[\mathbf{1}(p_J(\ell, k \mid \sigma, \sigma') > U) \mathbf{1}(\tau = S_\ell(\sigma, \sigma'_k)) + \mathbf{1}(p_J(\ell, k \mid \sigma, \sigma') \leq U) \mathbf{1}(\tau = \sigma) \Big]$$

and then express $\widehat{T}(p, q)(\tau) = \mathbb{E}_U[\widehat{T}(p, q \mid U)]$. We are using the elementary fact that $\mathbb{E}_U(\mathbf{1}\left(p_J(\ell, k \mid \sigma, \sigma') > U\right) = p_J(\ell, k \mid \sigma, \sigma')$, with the notation $\mathbb{E}_U$ that is the expectation with respect to the variable U. This explains the role of variable U in the Algorithm. It remains to show that on the event $\{U \leq \frac{1}{2} - \overline{J}\}$ holds

$$\widehat{T}(p, q \mid U) = \widehat{T}(p, q' \mid U), \quad \text{for} \quad m(q, \mathcal{A}) = m(q', \mathcal{A}). \tag{1.58}$$

If $U \leq \frac{1}{2} - \overline{J}$, as a consequence of Proposition 1.25 which guarantees the lower bound $p_J(\ell, k \mid \sigma, \sigma') \geq \frac{1}{2} - \overline{J}$, we have $p_J(\ell, k \mid \sigma, \sigma') \geq U$ almost surely for all ℓ, k in the same block of $\mathcal{A}$ and all $\sigma, \sigma' \in \Omega$, and therefore the proposed exchange of spins is accepted almost surely which means $\tau = S_\ell(\sigma, \sigma'_k)$. In this case, denoting with $p_B(\sigma_B) = \sum_{\sigma_{B^c}} p(\sigma_B \sigma_{B^c})$ the marginal of p over the subset $B \subset [n]$ where $\sigma_B = (\sigma_i, i \in B)$, it follows

$$\begin{aligned}
\widehat{T}(p, q \mid U)(\tau) &= \frac{1}{n} \sum_{\ell, k \in [n]} \mathcal{K}_{\mathcal{A}}(\ell, k) \sum_{\sigma_{[n]\setminus\{\ell\}}, \sigma'_k} p_{[n]\setminus\{\ell\}}(\sigma_{[n]\setminus\{\ell\}}) q_k(\sigma'_k) \mathbf{1}(\tau = S_\ell(\sigma, \sigma'_k)) \qquad (1.59) \\
&= \frac{1}{n} \sum_{\ell \in [n]} \sum_{k \in A(\ell)} \frac{1}{|A(\ell)|} p_{[n]\setminus\{\ell\}}(\tau_{[n]\setminus\{\ell\}}) q_k(\tau_\ell) \\
&= \frac{1}{n} \sum_{\ell \in [n]} p_{[n]\setminus\{\ell\}}(\tau_{[n]\setminus\{\ell\}}) \left(\frac{1 + m(q, A(\ell))}{2} \mathbf{1}(\tau_\ell = 1) + \frac{1 - m(q, A(\ell))}{2} \mathbf{1}(\tau_\ell = -1) \right).
\end{aligned}$$

which proves (1.58). In other words, conditioned on the event $\{U \leq \frac{1}{2} - \overline{J}\}$ the collision update is "oblivious", that is it does not take into account the interaction. In the last line of (1.59), we used the definition of magnetization vector in its coordinate $A(\ell)$

$$\frac{1}{|A(\ell)|} \sum_{k \in A(\ell)} q_k(\tau_\ell) = \frac{1 + m(q, A(\ell))}{2} \mathbf{1}(\tau_\ell = 1) + \frac{1 - m(q, A(\ell))}{2} \mathbf{1}(\tau_\ell = -1).$$

Thus, Algorithm 1 defines the wanted coupling.

Now, using the coupling defined by Algorithm 1, we show the contractive property expressed in inequality (1.57). Recall that for any pair of configurations $\sigma, \tau \in \Omega$, the set $D_{\sigma,\tau}$ denotes the set of sites $\ell \in [n]$ where discrepancies occur, that is $\sigma_\ell \neq \tau_\ell$.

Sample the random variables according to Algorithm 1: sample $U \in [0, 1], \ell, k \in [n]$, then derive the random configurations $\sigma^{p(i)}, \sigma^{q(i)}$ and τ^i with $i \in \{1, 2\}$, as prescribed. For simplicity, we denote the random probability acceptance $p^i_J(\ell, k) = p_J(\ell, k \mid \sigma^{p(i)}, \sigma^{q(i)})$ for $i \in \{1, 2\}$, then we have the following relation

$$\begin{aligned}
|D_{\tau^1,\tau^2}| - |D_{\sigma^{p(1)},\sigma^{p(2)}}| \leq &-\mathbf{1}\Big(U \leq \tfrac{1}{2} - \overline{J}\Big) \mathbf{1}\Big(\ell \in D_{\sigma^{p(1)},\sigma^{p(2)}}\Big) && \text{(A)} \\
&+ \mathbf{1}\Big(U \in \Big[\tfrac{1}{2} - \overline{J}, \tfrac{1}{2} + \overline{J}\Big]\Big) \mathbf{1}(k \in D_{\sigma^{q(1)},\sigma^{q(2)}}) && \text{(B)} \\
&+ \mathbf{1}\Big(\ell \notin D_{\sigma^{p(1)},\sigma^{p(2)}},\ k \notin D_{\sigma^{q(1)},\sigma^{q(2)}}, \\
&\quad U \in (p^1_J(\ell, k) \wedge p^2_J(\ell, k),\ p^1_J(\ell, k) \vee p^2_J(\ell, k))\Big) && \text{(C)}
\end{aligned}$$

where $a \vee b = \max\{a, b\}$ and $a \wedge b = \min\{a, b\}$. The previous inequality can be derived by a check of cases imposed by Algorithm 1. Since any new discrepancy in the pair τ^1, τ^2 can only be generated by the update at site $\ell \in [n]$,

$$|D_{\tau^1,\tau^2}| - |D_{\sigma^{p(1)},\sigma^{p(2)}}| = \mathbf{1}(\ell \in D_{\tau^1,\tau^2}) - \mathbf{1}(\ell \in D_{\sigma^{p(1)},\sigma^{p(2)}}). \qquad (1.61)$$

The inequality (1.60) can be verified by a case-by-case analysis of the difference on right-hand side of (1.61). We have the following three cases based on output of U:

- $\boxed{U \le \frac{1}{2} - \overline{J}}$

  By Proposition 1.25, $p_J^1(\ell,k) \wedge p_J^2(\ell,k) \ge U$ and therefore the update occurs in both cases: $\tau^1 = S_\ell(\sigma^{p(1)}, \sigma_k^{q(1)})$, $\tau^2 = S_\ell(\sigma^{p(2)}, \sigma_k^{q(2)})$. Since the algorithm prescribes $\sigma^{q(1)} = \sigma^{q(2)}$, we have $\ell \notin D_{\tau^1,\tau^2}$, which yields the contribution (A).

- $\boxed{U > \frac{1}{2} + \overline{J}}$

  Again, by Proposition 1.25, $p_J^1(\ell,k) \vee p_J^2(\ell,k) \le U$, so the update does not occur in either case and $\tau^1 = \sigma^{p(1)}$, $\tau^2 = \sigma^{p(2)}$. Hence, $\ell \in D_{\tau^1,\tau^2}$ if and only if $\ell \in D_{\sigma^{p(1)},\sigma^{p(2)}}$. In particular,

  $$\mathbf{1}(\ell \in D_{\tau^1,\tau^2}) - \mathbf{1}(\ell \in D_{\sigma^{p(1)},\sigma^{p(2)}}) = 0.$$

  Therefore, this case gives no contribution.

- $\boxed{U \in \left[\frac{1}{2} - \overline{J}, \frac{1}{2} + \overline{J}\right]}$

  In this case we use the bound

  $$|D_{\tau^1,\tau^2}| - |D_{\sigma^{p(1)},\sigma^{p(2)}}| \le \mathbf{1}(\ell \in D_{\tau^1,\tau^2},\, \ell \notin D_{\sigma^{p(1)},\sigma^{p(2)}}).$$

  If $k \in D_{\sigma^{q(1)},\sigma^{q(2)}}$, we simply bound this term by 1, which yields (B).

  Otherwise, if $k \notin D_{\sigma^{q(1)},\sigma^{q(2)}}$, a discrepancy appears if and only if exactly one update occurs while the other is rejected, that is,

  $$U \in \Big(p_J^1(\ell,k) \wedge p_J^2(\ell,k), p_J^1(\ell,k) \vee p_J^2(\ell,k)\Big) \subset \Big[\frac{1}{2} - \overline{J}, \frac{1}{2} + \overline{J}\Big],$$

  which gives the contribution (C).

To conclude the proof of (1.57), we estimate the expectation of $|D_{\tau^1,\tau^2}| - |D_{\sigma^{p(1)},\sigma^{p(2)}}|$, giving separately an upper bound for each term (A),(B) and (C). By the algorithm, $\sigma^{p(1)}, \sigma^{p(2)} \sim \pi^p_{\mathrm{opt}}$ and $\sigma^{q(1)}, \sigma^{q(2)} \sim \pi^q_{\mathrm{opt}}$ whenever $U \ge \frac{1}{2} - \overline{J}$. Hence,

$$\mathbb{E}[(\mathrm{A})] = \frac{1}{n}\Big(\frac{1}{2} - \overline{J}\Big) W_1(p(1), p(2)), \qquad \mathbb{E}[(\mathrm{B})] \le \frac{2\overline{J}}{n} W_1(q(1), q(2)).$$

We will prove that

$$\mathbb{E}[(\mathrm{C})] \le \frac{\overline{J}}{n}\Big(W_1(p(1), p(2)) + W_1(q(1), q(2))\Big). \tag{1.62}$$

Before giving a proof of estimate (1.62), we combine the estimates for (A), (B), and (C) and conclude the proof of inequality (1.57). By construction, the pair τ^1, τ^2 produced by

the algorithm is a coupling of $\widehat{T}(p(1),q(1))$ and $\widehat{T}(p(2),q(2))$, then

$$\begin{aligned} W_1(\widehat{T}(p(1),q(1)),\widehat{T}(p(2),q(2))) &\le \mathbb{E}[D_{\tau^1,\tau^2}] \\ &\le \mathbb{E}[D_{\sigma^{p(1)},\sigma^{p(2)}}] - \mathbb{E}[(\mathrm{A})] + \mathbb{E}[(\mathrm{B})] + \mathbb{E}[(\mathrm{C})] \\ &= \Big(1 - \frac{1}{n}(\frac{1}{2} - 2\overline{J})\Big) W_1(p(1),p(2)) + \frac{3\overline{J}}{n} W_1(q(1),q(2)). \end{aligned}$$

This concludes the proof of inequality (1.57).

It remains to prove the estimate (1.62) for term (C). By definition

$$\begin{aligned} \mathbb{E}[(\mathrm{C})] = \frac{1}{n} \sum_{\ell,k\in[n]} \frac{\mathbf{1}(k\in A(\ell))}{|A(\ell)|} \mathbb{E}\Big[\mathbf{1}(\ell \notin D_{\sigma^{p(1)},\sigma^{p(2)}},\, k \notin D_{\sigma^{q(1)},\sigma^{q(2)}}) \\ \times \mathbb{P}[U \in (p_J^1(\ell,k) \wedge p_J^2(\ell,k),\, p_J^1(\ell,k) \vee p_J^2(\ell,k))]\Big] \end{aligned}$$

We use the following consequence of the Lipschitz property stated in Proposition 1.25: once $\sigma^{p(1)}, \sigma^{q(1)}, \sigma^{p(2)}, \sigma^{q(2)}$ and $\ell, k \in [n]$ are fixed,

$$\mathbb{P}[U \in (p_J^1(\ell,k) \wedge p_J^2(\ell,k),\, p_J^1(\ell,k) \vee p_J^2(\ell,k))] \le \frac{1}{4} |E_{\ell,k}(\sigma^{p(1)},\sigma^{q(1)}) - E_{\ell,k}(\sigma^{p(2)},\sigma^{q(2)})|.$$

Since the event (C) is supported on $\ell \notin D_{\sigma^{p(1)},\sigma^{p(2)}}$ and $k \notin D_{\sigma^{q(1)},\sigma^{q(2)}}$, we have $\sigma_\ell^{p(1)} - \sigma_k^{q(1)} = \sigma_\ell^{p(2)} - \sigma_k^{q(2)}$. In particular,

$$\begin{aligned} |E_\ell(\sigma^{p(1)},\sigma^{q(1)}) - E_\ell(\sigma^{p(2)},\sigma^{q(2)})| &= 2\mathbf{1}(\sigma_\ell^{p(1)} \ne \sigma_k^{q(1)}) \Big| [J(\sigma^{p(1)} - \sigma^{p(2)})]_\ell + [J(\sigma^{q(2)} - \sigma^{q(1)})]_k \Big| \\ &\le 4\Big[\sum_{s:s\ne\ell} |J_{\ell,s}| \mathbf{1}(s \in D_{\sigma^{p(1)},\sigma^{p(2)}}) + \sum_{s:s\ne k} |J_{k,s}| \mathbf{1}(s \in D_{\sigma^{q(1)},\sigma^{q(2)}}) \Big]. \end{aligned} \tag{1.63}$$

Using in order (1.63), $\frac{\mathbf{1}(k\in A(\ell))}{|A(\ell)|} = \frac{\mathbf{1}(\ell\in A(k))}{|A(k)|}$ and $\sum_{s:s\ne i} |J_{i,s}| \le \overline{J}$ for all $i \in [n]$, we conclude

$$\begin{aligned} \mathbb{E}[(\mathrm{C})] &\le \frac{1}{n} \sum_{\ell\in[n]} \sum_{s:s\ne\ell} |J_{\ell,s}| \mathbb{P}(s \in D_{\sigma^{p(1)},\sigma^{p(2)}}) + \frac{1}{n} \sum_{k\in[n]} \sum_{s:s\ne k} |J_{k,s}| \mathbb{P}(s \in D_{\sigma^{q(1)},\sigma^{q(2)}}) \\ &\le \frac{\overline{J}}{n} \Big(\sum_{s\in[n]} \mathbb{P}(s \in D_{\sigma^{p(1)},\sigma^{p(2)}}) + \sum_{s\in[n]} \mathbb{P}(s \in D_{\sigma^{q(1)},\sigma^{q(2)}}) \Big) \\ &= \frac{\overline{J}}{n} \Big(W_1(p(1),p(2)) + W_1(q(1),q(2)) \Big). \end{aligned}$$

□

1.4 Relative entropy decay for the particle system

This section is concerned with the proof of Theorem 1.10. We start with a version of that theorem that covers the case of the mean field exchange dynamics in Example 1.2, that is for the kernel $\mathcal{K}(\ell, n) \equiv 1/n$, and then turn to the proof for the general multi-component case in Definition 1.3. A key intermediate step in these proofs will be a modified log-Sobolev inequality for a version of the so-called Down-Up walk for the canonical Ising measure.

1.4.1 Mean field exchange dynamics

Here we state the main result for the kernel $\mathcal{K}(\ell, n) \equiv 1/n$. In order to prepare for later applications to multi-component systems we need to extend the setup from Section 1.1.3 to include arbitrary external fields. Thus, for $h \in \mathbb{R}^{Nn}$, we write $h = (h(i), i \in [N])$, and $h(i) \in \mathbb{R}^n$ for all $i \in [N]$, and define

$$\mu_{N,J,h,\alpha} := \otimes_{i=1}^{N} \mu_{J,h(i)}(\cdot \mid \Omega_{N,\alpha}), \tag{1.64}$$

the measure obtained by taking products of $\mu_{J,h(i)}$ and then restricting to $\Omega_{N,\alpha}$. Recall the notation $\sigma^{i,\ell;j,k}$ for the configuration $\sigma \in \Omega^N$ obtained by exchanging the spins $\sigma_\ell(i)$ and $\sigma_k(j)$; see (1.12). We use the notation

$$r_{J,h,i,j}(\ell, k \mid \sigma(i), \sigma(j)) = \frac{\mu_{N,J,h,\alpha}(\sigma^{i,\ell;j,k})}{\mu_{N,J,h,\alpha}(\sigma) + \mu_{N,J,h,\alpha}(\sigma^{i,\ell;j,k})}\,.$$

Note that, because of the product structure, this quantity depends on σ only through $\sigma(i)$ and $\sigma(j)$. Moreover, it coincides with (1.10) when $h \equiv 0$. We define the Dirichlet form

$$\bar{\mathcal{E}}_{N,J,h,\alpha}(F, G) := \frac{1}{2Nn} \sum_{i,j=1}^{N} \sum_{\ell,k=1}^{n} \mu_{N,J,h,\alpha}\left[r_{J,h,i,j}(\ell, k \mid \sigma(i), \sigma(j)) \nabla^{i,\ell;j,k} F \nabla^{i,\ell;j,k} G \right],$$

for functions $F, G : \Omega_{N,\alpha} \to \mathbb{R}$. As usual, we let

$$\mathrm{Ent}_{\mu_{N,J,h,\alpha}}(F) = \mu_{N,J,h,\alpha}\left[F \log(F/\mu_{N,J,h,\alpha}[F])\right]$$

denote the entropy of $F \geq 0$ with respect to $\mu_{N,J,h,\alpha}$.

Theorem 1.26. *Let $h \in \mathbb{R}^{Nn}$, let $J \in \mathbb{R}^{n\times n}$ be a nonnegative definite matrix with largest eigenvalue $\lambda(J) < 1/2$, and let $\overline{J} = \max_{i\in[n]} \sum_{j\in[n]} |J_{i,j}|$, $\overline{h} = \max_{\ell\in[n], i\in[N]} |h_\ell(i)|$. For any $N \geq 2$, and $\alpha \in \mathcal{D}_{\mathcal{A},N}$,*

$$\bar{\mathcal{E}}_{N,J,h,\alpha}(F, \log F) \geq r\, \mathrm{Ent}_{\mu_{N,J,h,\alpha}}(F), \tag{1.65}$$

for all $F \geq 0$, with constant α such that

$$r \geq \tfrac{1}{4}(1 - 2\lambda(J))\, e^{-8(\overline{J}+\overline{h})}.$$

We note that Theorem 1.26 implies Theorem 1.10 in the special case $\mathcal{K}(\ell, n) \equiv 1/n$, since $(1 - 2\lambda(J)) \leq 1$ and for this choice of kernel, taking $h \equiv 0$, one has

$$\bar{\mathcal{E}}_{N,J,0,\alpha}(F, G) = n\mathcal{E}_{N,J,\alpha}(F, G)\,,$$

where $\mathcal{E}_{N,J,\alpha}(F, G)$ is given in (1.15) with $\mathcal{K}(\ell, n) \equiv 1/n$. The proof of Theorem 1.26 will be derived as a consequence of the following more general estimate for mean field exchange systems.

1.4.2 Down-Up Walk for Conservative Particle Systems

We introduce some new notation. Consider a generic Ising spin system, with L spins, for some $L \in \mathbb{N}$, with interaction given by an $L \times L$ symmetric matrix Λ and with external field $w \in \mathbb{R}^L$, that is

$$\mu_{\Lambda,w}(\eta) = \frac{1}{Z_{\Lambda,w}} \exp\left(\frac{1}{2}\langle \eta, \Lambda\eta\rangle + \langle \eta, w\rangle\right), \qquad \eta \in \{-1, 1\}^L\,, \tag{1.66}$$

where $\langle \eta, \Lambda\eta\rangle = \sum_{i,j=1}^L \eta_i \Lambda_{i,j} \eta_j$ and $\langle \eta, w\rangle = \sum_{i=1}^L w_i\eta_i$. Given an integer $M \in \{-L, \dots, L\}$, consider the state space

$$\Omega_{L,M} = \left\{\eta \in \{-1, 1\}^L,\ \textstyle\sum_{i=1}^L \eta_i = M\right\}.$$

We will assume that $L + M$ is even, so that $\Omega_{L,M}$ is not empty. We view a configuration $\eta \in \Omega_{L,M}$ as composed of K points in the set $[L] = \{1, \dots, L\}$, corresponding to the sites i where $\eta_i = +1$. Thus, η is a configuration of $K = (L + M)/2$ elements from $[L]$, which we view as sites occupied by "balls". The remaining sites are such that $\eta_i = -1$ and will be seen as empty or "holes". The Down-Up Walk on $\Omega_{L,M}$ for the Ising system $\mu_{\Lambda,w}$ is the continuous-time Markov chain where a ball at site i, with rate 1 is moved to a hole at j, chosen among all possible vacant locations (including going back to i itself) with probability proportional to $\mu(\eta^{i,j})$, where $\eta \to \eta^{i,j}$ is the configuration obtained by swapping the spins values at sites i and j. Using the notation $\nabla^{i,j}F(\eta) = F(\eta^{i,j}) - F(\eta)$, the infinitesimal generator of the process is

$$\mathcal{L}^{\mathrm{du}}F(\eta) = \sum_{i,j=1}^{L} c^{\mathrm{du}}(\eta, \eta^{i,j})\nabla^{i,j}F(\eta), \tag{1.67}$$

with $F : \Omega_{L,M} \mapsto \mathbb{R}$, and $c^{\mathrm{du}}(\eta, \eta^{i,j})$ is the transition rate

$$c^{\mathrm{du}}(\eta, \eta^{i,j}) = \mathbf{1}(\eta_i = +1, \eta_j = -1)\frac{\mu_{\Lambda,w}(\eta^{i,j})}{\sum_{k\in U_i(\eta)} \mu_{\Lambda,w}(\eta^{ik})}, \tag{1.68}$$

where $U_i(\eta) = \{j \in [L] : \eta_j = -1\} \cup \{i\}$ denotes the set of available sites. This generator defines a conservative dynamics, reversible with respect to $\mu_{\Lambda,w}$. It is not hard to check

that, for any M such that $L+M$ is even, the process is irreducible in the state space $\Omega_{L,M}$, and the restricted measure

$$\nu_{L,M} = \mu_{\Lambda,w}\left(\cdot \mid \Omega_{L,M}\right) \tag{1.69}$$

is reversible, and thus the unique invariant measure. Let $\mathcal{E}^{\mathrm{du}}(F,G)$ denote the associated Dirichlet form,

$$\mathcal{E}^{\mathrm{du}}(F,G) = \frac{1}{2}\sum_{i,j=1}^{L} \nu_{L,M}\left[c^{\mathrm{du}}(\eta,\eta^{i,j})\nabla^{i,j}F(\eta)\nabla^{i,j}G(\eta)\right], \tag{1.70}$$

for $F, G : \Omega_{L,M} \to \mathbb{R}$. Let also $\mathrm{Ent}_{\nu_{L,M}}(F)$ denote the entropy with respect to the canonical Ising measure $\nu_{L,M}$. The key to prove Theorem 1.26 will be the following MLSI for the Down-Up walk.

Theorem 1.27 (MLSI for the Down-Up walk)**.** *Suppose that Λ is a nonnegative definite matrix with largest eigenvalue $\lambda(\Lambda) < 1/2$. Then for all values of L, M, with $L+M$ even, for all $w \in \mathbb{R}^L$,*

$$\mathcal{E}^{\mathrm{du}}(F, \log F) \geq (1 - 2\lambda(\Lambda))\,\mathrm{Ent}_{\nu_{L,M}}(F)\,,$$

for every $F : \Omega_{L,M} \to \mathbb{R}_+$.

It is crucial that the estimate is uniform over the parameters L, M and the external fields $w \in \mathbb{R}^L$. The above theorem can be proven using the arguments in Theorem 2.2 and Lemma 2.7 in [BBD24a]. Moreover, following [BBD24a] it can be shown that if the largest eigenvalue of Λ is associated to a constant eigenvector, then one can impose the condition $\lambda(\Lambda) < 1/2$ on the second largest eigenvalue. This is natural, in light of the global constraint $\sum_i \eta_i = M$. Indeed, if e.g. $\Lambda_{i,j} \equiv c$, then the canonical measure $\nu_{L,M}$ becomes uniform and one might as well take $\Lambda = 0$. However, for the applications to our setting, the removal of the eigenvalue associated to constant eigenfunction will not be relevant. Below, we prove Theorem 1.27 with a slightly different argument, which will allow us to introduce the relevant tools to prove the more general statements for multi-component mean field exchange models, which are not covered by the results in [BBD24a]. Our proof will follow the factorization technique developed in [CCP25], which in turn was based on the stochastic localization approach from [CE22]. The proof of Theorem 1.27 will be given in the following subsections. Here we observe that it implies Theorem 1.26.

Proof of Theorem 1.26 assuming Theorem 1.27. Let $L = Nn$, where N is the number of particles and n the number of spins per particle. We use the correspondence

$$\eta \in \{-1,1\}^{Nn} \longrightarrow \sigma = (\sigma(i), i = 1,\dots,N), \quad \sigma(i) = (\eta_u,\, u = n(i-1)+1,\dots,ni), \tag{1.71}$$

that is the spins in the block of sites $\{(i-1)n+1,\dots,ni\}$ are associated with the i-th particle $\sigma(i) = (\sigma_1(i),\dots,\sigma_n(i))$. The $L \times L$ matrix Λ and the L-vector w are taken such that

$$\langle \eta, \Lambda\eta\rangle = \textstyle\sum_{i=1}^N \langle \sigma(i), J\sigma(i)\rangle, \qquad \langle \eta, w\rangle = \textstyle\sum_{i=1}^N \langle \sigma(i), h(i)\rangle$$

where the scalar products on the left hand side are over $\mathbb{R}^L$ while each of the scalar products on the right hand side is over $\mathbb{R}^n$. In this way, taking $M = (2\alpha - 1)Nn$, the canonical measure $\nu_{L,M}$ in (1.69) coincides with the probability measure $\mu_{N,J,h,\alpha}$. Since Λ consists of N blocks given by the $n \times n$ matrix J, one has $\lambda(\Lambda) = \lambda(J)$. Therefore, to prove (1.65) it remains to establish the comparison

$$\bar{\mathcal{E}}_{N,J,h,\alpha}(F, \log F) \geq \tfrac{1}{4}\, e^{-8(\overline{J}+\overline{h})} \mathcal{E}^{\mathrm{du}}(F, \log F), \tag{1.72}$$

for all nonnegative functions F. To prove this, observe that, if $M \leq 0$, that is $\alpha \leq 1/2$, then $|U_i(\eta)| \geq L/2 = Nn/2$, for all $\eta \in \Omega_{L,M}$, and therefore

$$\begin{aligned} c^{\mathrm{du}}(\eta, \eta^{i,j}) &\leq \frac{4}{Nn}\, e^{8(\overline{J}+\overline{h})} \frac{\mu_{\Lambda,w}(\eta^{i,j})}{\mu_{\Lambda,w}(\eta) + \mu_{\Lambda,w}(\eta^{i,j})} \\ &= \frac{4}{Nn}\, e^{8(\overline{J}+\overline{h})}\, r_{J,h,C_i,C_j}(\ell, k \mid \sigma(C_i), \sigma(C_j))\,, \end{aligned}$$

where $C_i, C_j \in [N]$, and $\ell, k \in [n]$, denote the labels such that $\sigma_\ell(C_i) = \eta_i$ and $\sigma_k(C_j) = \eta_j$ in the correspondence (1.71). Using this pointwise estimate in (1.70) we obtain (1.72). This ends the proof in the case $M \leq 0$. The case $M \geq 0$ can be handled in the same way with the change of variable $\eta \mapsto -\eta$. Indeed, this global flip changes the measure from $\mu_{N,J,h,\alpha}$ to $\mu_{N,J,-h,1-\alpha}$. Thus we may apply the previous argument to the new variables. This ends the proof of Theorem 1.26. □

1.4.3 Proof of Theorem 1.27

Let $(W_t)_{t\geq 0}$ be a standard Brownian motion in $\mathbb{R}^L$, adapted to a filtration $(\mathcal{F}_t)_{t\geq 0}$, and set $\tilde{W}_t = \Lambda^{1/2} W_t$. The stochastic localization process for the canonical Ising measure $\nu_{L,M}$ is given by

$$\varrho_t(\eta) = \frac{1}{Z_t} \exp\left(-\frac{t-1}{2}\langle \eta, \Lambda\eta\rangle + \langle y_t + w, \eta\rangle\right) \mathbf{1}(\eta \in \Omega_{L,M}), \quad t \geq 0\,, \tag{1.73}$$

where Z_t is the normalization so that $\varrho_t \in \mathcal{P}(\Omega_{L,M})$, and the random field $y_t \in \mathbb{R}^L, t \geq 0$ evolves according to the stochastic differential equation

$$dy_t = d\tilde{W}_t + \Lambda m(y_t)\, dt, \qquad y_0 = 0,$$

with $m(y_t)$ given by the mean vector $m(y_t) := \varrho_t[\eta]$, that is for all $i \in [L]$,

$$[m(y_t)]_i = \varrho_t[\eta_i] = \textstyle\sum_{\eta\in\Omega_{L,M}} \eta_i \varrho_t(\eta).$$

That the process is well defined can be checked by standard methods. Moreover, $(\varrho_t, t \geq 0)$ is an $\{\mathcal{F}_t\}$-martingale and therefore $\nu_{L,M} = \mathbb{E}[\varrho_t]$ for all $t \geq 0$, where $\mathbb{E}$ denotes the expectation with respect to the stochastic localization process. We refer to [CE22]; see also Lemma 3.1 in [CCP25].

At time $t = 1$, the measure (1.73) is a random product distribution restricted to $\Omega_{L,M}$,

$$\varrho_1(\eta) \propto \prod_{i=1}^{L} \exp\left(y_{1,i}\eta_i + w_i\eta_i\right) \mathbf{1}(\eta \in \Omega_{L,M}), \tag{1.74}$$

where $y_{1,i}$ denotes the i-th component of the vector y_1. That is, ϱ_1 is a Bernoulli product measure conditioned on having total magnetization M. Such measures are known to satisfy the Strong Rayleigh property; see [BBL09]. This implies strong log-concavity of its generating polynomial, which in turn yields a modified logarithmic Sobolev inequality (MLSI) for the associated Down-Up walk [CGM21]; see also [HS23; Ana+22; BBD24a]. We shall actually establish a stronger estimate, namely entropy factorization. This allows us to use the arguments in [CCP25] which yield a simple proof of Theorem 1.27.

To introduce the factorization estimate, we formulate the Down-Up walk for an arbitrary measure $\mu \in \mathcal{P}(\Omega_{L,M})$ as follows. As above, we say that there is a ball at $j \in [L]$ if $\eta_j = +1$. Thus, every $\eta \in \Omega_{L,M}$ has $K := (L+M)/2$ balls. According to the distribution μ, balls are not distinguished. However, we artificially distinguish them by taking a uniformly random labeling of the K balls, and write $X_k \in [L]$ for the position of the k-th ball. With slight abuse of notation we call again μ the induced distribution of the vector $X := (X_1, \dots, X_K)$. Let $\mu_i = \mu(\cdot \mid X_j,\, j \neq i)$ denote the distribution of X_i conditionally on the position of all remaining balls. Now, it is not hard to see that the Down-Up walk for the original measure μ, defined by the rates (1.68) with μ in place of $\mu_{\Lambda,w}$, coincides with the Markov chain where with rate 1 independently each variable X_i is updated by replacing it with a sample from μ_i, provided we observe the chain on functions that are symmetric under permutations of the ball labels. Thus, letting $\mathcal{S}$ be the set of symmetric functions $X \mapsto F(X)$, the generator of the chain becomes

$$\mathcal{L}^{\mathrm{du}} F(X) = \sum_{i=1}^{K} (\mu_i[F](X) - F(X)),$$

for any $F \in \mathcal{S}$, where $\mu_i[F](X) = \mu(F \mid X_j,\, j \neq i)$ is the conditional expectation of F given all X_j, $j \neq i$. It follows that the Dirichlet form is given by

$$\mathcal{E}^{\mathrm{du}}(F, G) = \sum_{i=1}^{K} \mu\left[\mathrm{Cov}_{\mu_i}(F, G)\right], \tag{1.75}$$

for any $F, G \in \mathcal{S}$, where $\mathrm{Cov}_{\mu_i}(F, G) = \mu_i[FG] - \mu_i[F]\mu_i[G]$ denotes the covariance with respect to μ_i. Moreover, let $\mathrm{Ent}_{\mu_i}(F) = \mu_i[F \log F] - \mu_i[F] \log \mu_i[F]$ denote the entropy functional with respect to μ_i. Note that this is a function of the variables X_j, $j \neq i$ and that for any realization of these variables Jensen's inequality implies

$$\mathrm{Ent}_{\mu_i} F \leq \mathrm{Cov}_{\mu_i}(F, \log F), \tag{1.76}$$

for all $F \geq 0$. In light of (1.75)-(1.76) the following factorization property implies the MLSI

for μ with constant 1. We say that $\mu \in \mathcal{P}(\Omega_{L,M})$ is strongly log-concave if its generating polynomial is strongly log-concave; see [CGM21].

Proposition 1.28. *Suppose $\mu \in \mathcal{P}(\Omega_{L,M})$ is strongly log-concave. Then, for any $F \in \mathcal{S}$,*

$$\operatorname{Ent}_\mu F \le \sum_{i=1}^{K} \mu \left[\operatorname{Ent}_{\mu_i} F\right]. \tag{1.77}$$

Proof. From [CGM21, Lemma 11], or equivalently [Ana+22, Theorem 5], any strongly log-concave μ satisfies the subadditivity inequality,

$$\operatorname{Ent}_\mu F \ge \sum_{i=1}^{K} \operatorname{Ent}_\mu F_i, \qquad F_i = \mu[F \mid X_i], \tag{1.78}$$

for any nonnegative symmetric function F. Note that by symmetry $\operatorname{Ent}_\mu F_i = \operatorname{Ent}_\mu F_1$ for all i. Moreover, it is known that the inequality (1.78) continues to hold with μ replaced by the conditional distribution $\mu(\cdot \mid X_1, \dots, X_k)$ for any $k = 1, \dots, K-1$ and any realization of the variables X_j, $j = 1, \dots, k$. The factorization statement (1.77) then follows from a simple recursive argument, see e.g. [Cap23, Section 4.1] □

Since the measure ϱ_1 in (1.74) is strongly log-concave [BBL09], we may apply Theorem 1.28 with $\mu = \varrho_1$. Taking the expectation with respect to the stochastic localization process,

$$\mathbb{E}\left[\operatorname{Ent}_{\varrho_1} F\right] \le \sum_{i=1}^{K} \mathbb{E}\left[\varrho_1\left[\operatorname{Ent}_{\varrho_{1,i}} F\right]\right], \tag{1.79}$$

where $\varrho_{1,i} = \varrho_1(\cdot \mid X_j,\, j \ne i)$. We would like to turn the statement (1.79) into a factorization statement for our original measure $\nu := \nu_{L,M} = \mathbb{E}[\varrho_1]$. The first step is a convexity estimate.

Lemma 1.29. *For any nonnegative symmetric function F, for all $i \in [K]$, one has*

$$\mathbb{E}\left[\varrho_1\left[\operatorname{Ent}_{\varrho_{1,i}} F\right]\right] \le \nu\left[\operatorname{Ent}_{\nu_i} F\right]. \tag{1.80}$$

Proof. Using the martingale property $\nu = \mathbb{E}[\varrho_1]$,

$$\mathbb{E}\left[\varrho_1\left[\operatorname{Ent}_{\varrho_{1,i}} F\right]\right] = \nu[F \log F] - \mathbb{E}\left[\varrho_1\left[\varrho_{1,i} F \log \varrho_{1,i} F\right]\right]. \tag{1.81}$$

We write

$$\varrho_{1,i} F(X) = \frac{\varrho_1(F; X_{\ne i})}{\varrho_1(X_{\ne i})}, \tag{1.82}$$

where we use the notation $X_{\ne i} = (X_j,\, j \ne i)$, $\varrho_1(X_{\ne i}) = \sum_{X_i} \varrho_1(X)$, and $\varrho_1(F; X_{\ne i}) = \sum_{X_i} F(X)\varrho_1(X)$, and therefore

$$\mathbb{E}\left[\varrho_1\left[\varrho_{1,i} F \log \varrho_{1,i} F\right]\right] = \sum_{X_{\ne i}} \mathbb{E}\left[\varrho_1(F; X_{\ne i}) \log \frac{\varrho_1(F; X_{\ne i})}{\varrho_1(X_{\ne i})}\right].$$

Using $\nu = \mathbb{E}[\varrho_1]$, and the joint convexity of the function $(x, y) \mapsto x \log(x/y)$ on $(0, \infty)^2$, by Jensen's inequality

$$\mathbb{E}\left[\varrho_1 \left[\varrho_{1,i} F \log \varrho_{1,i} F\right]\right] \geq \sum_{X_{\neq i}} \nu(F; X_{\neq i}) \log \frac{\nu(F; X_{\neq i})}{\nu(X_{\neq i})} = \nu\left[\nu_i F \log \nu_i F\right] ,$$

where we use the notation (1.82) for ν in place of ϱ_1. The desired bound (1.80) follows from (1.81). □

Next, we use the so-called entropic stability estimate from [CE22], which allows one to compare relative entropies with respect to ϱ_1 with relative entropies with respect to ν by using a covariance estimate. Let

$$\mathrm{Cov}[\mu](i, j) = \mu[\eta_i \eta_j] - \mu[\eta_i]\mu[\eta_j],$$

denote the $L \times L$ covariance matrix of any $\mu \in \mathcal{P}(\Omega_{L,M})$, and write $\lambda(\mathrm{Cov}[\mu])$ for its largest eigenvalue. If $v \in \mathbb{R}^L$, the tilted measure $\mathcal{T}_v \mu \in \mathcal{P}(\Omega_{L,M})$ is defined as

$$\mathcal{T}_v \mu(\eta) \propto \mu(\eta) e^{\langle v, \eta \rangle} .$$

The next lemma is a version of [CE22, Proposition 39]; see also [CCP25, Lemma 3.2]. The proof can be adapted to our conservative setting without modifications.

Proposition 1.30. *Suppose there exists a deterministic function $\kappa : [0, 1] \to \mathbb{R}_+$ such that $\lambda(\mathrm{Cov}[{}_v\alpha_t]) \leq \kappa(t)$ a.s., for all $t \in [0, 1]$, for all $v \in \mathbb{R}^L$. Then, for all functions $F \geq 0$,*

$$\mathbb{E}\left[\mathrm{Ent}_{\alpha_1} F\right] \geq \exp\left(-\lambda(\Lambda) \int_0^1 \kappa(s) ds\right) \mathrm{Ent}_\nu F .$$

We turn to a bound on the covariance. The following was established in [CE22, Lemma 50] for the non conservative setting, based on a decomposition introduced in [BB19]; see also [CCP25, Lemma 3.4]. We adapt the proof given in [CCP25] to the present conservative setting.

Proposition 1.31. *Let $\nu = \nu_{L,M}$ be the canonical Ising measure (1.69) with positive definite interaction matrix Λ and external fields $w \in \mathbb{R}^L$. If $\lambda(\Lambda) < 1/2$, then the largest eigenvalue of $\mathrm{Cov}[\nu]$ satisfies*

$$\lambda(\mathrm{Cov}[\nu]) \leq \frac{2}{1 - 2\lambda(\Lambda)} .$$

Proof. By gaussian integration, for any $\eta \in \Omega_{L,M}$,

$$\exp(\tfrac{1}{2}\langle \eta, \Lambda \eta \rangle) \propto \int_{\mathbb{R}^L} dy \exp\left(-\tfrac{1}{2}\langle y, \Lambda^{-1} y \rangle + \langle \eta, y \rangle\right), \tag{1.83}$$

where the implicit positive constant is independent of η. Define the function

$$f(y) := \exp\left(-\tfrac{1}{2}\langle y, \Lambda^{-1} y\rangle + \log\left(Z\left(y+w\right)\right)\right),$$

where $Z\left(v\right) := \sum_{\eta\in\Omega_{L,M}} \exp\left(\langle \eta, v\rangle\right)$, for any $v \in \mathbb{R}^L$. From (1.83) we obtain

$$\nu(\eta) \propto \int_{\mathbb{R}^L} dy f(y) Z\left(y+w\right)^{-1} \exp\left(\langle \eta, y+w\rangle\right) \mathbf{1}(\eta \in \Omega_{L,M}).$$

If μ denotes the uniform probability measure over $\Omega_{L,M}$ then for all $v \in \mathbb{R}^L$,

$$\mathcal{T}_v\mu(\eta) = Z\left(v\right)^{-1} \exp\left(\langle \eta, v\rangle\right) \mathbf{1}_{\{\eta\in\Omega_{L,M}\}}.$$

Therefore, ν is given by a mixture of tilted measures $\mathcal{T}_v\mu$,

$$\nu(\eta) = \int \mathcal{T}_{y+w}\mu(\eta)\hat{f}(y)dy,$$

where $\hat{f} := f/\int dy f(y)$ is a probability density on $\mathbb{R}^L$. By the law of total covariance,

$$\mathrm{Cov}[\nu] = \int dy \hat{f}(y)\, \mathrm{Cov}[\mathcal{T}_{y+w}\mu] \; + \mathcal{C}_\nu, \tag{1.84}$$

where $\mathcal{C}_\nu$ stands for the $L \times L$ matrix

$$\mathcal{C}_\nu(i,j) := \int \gamma_i(y)\gamma_j(y)\hat{f}(y)dy - \left(\int \gamma_i(y)\hat{f}(y)dy\right)\left(\int \gamma_i(y)\gamma_j(y)\hat{f}(y)dy\right),$$

and we use the notation $\gamma_i(y) := \sum_{\eta\in\Omega_{L,M}} \mathcal{T}_{y+w}\mu(\eta)\eta_i$, $i \in [L]$. Next, we observe that $\lambda(\mathrm{Cov}[\mathcal{T}_v\mu]) \le 2$ for any $v \in \mathbb{R}^L$. Indeed, $\mathcal{T}_v\mu$ is Strong Rayleigh [BBL09], which implies negative correlation, and thus for any $u \in \mathbb{R}^L$, letting $C = \mathrm{Cov}[\mathcal{T}_v\mu]$, one has

$$\sum_j |C_{i,j}| = C_{i,i} - \sum_{j:\, j\neq i} C_{i,j} = 2C_{i,i},$$

for any $i \in [L]$, where we use $C_{i,j} \le 0$ for $i \neq j$ and the fact that $\sum_j C_{i,j} = 0$. Since $\lambda(C) \le \max_i \sum_j |C_{i,j}|$ and $C_{i,i} \le 1$ for any $i \in [L]$ one has the desired bound $\lambda(\mathrm{Cov}[\mathcal{T}_v\mu]) \le 2$.

We turn to a bound on the matrix $\mathcal{C}_\nu$. For any $u \in \mathbb{R}^L$ we write

$$\langle u, \mathcal{C}_\nu u\rangle = \textstyle\sum_{i,j} u_i u_j \mathcal{C}_\nu(i,j) = \left(g_u(Y)\right),$$

where Y is the $\mathbb{R}^L$-valued random variable with density $\hat{f}$, and $g_u(y) := \sum_i u_i\gamma_i(y)$. We are going to use the Brascamp-Lieb inequality (Theorem 4.1 in [**brasclieb**]). Namely, suppose that $-\mathrm{Hess}\log\hat{f}(y)$ is positive definite, with lower bound $a\,\mathrm{Id}$ where $a > 0$ and Id is the $L \times L$ identity matrix. Then,

$$\left(g_u(X)\right) \le \frac{1}{a}\int dy \hat{f}(y)\|\nabla g_u(y)\|_2^2. \tag{1.85}$$

Note that,

$$\nabla_i g_u(y) = \sum_j u_j \frac{\partial \gamma_j(y)}{\partial y_i} = \sum_j u_i C_{i,j} \,,$$

where $C = \mathrm{Cov}[\mathcal{T}_{y+w}\mu]$. As discussed above, one has $\lambda(C) \leq 2$. Therefore, for all $y \in \mathbb{R}^L$,

$$\|\nabla g_u(y)\|_2^2 = \langle Cu, Cu\rangle \leq 4\langle u, u\rangle \,.$$

From (1.85) we conclude

$$\langle u, \mathcal{C}_\nu u\rangle \leq \frac{4}{a}\langle u, u\rangle \ .$$

It remains to establish the bound $-\mathrm{Hess}\log \hat{f}(y) \geq a\,\mathrm{Id}$. Since $\mathrm{Hess}\log\left(Z\left(y+v\right)\right) = \mathrm{Cov}[\mathcal{T}_{y+w}\mu]$, one has $\mathrm{Hess}\log\left(Z\left(y+w\right)\right) \leq 2\,\mathrm{Id}$. Moreover, $\mathrm{Hess}\langle y, \Lambda^{-1}y\rangle \geq 2\lambda(\Lambda)^{-1}\mathrm{Id}$. Therefore we obtain the quadratic form bound

$$-\mathrm{Hess}\log \hat{f}(y)) = \mathrm{Hess}\left(\tfrac{1}{2}\langle y, \Lambda^{-1}y\rangle - \log\left(Z\left(y+v\right)\right)\right) \geq \left(\frac{1}{\lambda(\Lambda)} - 2\right)\mathrm{Id}.$$

Summarizing, we may take $a = (1-2\lambda(\Lambda))/\lambda(\Lambda)$. Since $4/a + 2 = 2/(1-2\lambda(\Lambda))$, the desired conclusion follows from (1.84). □

We now have all ingredients to complete the proof of Theorem 1.27. From Proposition 1.31 applied with Λ replaced by $(1-t)\Lambda$, $t \in [0,1]$, and using the expression (1.73) for ϱ_t we see that the assumption of Proposition 1.30 holds with

$$\kappa(s) = \frac{2}{1 - 2(1-s)\lambda(\Lambda)}, \quad \forall s \in [0,1].$$

Integrating, $\int_0^1 \kappa(s)ds = -\frac{1}{\lambda(\Lambda)}\log(1-2\lambda(\Lambda))$. Therefore, Proposition 1.30 implies

$$\mathrm{Ent}_\nu F \leq \frac{1}{1-2\lambda(\Lambda)}\,\mathbb{E}\left[\mathrm{Ent}_{\varrho_1} F\right],$$

for all $F : \Omega_{L,M} \to \mathbb{R}_+$. This, together with (1.79) and (1.80), proves that

$$\mathrm{Ent}_\nu F \leq \frac{1}{1-2\lambda(\Lambda)} \sum_{i=1}^K \nu\left[\mathrm{Ent}_{\nu_i} F\right],$$

Using (1.76), it follows from (1.75) that

$$\mathrm{Ent}_\nu F \leq \frac{1}{1-2\lambda(\Lambda)}\,\mathcal{E}^{\mathrm{du}}(F, \log F)$$

This ends the proof of Theorem 1.27.

1.4.4 Multicomponent mean field exchange dynamics

We extend the previous analysis to the case where the system consists of several regions each of which is restricted to have a fixed magnetization. To this end, consider again the probability measure $\mu_{\Lambda,w}$ on $\{-1,+1\}^L$ given in (1.66). Let $\mathcal{B}$ denote a partition of $[L]$, that is $\cup_{B\in\mathcal{B}}B = [L]$ and $B \cap B' = \emptyset$ for all $B \neq B'$. Let also $\mathcal{M} = (M_B,\, B \in \mathcal{B})$, $M_B \in \{-|B|, \ldots, |B|\}$, denote the magnetization vector, and define

$$\Omega_{L,\mathcal{M}} := \left\{\eta \in \{-1,+1\}^L : \textstyle\sum_{i\in B}\eta_i = M_B, \quad \forall B \in \mathcal{B}\right\},$$

and the multicomponent canonical measure

$$\nu_{L,\mathcal{M}} := \mu_{\Lambda,w}\left(\cdot \mid \Omega_{L,\mathcal{M}}\right).$$

For every $B \in \mathcal{B}$ we define a Down-Up walk with rates

$$c_B^{\mathrm{du}}(\eta,\eta^{i,j}) = \mathbf{1}(i,j\in B)\mathbf{1}(\eta_i = +1, \eta_j = -1)\frac{\mu_{\Lambda,w}(\eta^{i,j})}{\sum_{k\in U_{B,i}(\eta)}\mu_{\Lambda,w}(\eta^{ik})}, \tag{1.86}$$

where $U_{B,i}(\eta) = \{j \in B : \eta_j = -1\} \cup \{i\}$. Consider now the Down-Up walk given by the generator (1.67) with rates

$$c^{\mathrm{du}}(\eta,\eta^{i,j}) := \sum_{B\in\mathcal{B}} c_B^{\mathrm{du}}(\eta,\eta^{i,j}).$$

It is not hard to check that, for any choice of the magnetization vector $\mathcal{M}$, the process is irreducible in the state space $\Omega_{L,\mathcal{M}}$, with unique reversible and invariant measure $\nu_{L,\mathcal{M}}$. The Dirichlet form is given by

$$\mathcal{E}_{\mathcal{B}}^{\mathrm{du}}(F,G) = \frac{1}{2}\sum_{B\in\mathcal{B}}\sum_{i,j\in B}\nu_{L,\mathcal{M}}\left[c_B^{\mathrm{du}}(\eta,\eta^{i,j})\nabla^{i,j}F(\eta)\nabla^{i,j}G(\eta)\right],$$

for $F, G : \Omega_{L,\mathcal{M}} \to \mathbb{R}$. Let also $\mathrm{Ent}_{\nu_{L,\mathcal{M}}}(F)$ denote the entropy with respect to the multi-canonical Ising measure $\nu_{L,\mathcal{M}}$. The multicomponent version of Theorem 1.27 is as follows.

Theorem 1.32 (MLSI for Multicomponent Down-Up walk)**.** *Let $\mathcal{B}$ be a fixed partition of $[L]$. Suppose that Λ is a nonnegative definite $L \times L$ matrix with largest eigenvalue $\lambda(\Lambda) < 1/2$. Then for all values of the vector $\mathcal{M}$, for all $w \in \mathbb{R}^L$,*

$$\mathcal{E}_{\mathcal{B}}^{\mathrm{du}}(F,\log F) \geq (1-2\lambda(\Lambda))^2\,\mathrm{Ent}_{\nu_{L,\mathcal{M}}}(F),$$

for every $F : \Omega_{L,\mathcal{M}} \to \mathbb{R}_+$.

The key to prove the above theorem is the following factorization statement. For any $B \in \mathcal{B}$ let $\nu_B = \nu_{L,\mathcal{M}}(\cdot \mid \eta_j, j \notin B)$ denote distribution inside B conditioned on the outside

variables, and let $\mathrm{Ent}_{\nu_B}(F) = \nu_B(F \log F) - \nu_B(F) \log \nu_B(F)$ be the associated entropy functional. Note that $\mathrm{Ent}_{\nu_B}(F)$ is a function of the outside variables $\eta_j, j \notin B$.

Proposition 1.33 (Approximate Factorization of Entropy)**.** *Under the assumptions of Theorem 1.32, for all* $F : \Omega_{L,\mathcal{M}} \to \mathbb{R}_+$,

$$\sum_{B \in \mathcal{B}} \nu_{L,\mathcal{M}} \left[\mathrm{Ent}_{\nu_B}(F) \right] \geq (1 - 2\lambda(\Lambda)) \, \mathrm{Ent}_{\nu_{L,\mathcal{M}}}(F) \, .$$

Proof. The proof can be obtained as in [CCP25, Theorem 2.8], which establishes a Shearer-type inequality for multicomponent systems. For the reader's convenience we sketch the main steps. Consider the measure ϱ_t defined as in (1.73) with the set $\Omega_{L,M}$ replaced by $\Omega_{L,\mathcal{M}}$. A key feature of the stochastic localization process in this case is that the measure ϱ_1 is now a product over $B \in \mathcal{B}$, namely

$$\varrho_1 = \bigotimes_{B \in \mathcal{B}} \varrho_{1,B} \, , \qquad \varrho_{1,B} =_{y_{1,B} + w_B} \mu_B \, ,$$

where μ_B is the uniform distribution over $\eta \in \{-1, +1\}^B$ such that $\sum_{j \in B} \eta_j = M_B$, and $y_{1,B}, w_B$ denote the projection of the L-vectors y_1, w onto B. This can be seen from (1.74) where now $\Omega_{L,M}$ is replaced by $\Omega_{L,\mathcal{M}}$.

By the tensorization property of product measures, see e.g. [Cap23], one has the factorization

$$\mathrm{Ent}_{\varrho_1} F \leq \sum_{B \in \mathcal{B}} \varrho_1 \left[\mathrm{Ent}_{\varrho_{1,B}} F \right] \, .$$

Moreover, we may apply the convexity argument in Lemma 1.29 to obtain, for any $B \in \mathcal{B}$,

$$\mathbb{E} \left[\varrho_1 \left[\mathrm{Ent}_{\varrho_{1,B}} F \right] \right] \leq \nu_{L,\mathcal{M}} \left[\mathrm{Ent}_{\nu_B} F \right] \, .$$

At this point we repeat the steps in Proposition 1.30 and Proposition 1.31. The fact that the measures ${}_v\mu_B$, $B \in \mathcal{B}$, are Strong Rayleigh for any $v \in \mathbb{R}^B$ implies that their product is also Strong Rayleigh, and thus the same covariance estimates can be used. In conclusion,

$$\sum_{B \in \mathcal{B}} \nu \left[\mathrm{Ent}_{\nu_B}(F) \right] \, \geq \, \mathbb{E}_\mu \left[\mathrm{Ent}_{\varrho_1}(F) \right] \, \geq \, (1 - 2\lambda(\Lambda)) \, \mathrm{Ent}_{\nu_{L,\mathcal{M}}}(F) \, .$$

□

Proof of Theorem 1.32. For every $B \in \mathcal{B}$, the measure ν_B is of the form of a canonical Ising measure as in (1.69) with $L = |B|$, Λ replaced by $\Lambda_B = (\Lambda_{i,j}, \, i, j \in B)$, the principal submatrix corresponding to the set B, and w replaced by some new external field $w(\eta_{B^c}) = (w_i(\eta_{B^c}), \, i \in B)$, which depends on the outside variables $\eta_{B^c} = (\eta_j, \, j \notin B)$ through

$$w_i(\eta_{B^c}) = w_i + \sum_{j \notin B} \Lambda_{i,j} \eta_j \, .$$

Thus we may apply Theorem 1.27 to this system. Noting that the rates $c_B^{\mathrm{du}}(\eta, \eta^{i,j})$ can be written as in (1.86) with $\mu_{\Lambda,w}$ replaced by ν_B, this gives, for all B, uniformly in η_{B^c},

$$\frac{1}{2}\sum_{i,j\in B}\nu_B\left[c_B^{\mathrm{du}}(\eta,\eta^{i,j})\nabla^{i,j}F(\eta)\nabla^{i,j}\log F(\eta)\right] \geq (1-2\lambda(\Lambda_B))\,\mathrm{Ent}_{\nu_B}(F)\,.$$

Summing over $B \in \mathcal{B}$, taking the expectation with respect to $\nu_{L,\mathcal{M}}$, and noting that $\lambda(\Lambda_B) \leq \lambda(\Lambda)$,

$$\mathcal{E}_{\mathcal{B}}^{\mathrm{du}}(F, \log F) \geq (1-2\lambda(\Lambda))\sum_{B\in\mathcal{B}}\nu_{L,\mathcal{M}}\left[\mathrm{Ent}_{\nu_B}(F)\right].$$

The statement in Theorem 1.32 then follows from Proposition 1.33. □

1.4.5 Proof of Theorem 1.10

We apply Theorem 1.32 with the notation given by the correspondence in (1.71), and with the following choice of blocks B. Given a partition $\mathcal{A}$ of $[n]$, we define the associated partition $\mathcal{B} = (B(A),\, A \in \mathcal{A})$ of $[Nn]$ where, for each $A \in \mathcal{A}$, $B(A)$ is the stripe $[N] \times A$, or more formally

$$B(A) := \{i \in [Nn] : i = \ell + qn,\ \ell \in A,\ q \in \mathbb{Z}_+,\ 1 \leq \ell + qn \leq Nn\}.$$

Note that $\{B(A), A \in \mathcal{A}\}$ forms a partition of $[Nn]$, and if $B = B(A)$, then η_B corresponds to the spins $(\sigma_\ell(i), \ell \in A, i \in [N])$. Let μ be the canonical Ising measure defined in (1.14), where $\alpha \in \mathcal{D}_{\mathcal{A},N}$, and for $B = B(A)$, write μ_B for the measure

$$\mu_B = \mu(\cdot \mid \sigma_\ell(i), \ell \notin A, i \in [N]),$$

that is μ_B is the measure obtained by conditioning on all variables outside of B. The factorization in Proposition 1.33 is equivalent to

$$\sum_{B\in\mathcal{B}}\mu\left[\mathrm{Ent}_{\mu_B}(F)\right] \geq (1-2\lambda(\Lambda))\,\mathrm{Ent}_{\mu}(F)\,. \tag{1.87}$$

Recalling that $\mathcal{K}$ is now given by $\mathcal{K}(\ell,k) = \frac{1}{|A|}$ if $\ell, k \in A$ and $\mathcal{K}(\ell,k) = 0$ otherwise, the Dirichlet form (1.15) becomes

$$\mathcal{E}_{N,J,\alpha}(F,G) = \frac{1}{n}\sum_{A\in\mathcal{A}}\mu\left[\mathcal{E}_{B(A)}(F,G)\right],$$

$$\mathcal{E}_{B(A)}(F,G) := \frac{1}{2N|A|}\sum_{i,j=1}^{N}\sum_{\ell,k\in A}\mu_{B(A)}\left[\hat{p}_J(\ell,k \mid \sigma(i),\sigma(j))\nabla^{i,\ell;j,k}F\nabla^{i,\ell;j,k}G\right],$$

for functions $F, G : \Omega_{N,\alpha} \to \mathbb{R}$. We note that every μ_B is of the form (1.64), with (n, J, h) replaced by $(|A|, J_A, h_A)$, where $J_A = (J_{i,j},\, i,j \in A)$, and $h_A \in \mathbb{R}^{N|A|}$ is the external field

induced by the boundary condition outside of B, which is given by $[h_A(i)]_\ell = \sum_{k\notin A} J_{\ell,k}\sigma_k(i)$, $i \in [N]$, for every $\ell \in A$. Since $\bar{h}_A \le \bar{J}$, and $\lambda(\Lambda) = \lambda(J) \ge \lambda(J_A)$, for all $A \subset [n]$, by Theorem 1.26, for all $A \in \mathcal{A}$,

$$\mathcal{E}_{B(A)}(F, \log F) \ge \tfrac{1}{4}(1 - 2\lambda(J))\, e^{-16\bar{J}}\, \mathrm{Ent}_{\mu_{B(A)}}(F),$$

for all $F \ge 0$. From (1.87) and (1.72) we conclude

$$\mathcal{E}_{N,J,\alpha}(F, \log F) \ge \tfrac{1}{4n}(1 - 2\lambda(J))^2\, e^{-16\bar{J}}\, \mathrm{Ent}_{\mu}(F)\,.$$

This ends the proof of Theorem 1.10.

1.5 Completing Kac's program

In this section we are going to prove Theorem 1.7. The main step consists in establishing a suitable nonlinear modified log-Sobolev inequality.

1.5.1 Nonlinear MLSI

Consider the quadratic system with kernel $\mathcal{Q}_J^{\mathcal{K}}$, and let $\mathcal{A}$ be the corresponding partition of $[n]$. Let also $\mathcal{D}_\mu$ denote the entropy dissipation functional from Lemma 1.15.

Definition 1.34. For all $h \in \Gamma_{\mathcal{A}}$, let $\mathcal{S}_h$ be the set of $f : \Omega \to \mathbb{R}_+$ such that $\mu_{J,h}[f] = 1$ and such that $f\mu_{J,h} \in \mathcal{P}(\Omega)$ has the same magnetization vector of $\mu_{J,h}$, that is

$$m_{\mathcal{A}}(\mu_{J,h}) = m_{\mathcal{A}}(f\mu_{J,h}).$$

The system is said to satisfy the *Nonlinear MLSI* with rate function $\mathcal{I} : \Gamma_{\mathcal{A}} \to \mathbb{R}_+$ if for all $h \in \Gamma_{\mathcal{A}}$,

$$\mathcal{D}_{\mu_{J,h}}(f) \ge \mathcal{I}(h)\, \mu_{J,h}\,[f \log f]\ , \qquad f \in \mathcal{S}_h\,.$$

Recall the notation $\mathcal{P}_{\mathcal{A}}(\Omega)$ for the set of probability measures with nondegenerate magnetization, see Definition 1.4, and recall the uniqueness statement from Lemma 1.14.

Lemma 1.35. *For any $p \in \mathcal{P}_{\mathcal{A}}(\Omega)$, let $h = h(m_{\mathcal{A}}(p))$ denote the unique $h \in \Gamma_{\mathcal{A}}$ such that $m_{\mathcal{A}}(\mu_{J,h}) = m_{\mathcal{A}}(p)$, and let p_t denote the solution of (1.5) with initial datum p. If the system satisfies the Nonlinear MLSI with rate function $\mathcal{I}$, then*

$$H(p_t|\mu_{J,h}) \le e^{-\mathcal{I}(h)t} H(p|\mu_{J,h})\,, \qquad t \ge 0. \tag{1.88}$$

Proof. Let $\mu = \mu_{J,h}$ and write $f_t = p_t/\mu$. From Lemma 1.17 it follows that $f_t \in \mathcal{S}_h$ for all $t \ge 0$. By Lemma 1.15, and using the Nonlinear MLSI assumption, the function $\varphi(t) := H(p_t|\mu) = \mu[f_t \log f_t]$ satisfies

$$\frac{d}{dt}\,\varphi(t) = -\mathcal{D}_\mu(f_t) \le -\mathcal{I}(h)\,\mu\,[f_t \log f_t] = -\mathcal{I}(h)\varphi(t)\,.$$

Integrating, we obtain (1.88). □

To prove Theorem 1.7 it is then sufficient to prove the following *uniform* Nonlinear MLSI.

Theorem 1.36. *In the setting of Theorem 1.7, the system satisfies the Nonlinear MLSI with rate function $\mathcal{I}$ such that $\mathcal{I}(h) \geq r$, for all $h \in \Gamma_{\mathcal{A}}$, where $r \geq \frac{1}{8n}(1 - 2\lambda(J))^2 e^{-16\bar{J}}$.*

The rest of this section is devoted to the proof of Theorem 1.36. The estimate will be derived using Theorem 1.10 together with Entropic Chaos and Fisher Chaos estimates that are at the heart of Kac's program. We start with some preliminaries concerning the local central limit theorem.

1.5.2 Local Central Limit Theorem and Kac Chaos

Let $\nu \in \mathcal{P}(\Omega)$ be a probability measure on $\Omega = \{-1, +1\}^n$, and let $\mathcal{A}$ be a partition of the set $[n]$. Define the $|\mathcal{A}|$-dimensional vector $M = (M_A,\, A \in \mathcal{A})$, where

$$M_A = \sum_{\ell \in A} \sigma_\ell, \qquad A \in \mathcal{A}\,.$$

Definition 1.37. We say that $\nu \in \mathcal{P}(\Omega)$ is *irreducible* if for all $A \in \mathcal{A}$, there exists an $|\mathcal{A}|$-dimensional vector S such that $\nu(M = S) > 0$ and $\nu(M = S + 2\delta_A) > 0$, where δ_A denotes the standard unit vector in the A-th coordinate.

Let $V_\nu^{\mathcal{A}}$ denote the *covariance matrix* of the vector M,

$$V_\nu^{\mathcal{A}}(A, A') = \nu(M_A M_{A'}) - \nu(M_A)\nu(M_{A'}), \qquad A, A' \in \mathcal{A}.$$

If ν is irreducible then

$$\det V_\nu^{\mathcal{A}} > 0.$$

Indeed, suppose that $\det V_\nu^{\mathcal{A}} = 0$. Then there exists $\overline{A} \in \mathcal{A}$ such that $M_{\overline{A}}$ is almost surely a non-trivial linear combination of $\{M_A : A \neq \bar{A}\}$ under ν. Consequently, no $S \in \mathbb{R}^{\mathcal{A}}$ can exist such that both S and $S + 2\delta_A$ have positive probability under ν, yielding a contradiction. The following is a version of the local central limit theorem; see [CP24, Proposition 2.2] for a proof.

Proposition 1.38. *Let $\nu \in \mathcal{P}(\Omega)$ be irreducible, and let $M^{(N)}$ be the sum of N independent copies of M, each distributed according to ν. Then, there exists a constant $K(\nu) > 0$ such that for all $N \geq 1$,*

$$\max_S \left| \mathbb{P}\big(M^{(N)} = S\big) - \frac{e^{-\frac{1}{2}\langle z_N, z_N\rangle}}{(2\pi N)^{|\mathcal{A}|/2}\sqrt{\det V_\nu^{\mathcal{A}}}} \right| \leq \frac{K(\nu)}{(2\pi N)^{|\mathcal{A}|/2+1}},$$

where z_N is the $|\mathcal{A}|$-dimensional vector defined as

$$z_N = \frac{(V_\nu^{\mathcal{A}})^{-1/2}}{\sqrt{N}} \left(S - \mathbb{E}[M^{(N)}] \right),$$

and the maximum is taken over all $|\mathcal{A}|$-tuples $S = (S_A, A \in \mathcal{A})$ such that for each $A \in \mathcal{A}$, $S_A = 2X_A - N|A|$ for some integer $X_A \in \{0, \dots, N|A|\}$.

We apply the above statement with n and $\mathcal{A}$ fixed, while we take N arbitrarily large. In particular, we use Proposition 1.38 to obtain a lower bound on the probability of the event $\Omega_{N,\alpha}$ defined in (1.13), for suitable sequences $\alpha \in \mathcal{D}_{\mathcal{A},N}$, when the spin configuration $(\sigma(i),\, i \in [N])$ is distributed according to the product measure $\nu^{\otimes N}$.

Corollary 1.39. *Let $\nu \in \mathcal{P}(\Omega)$ be irreducible and assume that $\alpha \in \mathcal{D}_{\mathcal{A},N}$ satisfies the condition*

$$\langle \alpha - \alpha_{\mathcal{A}}(\nu), \alpha - \alpha_{\mathcal{A}}(\nu) \rangle = O\left(1/N\right), \tag{1.89}$$

where $\alpha_{\mathcal{A}}(\nu) = (1 + m_{\mathcal{A}}(\nu))/2$ is the density profile associated to ν. Then,

$$\nu^{\otimes N}\left(\Omega_{N,\alpha}\right) \geq \frac{c}{N^{|\mathcal{A}|/2}},$$

where $c = c(\nu) > 0$ is a constant independent of N. In particular,

$$\lim_{N\to\infty} \frac{1}{N} \log\left(\nu^{\otimes N}\left(\Omega_{N,\alpha}\right) \right) = 0.$$

Proof. The event $\Omega_{N,\alpha}$ can be written as $M^{(N)} = S$ with $S_A = (2\alpha(A) - 1)N|A|$, $A \in \mathcal{A}$. Moreover, for each A, the A-component of $\mathbb{E}[M^{(N)}]$ is given by $N|A|m_{\mathcal{A}}(\nu)$. The results follow from Proposition 1.38 by noting that (1.89) is equivalent to $\langle z_N, z_N \rangle = O(1)$. □

Given $\nu \in \mathcal{P}(\Omega)$, it is always possible to find $\alpha \in \mathcal{D}_{\mathcal{A},N}$ such that (1.89) is satisfied. In fact, choosing

$$\alpha(A) = \frac{1}{N|A|} \lfloor N|A|(1 + m(A,\nu))/2 \rfloor, \qquad A \in \mathcal{A} \tag{1.90}$$

one has $\alpha(A) - \frac{1}{2}(1 + m(\nu, A)) = O(1/N)$, for all $A \in \mathcal{A}$, and therefore (1.89) actually holds with $O(1/N^2)$ in place of $O(1/N)$.

Definition 1.40 (Canonical tensor product)**.** The choice (1.90) of $\alpha \in \mathcal{D}_{\mathcal{A},N}$ will be called the *canonical density* associated to ν. We also define the *canonical tensor product* associated to $\nu \in \mathcal{P}(\Omega)$ as the product measure $\nu^{\otimes N}$ conditioned on $\Omega_{N,\alpha}$, where α is the canonical density for ν,

$$\gamma_N(\nu) = \nu^{\otimes N}\left(\cdot \mid \Omega_{N,\alpha} \right).$$

Kac's Chaos refers to the fact that marginals of the canonical tensor product converge as $N \to \infty$ to independent copies of ν. The next proposition quantifies this statement. For all fixed $k \in \mathbb{N}$, we write $P_k\gamma$ for the marginal of a measure $\gamma \in \mathcal{P}(\Omega^N)$ on the first k coordinates Ω^k, for $k \leq N$. We refer to [CP24, Theorem 2.6] for a proof.

Proposition 1.41 (Kac Chaos). *Let $\nu \in \mathcal{P}(\Omega)$ be irreducible. Then, for all $N \geq k \geq 1$,*

$$\|P_k\gamma_N(\nu) - \nu^{\otimes k}\|_{\mathrm{TV}} \leq \frac{Ck}{N},$$

for some constant $C = C(\nu)$.

1.5.3 Entropic and Fisher Chaos

We turn to the implications of the above estimates for the relative entropy and the Dirichlet form of the exchange dynamics.

Proposition 1.42 (Entropic Chaos). *Suppose $\nu_1, \nu_2 \in \mathcal{P}(\Omega)$ are irreducible measures such that $\nu_1 \ll \nu_2$ and $m_{\mathcal{A}}(\nu_1) = m_{\mathcal{A}}(\nu_2)$. Then,*

$$\lim_{N\to+\infty} \frac{1}{N} H_N\left(\gamma_N(\nu_1) \mid \gamma_N(\nu_2)\right) = H(\nu_1 \mid \nu_2).$$

Proof. By definition of the canonical tensor products, calling α the common canonical density for ν_1, ν_2,

$$H_N\left(\gamma_N(\nu_1) \mid \gamma_N(\nu_2)\right) = \gamma_N(\nu_1)\left[\log\left(\frac{\nu_1^{\otimes N}}{\nu_2^{\otimes N}}\right)\right] - \log\left(\frac{\nu_1^{\otimes N}(\Omega_{N,\alpha})}{\nu_2^{\otimes N}(\Omega_{N,\alpha})}\right).$$

Corollary 1.39 implies

$$\lim_{N\to+\infty} \frac{1}{N} \log\left(\frac{\nu_1^{\otimes N}(\Omega_{N,\alpha})}{\nu_2^{\otimes N}(\Omega_{N,\alpha})}\right) = 0.$$

Moreover, by symmetry

$$\frac{1}{N}\gamma_N(\nu_1)\left[\log\left(\frac{\nu_1^{\otimes N}}{\nu_2^{\otimes N}}\right)\right] = P_1\gamma_N(\nu_1)\left[\log\left(\frac{\nu_1}{\nu_2}\right)\right],$$

and the conclusion follows by Proposition 1.41. □

Fisher Chaos is the convergence of the Dirichlet form (1.15) when the function F is the density of a canonical tensor product. Recalling the notation from Definition 1.34, this can be formulated as follows.

Proposition 1.43 (Fisher Chaos). *Fix $h \in \Gamma_{\mathcal{A}}$, and let $\mu := \mu_{J,h}$. Suppose $f \in \mathcal{S}_h$ is such that the measure $\nu := f\mu \in \mathcal{P}(\Omega)$ is irreducible. Let α denote the common canonical density of ν and μ, and define*

$$F_N := \frac{\gamma_N(\nu)}{\gamma_N(\mu)}. \tag{1.91}$$

Then,

$$\lim_{N\to\infty} \frac{1}{N}\,\mathcal{E}_{N,J,\alpha}(F_N, \log F_N) = 2\mathcal{D}_\mu(f). \tag{1.92}$$

Proof. For any $F : \Omega^N \to \mathbb{R}_+$ we write $\mathcal{E}_{N,J,\alpha}(F, \log F) = \frac{1}{n}\sum_{\ell,k\in[n]} \mathcal{K}(\ell,k)\mathcal{E}^{\ell,k}_{N,J,\alpha}(F)$, where

$$\mathcal{E}^{\ell,k}_{N,J,\alpha}(F) := -\frac{1}{N}\sum_{i,j=1}^{N} \mu_{N,J,\alpha}\left[\hat{p}_J(\ell,k|\sigma(i),\sigma(j))\nabla^{i,\ell;j,k}F\log F\right].$$

Similarly, $\mathcal{D}_\mu(f) = \frac{1}{n}\sum_{\ell,k\in[n]} \mathcal{K}(\ell,k)\mathcal{D}^{\ell,k}_\mu(f)$, where

$$\begin{aligned}\mathcal{D}^{\ell,k}_\mu(f) := -\frac{1}{2}\sum_{\tau,\tau'\in\Omega} &\mu(\tau)\mu(\tau')p_J(\ell,k|\tau,\tau')\times\\ &\times\ \big(f(S_\ell(\tau,\tau'_k))f(S_k(\tau',\tau_\ell)) - f(\tau)f(\tau')\big)\log\big(f(\tau)f(\tau')\big).\end{aligned}$$

Therefore, to prove (1.92) it is sufficient to show that for all fixed n and $\ell,k\in[n]$, with F_N as in (1.91),

$$\lim_{N\to\infty}\frac{1}{N}\,\mathcal{E}^{\ell,k}_{N,J,\alpha}(F_N) = 2\mathcal{D}^{\ell,k}_\mu(f).$$

Since F_N and $\mu_{N,J,\alpha}$ are invariant under particle permutations,

$$\frac{1}{N}\mathcal{E}^{\ell,k}_{N,J,\alpha}(F_N) = \frac{N-1}{N}\mathcal{E}_{1,2}(F_N) + \frac{1}{N}\mathcal{E}_1(F_N)\,,$$

where

$$\begin{aligned}\mathcal{E}_{1,2}(F_N) &= -\mu_{N,J,\alpha}\left[p_J(\ell,k|\sigma(1),\sigma(2))\nabla^{1,\ell;2,k}F_N\log F_N\right],\\ \mathcal{E}_1(F_N) &= -\mu_{N,J,\alpha}\left[\hat{p}_J(\ell,k|\sigma(1),\sigma(1))\nabla^{1,\ell;1,k}F_N\log F_N\right].\end{aligned}$$

Thus, it suffices to show that

$$\lim_{N\to\infty}\mathcal{E}_{1,2}(F_N) = 2\mathcal{D}^{\ell,k}_\mu(f)\,, \tag{1.94}$$

$$\lim_{N\to\infty}\frac{1}{N}\,\mathcal{E}_1(F_N) = 0\,. \tag{1.95}$$

For any configuration $\sigma\in\Omega_{N,\alpha}$,

$$F_N(\sigma) = \frac{\mu^{\otimes N}(\Omega_{N,\alpha})}{\nu^{\otimes N}(\Omega_{N,\alpha})}\prod_{s=1}^{N} f(\sigma(s)),$$

so that

$$\begin{aligned}\mathcal{E}_{1,2}(F_N) = -\sum_{\sigma\in\Omega_{N,\alpha}} &\frac{\mu(\sigma(1))\mu(\sigma(2))\prod_{s\neq1,2}\nu(\sigma(s))}{\nu^{\otimes N}(\Omega^{\mathcal{A}}_{\alpha_N})}\,p_J(\ell,k\mid\sigma(1),\sigma(2))\,\log\left(F_N(\sigma)\right)\,\times\\ &\times\left[f(S_\ell(\sigma(1),\sigma(2)_k))f(S_k(\sigma(2),\sigma(1)_\ell)) - f(\sigma(1))f(\sigma(2))\right].\end{aligned}$$

We decompose $\Omega_{N,\alpha}$ by fixing the values $\tau = \sigma(1), \tau' = \sigma(2)$,

$$\Omega_{N,\alpha} = \cup_{\tau,\tau' \in \Omega} \Omega_{N,\alpha}(\tau, \tau'), \tag{1.96}$$

where $\Omega_{N,\alpha}(\tau, \tau')$ denotes the set of configurations $(\sigma(s))_{s \neq 1,2} \in \Omega^{N-2}$ such that

$$\sum_{\ell \in A} \sum_{s \neq 1,2} \sigma_\ell(s) = (N-2)|A| \left(2\widetilde{\alpha}(A \mid \tau, \tau') - 1\right), \qquad A \in \mathcal{A},$$

where the adjusted density vector $\widetilde{\alpha}(\cdot \mid \tau, \tau') \in \mathcal{D}_{\mathcal{A},N-2}$ satisfies

$$(2\widetilde{\alpha}(A \mid \tau, \tau') - 1) = \frac{N}{N-2}(2\alpha(A) - 1) - \frac{1}{(N-2)|A|} \sum_{\ell \in A} (\tau_\ell + \tau'_\ell).$$

Note that $\widetilde{\alpha}(A \mid \tau, \tau') = \alpha(A) + O(1/N)$ for all A and $\tau, \tau' \in \Omega$. In particular, by Proposition 1.38, for all $\tau, \tau' \in \Omega$,

$$\lim_{N \to \infty} \frac{\nu^{\otimes N-2}(\Omega_{N,\alpha}(\tau, \tau'))}{\nu^{\otimes N}(\Omega_{N,\alpha})} = 1. \tag{1.97}$$

From the decomposition (1.96),

$$\mathcal{E}_{1,2}(F_N) = - \sum_{\tau,\tau' \in \Omega} \mu(\tau)\mu(\tau') p_J(\ell, k \mid \tau, \tau') \left[f(S_\ell(\tau, \tau'_k)) f(S_k(\tau', \tau_\ell)) - f(\tau) f(\tau') \right] \times$$
$$\times \sum_{\eta \in \Omega_{N,\alpha}(\tau,\tau')} \frac{\prod_{s \neq 1,2} \nu(\eta(s))}{\nu^{\otimes N}(\Omega_{N,\alpha})} \left[\log \frac{\mu^{\otimes N}(\Omega_{N,\alpha})}{\nu^{\otimes N}(\Omega_{N,\alpha})} + \log(f(\tau) f(\tau')) + \log \left(\prod_{s \neq 1,2} f(\eta(s)) \right) \right].$$

Using reversibility as in Lemma 1.13, we see that the function

$$\phi(\tau, \tau') := \mu(\tau)\mu(\tau') p_J(\ell, k \mid \tau, \tau') \left[f(S_\ell(\tau, \tau'_k)) f(S_k(\tau', \tau_\ell)) - f(\tau) f(\tau') \right]$$

is antisymmetric in the exchange $(\tau, \tau') \leftrightarrow (S_\ell(\tau, \tau'_k), S_k(\tau', \tau_\ell))$. On the other hand,

$$\Omega_{N,\alpha}(\tau, \tau') = \Omega_{N,\alpha}(S_\ell(\tau, \tau'_k), S_k(\tau', \tau_\ell)).$$

It follows that

$$\mathcal{E}_{1,2}(F_N) = - \sum_{\tau,\tau' \in \Omega} \frac{\nu^{\otimes N-2}(\Omega_{N,\alpha}(\tau, \tau'))}{\nu^{\otimes N}(\Omega_{N,\alpha})} \phi(\tau, \tau') \log(f(\tau) f(\tau')).$$

From (1.97) we obtain

$$\lim_{N \to \infty} \mathcal{E}_{1,2}(F_N) = - \sum_{\tau,\tau' \in \Omega} \phi(\tau, \tau') \log(f(\tau) f(\tau')) = 2\mathcal{D}_\mu(f).$$

This proves (1.94). To prove (1.95) we proceed in a similar way. We decompose $\Omega_{N,\alpha} = \cup_{\tau \in \Omega} \Omega_{N,\alpha}(\tau)$ by fixing $\tau = \sigma(1)$, where $\Omega_{N,\alpha}(\tau) = \Omega_{N-1,\widetilde{\alpha}(\cdot|\tau)}$ denotes the set of configurations $(\sigma(s))_{s \neq 1} \in \Omega^{N-1}$ with fixed density $\widetilde{\alpha}(\cdot \mid \tau)$ such that $\widetilde{\alpha}(A \mid \tau) = \alpha(A) + O(1/N)$

for all A and $\tau \in \Omega$. Reasoning as above, it follows that

$$\mathcal{E}_1(F_N) = -\sum_{\tau\in\Omega} \frac{\nu^{\otimes N-1}(\Omega_{N,\alpha}(\tau))}{\nu^{\otimes N}(\Omega_{N,\alpha})}\, \phi(\tau)\log(f(\tau)),$$

where

$$\phi(\tau) = \mu(\tau)\hat{p}_J(\ell, k \mid \tau, \tau)\left[f(\tau^{\ell,k}) - f(\tau)\right].$$

Therefore,

$$\lim_{N\to\infty} \mathcal{E}_1(F_N) = -\sum_{\tau\in\Omega} \phi(\tau)\log(f(\tau))\,.$$

This implies (1.95). □

1.5.4 Proof of Theorem 1.7

As we have seen in Section 1.5.1, it is sufficient to prove Theorem 1.36. Take $f \in \mathcal{S}_h$, for a fixed $h \in \Gamma_{\mathcal{A}}$, and write F_N for the function in Proposition 1.43. Set $\mu = \mu_{J,h}$, and assume that $\nu = f\mu$ is irreducible in the sense of Definition 1.37. Therefore, we know that

$$2\,\mathcal{D}_\mu(f) = \lim_{N\to\infty} \frac{1}{N}\,\mathcal{E}_{N,J,\alpha}(F_N, \log F_N)\,.$$

On the other hand, by Theorem 1.10, $\mathcal{E}_{N,J,\alpha}(F_N, \log F_N) \geq r_0\ \mathrm{Ent}_{N,J,\alpha}\, F_N$, for all N, with $r_0 \geq \frac{1}{4n}(1-2\lambda(J))^2 e^{-16\bar{J}}$. By Proposition 1.42 with $\nu_1 = \nu$ and $\nu_2 = \mu$, one has the convergence

$$\lim_{N\to\infty} \frac{1}{N}\ \mathrm{Ent}_{N,J,\alpha}\, F_N = \mu[f\log f]\,.$$

It follows that

$$\mathcal{D}_\mu(f) \geq r\,\mu[f\log f]\,,$$

with $r = r_0/2$. This establishes the desired bound for all $f \in \mathcal{S}_h$ such that $\nu = f\mu$ is irreducible. If ν is not irreducible, we proceed by regularising it as follows. For any $\varepsilon \in (0,1)$, define

$$f_\varepsilon = (1-\varepsilon)f + \varepsilon.$$

Note that $f_\varepsilon \in \mathcal{S}_h$ whenever $f \in \mathcal{S}_h$. Moreover $\nu_\varepsilon := f_\varepsilon\mu$ is irreducible since $\nu_\varepsilon(\tau) \geq \varepsilon\mu(\tau) > 0$ for all $\tau \in \Omega$. Thus, applying the above reasoning to f_ε we obtain

$$\mathcal{D}_\mu(f_\varepsilon) \geq r\,\mu[f_\varepsilon \log f_\varepsilon]\,.$$

Taking the limit as $\varepsilon \to 0$, by continuity we obtain the desired inequality for f. This proves $\mathcal{I}(h) \geq \alpha$ for all $h \in \Gamma_{\mathcal{A}}$, which ends the proof.

1.5.5 Open problem: exponential ergodicity at all temperatures?

We conclude the chapter with a discussion of some open questions. The exponential decay established in Theorem 1.7 holds under the assumption that the interaction matrix J is nonnegative definite and has spectral radius $\lambda(J) < 1/2$. The positivity assumption is, in principle, not restrictive, since one may freely modify the diagonal entries of J without affecting the corresponding Ising measure. The spectral radius condition, however, imposes that the interaction be sufficiently weak, which we interpret as a high-temperature assumption.

For instance, in the case of the Curie-Weiss model, where $J_{i,j} = \beta/n$ for all $i, j \in [n]$, the condition $\lambda(J) < 1/2$ imposes $\beta < 1/2$ which is a factor 2 off with respect to the critical value $\beta_c = 1$. It would be certainly desirable to have exponential decay with rate of order $1/n$ throughout the whole region $\beta \in (0, \beta_c)$ in this case. On the other hand, understanding the behavior of the system in the low temperature region $\beta > \beta_c$ seems to be a very challenging open problem. In particular, it is not even clear whether one should expect a slowdown or not for the nonlinear equation.

Nevertheless, by analogy with linear evolutions such as low-temperature Glauber or Kawasaki dynamics [Mar97], one may ask whether exponential decay to equilibrium still holds for every interaction J, possibly with rates that deteriorate exponentially in n. Establishing such a result would amount to an exponential ergodicity property for the nonlinear equation, providing a quantitative counterpart to the plain convergence in Theorem 1.5.

A possible strategy toward this goal is to revisit Kac's program and show that a statement analogous to Theorem 1.10 remains valid for arbitrary interactions J, with some rate $\alpha \geq c(J)$ for a constant $c(J) > 0$ that may depend on J, possibly exponentially small in n. By repeating the reasoning in Section 1.4, and in particular, invoking Proposition 1.30, such a result would follow if one could prove that the tilted canonical Ising measure $\mathcal{T}_v \mu_{N,\alpha}$ satisfies a covariance bound of the form

$$\lambda(\mathrm{Cov}[\mathcal{T}_v \mu_{N,\alpha}]) \leq C(J),$$

for some constant $C(J)$ that may grow with n but remains uniformly bounded in both the linear tilt v and the number of particles N. We conjecture that such a bound indeed holds for all interactions J. Using equivalence-of-ensembles estimates, as in Proposition 1.41, one can obtain such bounds for each fixed choice of the tilt v, uniformly in N. However, establishing estimates that are also uniform in the tilt parameter appears to be a rather challenging problem.

Chapter 2

Nonlinear Mean-Field Hardcore Model

This chapter is entirely devoted to the analysis of nonlinear exchange dynamics in mean field case for the hardcore model. The content presented here is based on the article [Min+26]. For the sake of completeness, we briefly recall the mathematical definition of the model. Let $G = (V, E)$ be an undirected graph with $V = [n]$ and maximum degree $\Delta \geq 3$, and let $\Omega_G \subseteq \{0,1\}^n$ denote the set of independent sets of G, which are the configurations $\sigma = (\sigma_i,\, i \in [n])$ with the constraint $\sigma_i \sigma_j = 0$ for all $\{i,j\} \in E$. The hardcore model on G at fugacity $\lambda > 0$ is defined through the Gibbs measure

$$\mu_{G,\lambda}(\sigma) = \frac{\lambda^{|\sigma|}}{Z_G(\lambda)}\, \mathbf{1}(\sigma \in \Omega_G), \qquad |\sigma| = \sum_{i=1}^{n} \sigma_i, \tag{2.1}$$

where $|\sigma|$ denotes the size of the independent set σ and $Z_G(\lambda)$ is the normalizing partition function. We refer to the Introduction of this thesis for a discussion of the hardcore model, the uniqueness threshold $\lambda_c(\Delta)$, the critical density $\alpha_c(\Delta)$, and for the relevant background and motivation.

The results discussed in this chapter were already presented informally in the introduction of this thesis. Here, we expose them in full detail, providing the complete statements of the theorems together with their proofs.

2.1 Introduction and Main results

2.1.1 The collision kernel

A collision between two spin configurations $(\sigma, \sigma') \in \Omega \times \Omega$ consists in the exchange of the ℓ-th spin of σ with the k-th spin of σ', that is

$$(\sigma, \sigma') \mapsto \left(S_\ell(\sigma, \sigma'_k),\ S_k(\sigma', \sigma_\ell)\right), \tag{2.2}$$

where $S_\ell : \Omega \times \{0,1\} \to \Omega$ is the configuration defined by

$$S_\ell(\sigma, a)_i = \begin{cases} \sigma_i & \text{if } i \neq \ell, \\ a & \text{if } i = \ell. \end{cases}$$

This mechanism is well defined for every pair ℓ, k for which configuration $S_\ell(\sigma, \sigma'_k)$ and $S_k(\sigma', \sigma_\ell)$ are still independent sets of G. In this case we say the exchange is *admissible*, and we define $(\ell, k)_{\sigma,\sigma'}$ as the event that the exchange (ℓ, k) is admissible for the pair of independent sets $\sigma, \sigma' \in \Omega_G$, and by $\widehat{(\ell, k)}_{\sigma,\sigma'}$ its complementary event.

We define the *mean-field* collision kernel $\mathcal{Q}$ by

$$\mathcal{Q}(\sigma, \sigma'; \tau, \tau') = \frac{1}{n^2} \sum_{\ell,k \in [n]} Q_{\ell,k}(\sigma, \sigma'; \tau, \tau'), \tag{2.3}$$

with

$$Q_{\ell,k}(\sigma, \sigma'; \tau, \tau') = \mathbf{1}\{(\ell, k)_{\sigma,\sigma'}\}\, \mathbf{1}\{(\tau, \tau') = (\sigma, \sigma')^{\ell,k}\} + \mathbf{1}\{\widehat{(\ell, k)}_{\sigma,\sigma'}\}\, \mathbf{1}\{(\tau, \tau') = (\sigma, \sigma')\}.$$

Here $(\sigma, \sigma')^{\ell,k}$ is a shorthand for the right hand side of (2.2). In words, given a pair (σ, σ'), vertices (ℓ, k) are selected uniformly at random and an exchange between sites ℓ and k is performed if and only if it results in a new pair of independent sets.

More generally, one can define a collision kernel associated to any symmetric stochastic matrix $\mathcal{K}$:

$$\mathcal{Q}^{\mathcal{K}}(\sigma, \sigma'; \tau, \tau') := \frac{1}{n} \sum_{\ell,k=1}^{n} \mathcal{K}(\ell, k)\, Q_{\ell,k}(\sigma, \sigma'; \tau, \tau').$$

The definition in (2.3) corresponds to the choice $\mathcal{K}(\ell, k) = 1/n$ for all $\ell, k \in [n]$, which we refer to as the mean-field case. Another case of interest is the *single-site* kernel $\mathcal{K}(\ell, k) = \mathbf{1}(\ell = k)$, which will be studied in [Min+26]. In this work, however, we analyse only the nonlinear dynamics for the hardcore model in the mean-field case.

2.1.2 Nonlinear dynamics and convergence theorem

Let $\mathcal{P}(\Omega_G)$ denote the set of probability measures on Ω_G. Given the collision kernel $\mathcal{Q}$ defined in (2.3), the nonlinear evolution on $\mathcal{P}(\Omega_G)$ is defined by the dynamical system

$$\frac{dp_t}{dt} = p_t \circ p_t - p_t, \qquad p_0 = p, \tag{2.4}$$

for some initial value $p \in \mathcal{P}(\Omega_G)$, where $(p \circ q)(\tau)$, for $p, q \in \mathcal{P}(\Omega_G)$, denotes the commutative collision product

$$(p \circ q)(\tau) = \frac{1}{2} \sum_{\sigma,\sigma',\tau' \in \Omega_G} \big(p(\sigma)q(\sigma') + p(\sigma')q(\sigma)\big)\, \mathcal{Q}(\sigma, \sigma'; \tau, \tau'). \tag{2.5}$$

Definition 2.1 (Density)**.** Given $p \in \mathcal{P}(\Omega_G)$, the *density* $\alpha(p)$ is the expected density of

an independent set sampled according to p:

$$\alpha(p) = \frac{1}{n}\mathbb{E}_p\left[|\sigma|\right].$$

For example, the density of the hardcore model with fugacity $\lambda > 0$ coincides with its occupancy ratio:

$$\alpha(\mu_{G,\lambda}) = \alpha_G(\lambda) := \frac{1}{n}\frac{Z'_G(\lambda)}{Z_G(\lambda)}.$$

We denote by $\mathcal{P}_+(\Omega_G)$ the family of distributions fully supported on the set of independent sets Ω_G; that is, $p \in \mathcal{P}_+(\Omega_G)$ if and only if $p \in \mathcal{P}(\Omega_G)$ and $p(\sigma) > 0$ for all $\sigma \in \Omega_G$. In the next result we formulate our main convergence statement for the mean-field nonlinear process.

Theorem 2.2. *A distribution $p \in \mathcal{P}_+(\Omega_G)$ is a stationary state if and only if $p = \mu_{G,\lambda}$ for some $\lambda > 0$. Moreover, the density $\alpha(p_t)$ is independent of $t \geq 0$. Finally, for any initial state $p \in \mathcal{P}(\Omega_G)$ with $\alpha(p) \in \left(0, \frac{1}{1+\Delta}\right)$,*

$$\|p_t - \mu_{G,\lambda}\|_{TV} \to 0, \quad t \to +\infty,$$

where $\lambda = \lambda(p)$ is the unique fugacity such that $\alpha(\mu_{G,\lambda(p)}) = \alpha(p)$.

As mentioned in the main introduction, the range of convergence stated in Theorem 2.2 is optimal: it is not possible in general to extend this result beyond the critical value $1/(\Delta + 1)$; see the example proposed after Theorem 0.2.

2.1.3 Rates of convergence

Theorem 2.2 does not give any information on rates of convergence. Here we discuss our quantitative results. We have two sets of results: exponential decay in Wasserstein distance, and relative entropy decay. As discussed in Remarks 1.8 and 1.12, both exponential decay of Wasserstein distance and relative entropy imply the same exponential decay rate (up to a constant factor) for the total variation distance, uniformly in the initial condition.

Wasserstein contraction

Recall the Wasserstein distance W_1 defined in (0.8).

Theorem 2.3. *For all initial states $p \in \mathcal{P}(\Omega_G)$ with $\alpha(p) \in (0, 1/(3(\Delta + 1)))$, the solution p_t of (2.4) satisfies*

$$W_1(p_t, \mu_{G,\lambda}) \leq W_1(p, \mu_{G,\lambda})\, e^{-\varepsilon t/n},$$

where $\varepsilon = 1 - 3(\Delta+1)\alpha(p) > 0$ and $\lambda = \lambda(p)$ is the unique fugacity such that $\alpha(\mu_{G,\lambda}) = \alpha(p)$.

Entropy decay

Recall the definition of relative entropy given in (0.6). We also recall the definition of the critical density

$$\alpha_c(\Delta) = \frac{\lambda_c(\Delta)}{1 + (\Delta + 1)\lambda_c(\Delta)}, \tag{2.6}$$

which serves as the critical threshold governing the range of entropy decay in the next result. It is important to relate this quantity to the upper bounds appearing in other convergence theorems, namely $\frac{1}{\Delta+1}$ in Theorem 2.2 and $\frac{1}{3(\Delta+1)}$ in Theorem 2.3. In particular, we have

$$\frac{1}{3(\Delta+1)} < \alpha_c(\Delta) < \frac{1}{\Delta+1}.$$

Theorem 2.4 (Informal version of Theorem 2.36)**.** *For any $\delta > 0$, if $\alpha(p) \in (\delta, (1-\delta)\alpha_c(\Delta))$, then the convergence stated in Theorem 2.2 can be strengthened to exponential decay in relative entropy:*

$$H(p_t \mid \mu_{G,\lambda}) \leq e^{-rt/n} H(p \mid \mu_{G,\lambda}),$$

with rate $r = r(\gamma, \Delta) > 0$.

We conjecture that the rate r should not depend on the lower bound on the density. However, our result in Theorem 2.4 is not uniform in γ. On the other hand, note that the bound in Theorem 2.3 does not deteriorate as the density becomes samall. Moreover, it is known that for Glauber dynamics, entropy decay holds uniformly in the fugacity, provided it is sufficiently close to zero; see [CE22]. It is however interesting to note that Theorem 2.4 allows one to reach densities that are arbitrarily close to the critical value $\alpha_c(\Delta)$ whereas Theorem 2.3 requires a cruder upper bound on the density.

Theorem 2.4 will be proved using Kac's program, namely by establishing a uniform modified log-Sobolev inequality for the associated linear interacting particle system, which corresponds to an exchange dynamics on fixed-size independent sets on a disjoint union of N copies of G. For this step we rely on the analysis in [Jai+23]. Completing Kac's program requires one further step, involving entropic chaos and Fisher chaos, based on a local central limit theorem. For that part we can proceed as in the context of the nonlinear process for the Ising model; see the previous chapter, and [CP24]. It is straightforward to check that the same proofs of Section 1.5 can be adapted to the present setting of the Hardcore model. For the reader's convenience, we reproduce the main part of the argument in Section 2.5.

2.1.4 Kac's program

We consider a system of $N \in \mathbb{N}$ *particles* $\{\sigma(i), i = 1, \dots, N\}$, where each $\sigma(i) = (\sigma_1(i), \dots, \sigma_n(i)) \in \Omega_G$. Alternatively, $\sigma \in \Omega^N$ can be seen as an independent set over the graph $\mathcal{G}_N$ given by the disjoint union of N copies $G(i)$ of G. The process is described as follows. Each ordered pair (i, j), is updated at the arrival times of an independent

Poisson process with rate $1/N$. When an update occurs at (i,j), a pair of sites ℓ, k is picked uniformly at random and, if $i \neq j$ and $(S_\ell(\sigma(i),\sigma(j)_k), S_k(\sigma(j),\sigma(i)_\ell)) \in \Omega_G \times \Omega_G$, the ℓ-th spin of particle i is exchanged with the k-th spin of particle j. Note that the updates conserve the independent structure of each $\sigma(i)$, with $i \in [N]$ and the process is well-defined and they involve every time a pair of different particles $i \neq j$. In particular, the exchange mechanism which defines the particle system is the same of nonlinear dynamics mean field, once $i, j \in [N]$ are fixed, then $\ell, k \in [n]$ are sampled uniformly at random and the current configuration $\sigma = (\sigma(s),\ s \in [N])$ is updated in $s = i, j$ with the exchange rule

$$(\sigma(i),\sigma(j)) \mapsto \left(S_\ell(\sigma(i),\sigma(j)_k)\,,\ S_k(\sigma(j),\sigma(i)_\ell)\right), \tag{2.7}$$

provided that such exchange is admissible, which we denote through the event $(\ell,k)_{\sigma(i),\sigma(j)}$ that is $S_\ell(\sigma(i),\sigma(j)_k)$ and $S_k(\sigma(j),\sigma(i)_\ell)$ are independent set. The process described above has a precise expression in term of continuous-time Markov chain with generator

$$\mathcal{L}_N F(\sigma) = \frac{1}{Nn^2} \sum_{i,j=1}^{N} \sum_{\ell,k\in[n]} \mathbf{1}\{(\ell,k)_{\sigma(i),\sigma(j)}\}\, \nabla^{i,\ell;j,k} F(\sigma),$$

where the gradient operator $\nabla^{i,\ell;j,k}$ is defined by

$$\nabla^{i,\ell;j,k} F(\sigma) := F(\sigma^{i,\ell;j,k}) - F(\sigma), \quad \sigma^{i,\ell;j,k}(u) := \begin{cases} \sigma(u) & u \neq i,j, \\ S_\ell(\sigma(i),\sigma(j)_k) & u = i, \\ S_k(\sigma(j),\sigma(i)_\ell) & u = j, \end{cases}$$

with the convention that when $i = j$, $\sigma^{i,\ell;j,k} = \sigma$. Our main purpose in considering this particle system is to study its convergence in terms of the functional inequality MLSI, and to relate it to the entropy decay of the nonlinear process in the spirit of Kac's program. An interesting aspect of Kac's program that we will not develop here is the *propagation of chaos*, which compares the evolution of the particle system with that of the nonlinear process in the limit $N \to \infty$. We will partially encounter this aspect in the application presented in the next subsection, where the connection between these two systems yields a poly(n) algorithm for the hardcore model.

We now present the main ingredients to study the convergence of the linear system. First, we observe that the mechanism in (2.7) conserves the total number of 1's, or equivalently the density of the system. In order to guarantee convergence, we therefore fix the density to a prescribed value.

Definition 2.5. We write D_N for the set of densities $\alpha \in \left(0, \frac{1}{\Delta+1}\right]$ such that $\alpha N n$ is a nonnegative integer. Given $\alpha \in D_N$, define

$$\Omega_{N,\alpha} := \left\{ \sigma \in \Omega_G^N : \sum_{i=1}^{N} \sum_{\ell=1}^{n} \sigma_\ell(i) = Nn\alpha \right\}.$$

Note that the upper bound $\frac{1}{\Delta+1}$ in the definition of D_N ensures that $\Omega_{N,\alpha}$ is non-empty for every $\alpha \in D_N$. Once $\alpha \in D_N$ is fixed and the system has the configuration $\sigma \in \Omega_{N,\alpha}$ the next state is chosen uniformly at random between possible states available with one step of the mechanism (2.7). Thus, this makes the uniform measure on the set $\Omega_{N,\alpha}$ a reversible measure for the process defined by $\mathcal{L}_N$. Namely, for every $\alpha \in D_N$ the distribution

$$\mu_{N,\alpha}(\sigma) = \frac{\mathbf{1}\left(\sigma \in \Omega_{N,\alpha}\right)}{|\Omega_{N,\alpha}|}, \quad \sigma \in \Omega_G^N$$

is a invariant measure for $\mathcal{L}_N$ and the pair $(\mathcal{L}_N, \mu_{N,\alpha})$ is reversible. To present the convergence statement in terms of the Modified Log Sobolev inequality we need to introduce the relevant quantities: given $\alpha \in D_N$ the Dirichlet Form associated to the pair $(\mathcal{L}_N, \mu_{N,\alpha})$ is

$$\mathcal{E}_{N,\alpha}(F,G) := \frac{1}{2Nn^2} \sum_{i,j=1}^{N} \sum_{\ell,k=1}^{n} \mu_{N,\alpha}\Big[\mathbf{1}\{(\ell,k)_{\sigma(i),\sigma(j)}\} \nabla^{i,\ell;j,k} F \, \nabla^{i,\ell;j,k} G\Big],$$

for all $F, G : \Omega_{N,\alpha} \to \mathbb{R}$.

Theorem 2.6 (MLSI for the particle system)**.** *Suppose G has maximum degree $\Delta \geq 3$ and n vertices, and let $N \geq n^2$. For any $\gamma > 0$, if $\alpha \in D_N$ satisfies $\alpha \in (\gamma, (1-\delta)\alpha_c(\Delta))$, then*

$$\mathcal{E}_{N,\alpha}(F, \log F) \geq \frac{r}{n} \operatorname{Ent}_{\mu_{N,\alpha}}(F) \qquad \forall\, F \geq 0, \tag{2.8}$$

with $r = \Omega_{\gamma,\delta,\Delta}(1)$.

Here, we used the notation $\operatorname{Ent}_{\mu_{N,\alpha}}(F) = \mathbb{E}_{\mu_{N,\alpha}}\left[F \log\left(\frac{F}{\mathbb{E}_{\mu_{N,\alpha}}[F]}\right)\right]$, and the subscripts γ, δ, Δ in $\Omega_{\gamma,\delta,\Delta}(\cdot)$ indicate that the implicit constant depends only on γ, δ, Δ. As in the Ising case discussed above for the particle system, the MLSI in (2.8) implies exponential decay of the heat semigroup $e^{t\mathcal{L}_N}$ in relative entropy: for any initial distribution $\nu_0 \in \mathcal{P}(\Omega_{N,\alpha})$, letting $\nu_t = e^{t\mathcal{L}_N}\nu_0$ denote the law of the process at time $t \in \mathbb{R}_+$, one has

$$H(\nu_t \mid \mu_{N,\alpha}) \leq e^{-rt/n} H(\nu_0 \mid \mu_{N,\alpha}).$$

Theorem 2.6 is proved by adapting the argument of Theorem 2 in [Jai+23] to the Kac particle system. We note that the result in [Jai+23] is in fact stronger than we need, as it necessarily establishes a full log-Sobolev inequality rather than a MLSI. This also explains the spurious dependance on the parameter γ. Unfortunately, the proof uses heavily comparison tools that only seem to apply to LSI, and not directly to MLSI. Since the proof is very technical, and there are some simplifications that can be exploited in our setting, we have opted to carry out the proof in detail, adapting their approach directly to our setting.

Moreover, it should be mentioned that a key part of the argument in [Jai+23] was made fully rigorous only through the recent results by the same authors in [JPV26], which appeared during the final stages of our work.

Another motivation for giving a full proof is that the product structure of $\mathcal{G}_N = G^{\sqcup N}$, which is specific to the Kac system in our setting, allows one to track explicitly the joint

dependence of the log-Sobolev constant on both N and n, a dependence that cannot be read off directly from the general statement of [Jai+23], which applies in principle to fixed size independent sets on any graph. As we have already mentioned, the key feature of the bound is its uniformity in N.

2.1.5 Plan of the chapter

The structure of this chapter is parallel to that of Chapter 1. Some parts are omitted to avoid repetition, and the proof that transfers the MLSI constant of the particle system to the nonlinear process is not reproduced here, since it follows essentially the same argument as in Section 1.5.

Section 2.2 establishes the structural properties of the nonlinear hardcore dynamics, leading to the proof of Theorem 2.2. Section 2.3 contains the proof of Wasserstein contraction (Theorem 2.3) via an explicit coupling. The chapter concludes with Section 2.4 that is devoted to the uniform MLSI for the hardcore particle system (Theorem 2.6), proved via a comparison with the down-up walk and the local central limit theorem.

2.2 Structural properties of nonlinear equation

In this section, we present general properties of the mean-field nonlinear process associated with the hardcore model on a graph G, whose vertices are indexed by the set $[n] = \{1, \cdots, n\}$ and have maximum degree Δ. More precisely, we show that all strictly positive stationary measures for (2.4) are given by hardcore measures $\mu_{G,\lambda}$ with arbitrary fugacity $\lambda > 0$. We also prove that the density (2.1) is preserved along the evolution, and we conclude with the proof of the convergence result stated in Theorem 2.2.

We begin with the definition of a stationary measure. We denoted with Ω_G the set of all independent sets of G.

Definition 2.7. A measure $p \in \mathcal{P}(\Omega_G)$ is said to be stationary for the nonlinear dynamics defined in (2.4) if it satisfies

$$p \circ p = p.$$

Equivalently, p is stationary if and only if $p_t = p$ for all $t \geq 0$.

In particular, any probability measure p such that the product measure $p \otimes p$ is reversible with respect to $\mathcal{Q}$ is stationary. By the following lemma, it follows that every hardcore measure $\mu_{G,\lambda}$, for any $\lambda > 0$, is stationary.

Lemma 2.8. *Consider the collision kernel $\mathcal{Q}$ defined in (2.3). For every fugacity $\lambda > 0$, the product measure $\mu_{G,\lambda} \otimes \mu_{G,\lambda}$ is reversible with respect to $\mathcal{Q}$.*

Proof. It suffices to observe that if $\mathcal{Q}(\sigma, \sigma'; \tau, \tau') > 0$, then the kernel is symmetric in these pairs and satisfies $|\sigma| + |\sigma'| = |\tau| + |\tau'|$, since the mechanism (1.1) is conservative. Moreover, the measure $\mu_{G,\lambda} \otimes \mu_{G,\lambda}$ is constant on pairs with the same total cardinality. □

A complete characterization of stationary measures relies on entropy methods and is presented in the next subsection.

2.2.1 Entropy and stationary states

To characterize all stationary states, we use the relative entropy functional. The following result is a version of Boltzmann's H-theorem from kinetic theory. It holds in a more general setting (namely, reversible quadratic systems; see, for example, [CS18]). We have already proved it in Lemma 1.15 for the exchange model of the Ising system. We state it again here for the reader's convenience and for completenes.

Lemma 2.9. *Let p_t denote the solution of (2.4) with initial datum $p \in \mathcal{P}(\Omega_G)$, and let $\mu = \mu_{G,\lambda}$ for some arbitrary, fixed fugacity $\lambda > 0$. Then,*

$$\frac{d}{dt} H(p_t \mid \mu) = -\mathcal{D}_\mu(f_t, f_t), \qquad t \geq 0, \tag{2.9}$$

where $f_t := p_t/\mu$, and the entropy dissipation functional $\mathcal{D}_\mu$ is defined by

$$\mathcal{D}_\mu(f,f) = \frac{1}{4} \sum_{\sigma,\sigma',\tau,\tau'} \mu(\tau)\mu(\tau')\mathcal{Q}(\tau,\tau';\sigma,\sigma') \left(f(\tau)f(\tau') - f(\sigma)f(\sigma')\right) \log\left(\frac{f(\tau)f(\tau')}{f(\sigma)f(\sigma')}\right).$$

In particular, $H(p_t \mid \mu)$ is non-increasing in $t \geq 0$. Moreover, $p \in \mathcal{P}(\Omega)$ is stationary iff the density $f = p/\mu$ satisfies $f(\tau)f(\tau') = f(\sigma)f(\sigma')$ whenever $\mathcal{Q}(\tau,\tau';\sigma,\sigma') > 0$.

Denote with $\mathcal{P}_+(\Omega)$ the set of all probability measures on Ω strictly positive.

Lemma 2.10. *A measure $p \in \mathcal{P}_+(\Omega_G)$ is stationary if and only if $p = \mu_{G,\lambda}$ for some fugacity $\lambda > 0$.*

Proof. By Lemma 2.8 we know that the measures $\mu_{G,\lambda}$, for $\lambda > 0$, are stationary. We only need to show that if $p \in \mathcal{P}_+(\Omega_G)$ is stationary, then it must coincide with one of these hardcore measures.
Fix any $\lambda > 0$ and let

$$f = \frac{p}{\mu_{G,\lambda}} > 0$$

be the relative density of p with respect to $\mu_{G,\lambda}$. By the H-theorem of Lemma 2.9,

$$f(\sigma)\, f(\sigma') = f\big(S_\ell(\sigma, \sigma'_k)\big)\, f\big(S_k(\sigma', \sigma_\ell)\big), \tag{2.10}$$

for all $\sigma, \sigma' \in \Omega_G$ and all $\ell, k \in [n]$ such that the exchange in (1.1) produces a new pair in $\Omega_G \times \Omega_G$. Here, it is convenient to use the set notation for independent sets, in which a configuration is identified with the set of vertices assigned the value 1. Fix $\sigma = \{v_1, \dots, v_j\} \in \Omega_G$ with $v_i \in [n]$ for all $i \leq j$, and define the recursive sequence

$$\sigma(i) = \begin{cases} \sigma, & i = 0, \\ \sigma(i-1) \setminus \{v_i\}, & 0 < i \leq j. \end{cases}$$

This yields a telescopic representation of $f(\sigma)$:

$$f(\sigma) = f(\emptyset) \prod_{i=0}^{j-1} \frac{f(\sigma(i))}{f(\sigma(i+1))}. \tag{2.11}$$

Applying (2.10) with $\ell = k = v_{i+1}$ to the pair $(\sigma(i), \emptyset)$, which is admissible since $v_{i+1} \in \sigma(i)$ by construction, we obtain

$$\frac{f(\sigma(i))}{f(\sigma(i+1))} = \frac{f(\{v_{i+1}\})}{f(\emptyset)}.$$

Moreover, applying (2.10) again with $\ell = v_i$ and $k = v_{i'}$ for any $i, i' \leq j$ shows that

$$\frac{f(\{v_i\})}{f(\emptyset)} = \frac{f(\{v_{i'}\})}{f(\emptyset)},$$

so there exists a constant $C > 0$ such that all factors in the telescopic product (2.11) are equal to C. We deduce that

$$f(\sigma) = c\, C^{|\sigma|},$$

where $c = f(\emptyset)$ and $C > 0$ are constants independent of σ.
Since $f = p/\mu_{G,\lambda}$, this implies that p is itself a hardcore measure with fugacity $C\lambda$; that is, $p = \mu_{G,C\lambda}$. This proves the claim. □

Conserved quantity

Recall that the density of a given distribution $p \in \mathcal{P}(\Omega_G)$ is defined by $\alpha(p) = \frac{1}{n}\mathbb{E}_p(|\sigma|)$.

Lemma 2.11. *For all $p, q \in \mathcal{P}(\Omega_G)$ we have*

$$\alpha(p \circ q) \;=\; \frac{1}{2}\,\alpha(p) \;+\; \frac{1}{2}\,\alpha(q).$$

In particular, if $(p_t)_{t\geq 0}$ is the solution of (2.4), then $\frac{\mathrm{d}}{\mathrm{d}t}\alpha(p_t) = 0$.

Proof. Fix $p, q \in \mathcal{P}(\Omega)$. Recall $(\ell, k)_{\sigma,\sigma'}$ denotes the case when the ℓ, k-exchange is possible between σ and σ', and $\widehat{(\ell, k)}_{\sigma,\sigma'}$ is the complement. By definition, the density of distribution given by the collision (2.5) of $p, q \in \mathcal{P}(\Omega_G)$ is

$$\alpha(p \circ q) = \frac{1}{n} \sum_{\sigma,\sigma'} \frac{1}{2}(p(\sigma)q(\sigma') + p(\sigma')q(\sigma)) \sum_{\tau,\tau'} \mathcal{Q}(\sigma, \sigma'; \tau, \tau')\, |\tau| \tag{2.12}$$

As a first step, one has the computation

$$\sum_{\tau,\tau'} \mathcal{Q}(\sigma, \sigma'; \tau, \tau')\, |\tau| = |\sigma| + \frac{1}{n^2} \sum_{\ell,k=1}^{n} \mathbf{1}\{(\ell, k)_{\sigma,\sigma'}\}\, (\sigma'_k - \sigma_\ell). \tag{2.13}$$

Now, we rewrite (2.12) and apply (2.13)

$$\begin{aligned}\alpha(p\circ q) &= \frac{1}{n}\sum_{\sigma,\sigma'} p(\sigma)q(\sigma')\sum_{\tau,\tau'}\frac{1}{2}\Big(\mathcal{Q}(\sigma,\sigma';\tau,\tau')+\mathcal{Q}(\sigma',\sigma;\tau,\tau')\Big)|\tau| \\ &= \frac{1}{n}\sum_{\sigma,\sigma'} p(\sigma)q(\sigma')\left[\frac{|\sigma|}{2}+\frac{|\sigma'|}{2}+\frac{1}{2}\sum_{\ell,k=1}^{n}\Big(\mathbf{1}\{(\ell,k)_{\sigma,\sigma'}\}(\sigma'_k-\sigma_\ell)+\mathbf{1}\{(\ell,k)_{\sigma',\sigma}\}(\sigma_k-\sigma'_\ell)\Big)\right] \\ &= \frac{1}{2}\,\alpha(p)+\frac{1}{2}\,\alpha(q),\end{aligned}$$

where in the last equality we used the identity

$$\sum_{\ell,k=1}^{n}\mathbf{1}\{(\ell,k)_{\sigma,\sigma'}\}(\sigma'_k-\sigma_\ell) = -\sum_{\ell,k=1}^{n}\mathbf{1}\{(\ell,k)_{\sigma',\sigma}\}(\sigma_k-\sigma'_\ell),$$

which follows from the change of variables $\ell \leftrightarrow k$ and the symmetry $(\ell,k)_{\sigma,\sigma'} = (k,\ell)_{\sigma',\sigma}$. This concludes the proof. □

The next lemma establishes that, in the regime of convergence $(0, \frac{1}{\Delta+1})$ described in Theorem 2.2, each hardcore model is uniquely parametrized by its fugacity and corresponds to a unique density, and vice versa.

Combined with the previous lemma, which shows that the density is conserved along the dynamics, this justifies the uniqueness in Theorem 2.2.

Lemma 2.12. *Consider a graph G on n vertices with maximum degree $\Delta \geq 3$. For every density $\bar{\alpha} \in (0, \frac{1}{\Delta+1})$ there exists a unique fugacity $\lambda \in (0,+\infty)$ such that $\alpha(\mu_{G,\lambda}) = \bar{\alpha}$. In particular, the density function $\lambda \mapsto \alpha(\mu_{G,\lambda})$ is a bijection from $(0,+\infty)$ onto $(0, \frac{1}{\Delta+1})$ which restricts to a bijection between the subcritical intervals $(0,\lambda_c(\Delta))$ and $(0,\alpha_c(\Delta))$.*

Proof. We first establish the bijection between $(0,\lambda_c(\Delta))$ and $(0,\alpha_c(\Delta))$, and then extend it to the full range.

Strict monotonicity. The identity

$$\alpha'_G(\lambda) = \frac{1}{\lambda n}\operatorname{Var}_{\mu_{G,\lambda}}(|\sigma|) > 0, \qquad \forall \lambda > 0,$$

shows that α_G is strictly increasing on $(0,+\infty)$, hence injective.

Surjectivity onto $(0,\alpha_c(\Delta))$. We use the fact [CR14, Theorem 12] that, among all graphs of maximum degree Δ, the occupancy ratio is minimized by the clique $K_{\Delta+1}$: for all $\lambda > 0$,

$$\alpha_G(\lambda) \geq \alpha_{K_{\Delta+1}}(\lambda) = \frac{\lambda}{1+\lambda(\Delta+1)}. \tag{2.14}$$

In particular, $\alpha_c(\Delta) = \alpha_{K_{\Delta+1}}(\lambda_c(\Delta))$. Fix $\alpha \in (0,\alpha_c(\Delta))$ and set

$$\lambda^* = \frac{\alpha}{1-\alpha(\Delta+1)},$$

which is well defined since $\alpha < \alpha_c(\Delta) < \frac{1}{\Delta+1}$ and satisfies $\alpha_{K_{\Delta+1}}(\lambda^*) = \alpha$. By (2.14), $\alpha_G(\lambda^*) \geq \alpha$, and since $\alpha < \alpha_c(\Delta)$ we also have $\lambda^* < \lambda_c(\Delta)$. As α_G is continuous with $\alpha_G(0^+) = 0$, the intermediate value theorem yields $\bar{\lambda} \in (0, \lambda^*] \subset (0, \lambda_c(\Delta))$ such that $\alpha_G(\bar{\lambda}) = \alpha$.

Extension to $(0, \frac{1}{\Delta+1})$. The function $\alpha_{K_{\Delta+1}}$ is a strictly increasing continuous bijection from $(0, +\infty)$ onto $(0, \frac{1}{\Delta+1})$. Hence, for every $\bar{\alpha} \in [\alpha_c(\Delta), \frac{1}{\Delta+1})$, there exists $\lambda^* \geq \lambda_c(\Delta)$ with $\alpha_{K_{\Delta+1}}(\lambda^*) = \bar{\alpha}$. Arguing as above via (2.14) and the intermediate value theorem, we obtain $\bar{\lambda} \in (0, \lambda^*]$ such that $\alpha_G(\bar{\lambda}) = \bar{\alpha}$. □

2.2.2 Proof of Theorem 2.2

In Section 1.2.4 of Chapter 1, we have shown that for the nonlinear spin system associated with the Ising model, convergence to the stationary solution can be proved by establishing irreducibility, that is, eventual positivity of the solution in time. A careful check of that proof reveals that the same argument applies in our current context of the hardcore model. In particular, we now focus only on proving the irreducibility statement below.

Theorem 2.13 (Irreducibility)**.** *Assume* $\alpha \in (0, \frac{1-\delta}{\Delta+1})$*, with* $\alpha \in (0, 1)$*. Then there exist a time* $t_0 = t_0(\alpha, \Delta, \delta)$ *and a positive constant* $\varepsilon_0 = \varepsilon_0(\alpha, \Delta, \delta)$ *such that for all* $t \geq t_0$ *the solution* p_t *of the nonlinear spin system* (2.4) *in the mean-field case satisfies*

$$p_t(\tau) \geq \varepsilon_0, \quad \forall \tau \in \Omega_G.$$

In contrast with the case of the Ising model, for the Hardcore model there is no simple non-interacting system to compare with, as in the Ising model case, where a statement as in Theorem 2.13 was obtained by comparison with the $J = 0$ case. This makes the problem for the Hardcore model more challenging. Nevertheless, part of the strategy can still be applied. In particular, a reduction to the discrete-time case is still possible. A straightforward adaptation of Lemma 1.21, in particular, see inequalities (1.32), (1.33), and (1.34), shows that it is sufficient to show that there exist a positive constant $\varepsilon := \varepsilon(\alpha, \Delta, \alpha)$ and a positive integer $u := u(n)$ such that

$$\Phi_u(p_1, \ldots, p_{2^u})(\tau) \geq \varepsilon, \quad \forall \tau \in \Omega_G, \tag{2.15}$$

for every choice of $(p_1, \ldots, p_{2^u})$ such that $\alpha(p_i) = \alpha$. The definition of operator Φ_u was given in Subsection 1.2.6.

Lemma 2.14. *The statement* (2.15) *holds with* $u = 2n - 1$ *and* $\varepsilon = n! \left(\frac{\alpha\delta}{2n^2}\right)^n \frac{\Delta+\delta}{\Delta+1}$.

Proof. We divide the proof into several steps. Informally, we first show in Step 2 that there is a uniformly positive probability of reaching the empty configuration after n steps of the discrete evolution. We then prove in Step 3 that, starting from the empty set, any prescribed configuration can be reached with uniformly positive probability in $n-1$ further steps.

For every $m \in \{1, \dots, n\}$, let A_m be a uniformly random subset of $[n]$ of cardinality m, and write $\mathbb{E}_{A_m}$ for the expectation with respect to A_m. We write $\mathcal{P}_\alpha(\Omega)$ for the set of distributions $p \in \mathcal{P}(\Omega)$ with density α.

Step 1. For every $t \in \mathbb{N}$ and every $\tau \sim \Phi_t(p_1, \dots, p_{2^t})$ with $\alpha(p_i) = \alpha$,

$$\mathbb{E}_{A_1}\big[\mathbb{P}(\tau_{N^+(A_1)} = 0)\big] \geq \delta.$$

By the conservation law, $\alpha(\Phi_t(p_1, \dots, p_{2^t})) = \alpha$, and therefore

$$\mathbb{E}_{A_1}\big[\mathbb{P}(\exists\, u \in N^+(A_1) : \tau_u = 1)\big] \leq \frac{1}{n} \sum_{\ell \in [n]} \sum_{u \in N^+(\ell)} \mathbb{P}(\tau_u = 1) \leq (\Delta + 1)\alpha \leq 1 - \delta.$$

Step 2. For all $m \in \{1, \dots, n\}$ and all $p_1, \dots, p_{2^m} \in \mathcal{P}_\alpha(\Omega_G)$, if $\tau \sim \Phi_m(p_1, \dots, p_{2^m})$, then

$$\mathbb{E}_{A_m}\big[\mathbb{P}(\tau_{A_m} = \mathbf{0})\big] \geq \Big(\frac{\delta}{n}\Big)^m m! \, \frac{\Delta + \delta}{\Delta + 1}. \tag{2.16}$$

In particular, for $m = n$ we have $A_n = [n]$, and $\tau \sim \Phi_{n-1}(p_1, \dots, p_{2^{n-1}})$ satisfies

$$\mathbb{P}(\tau = 0) \geq \Big(\frac{\delta}{n}\Big)^n n! \, \frac{\Delta + \delta}{\Delta + 1}.$$

We prove (2.16) by induction on m. For $m = 1$, let $\tau \sim p$ with $\alpha(p) = \alpha$; then

$$\mathbb{E}_{A_1}\big[\mathbb{P}(\tau_{A_1} = 0)\big] = 1 - \alpha \geq \frac{\Delta + \delta}{\Delta + 1},$$

which proves the case $m = 1$.

Suppose now $\tau \sim \Phi_1(p, q) = p \circ q$, and draw independently $\sigma \sim p$ and $\sigma' \sim q$, both with density α. By the definition of the collision operator (2.5), with probability $1/2$ the configuration τ is obtained from σ by replacing σ_ℓ with σ'_k, where $\ell, k \in [n]$ are chosen uniformly at random and the swap is performed only if admissible; otherwise $\tau = \sigma$. With the remaining probability $1/2$, the roles of σ and σ' are exchanged.
Fix $A \subseteq [n]$. Conditionally on ℓ and k, the event $\{\tau_A = 0\}$ contains $\{\sigma_{A \setminus \{\ell\}} = 0\} \cap \{\sigma'_{N^+(k)} = 0\}$ in the first case, and the analogous event with σ and σ' interchanged in the second. Averaging over ℓ and k, both uniform on $[n]$ and independent of σ, σ',

$$\mathbb{P}(\tau_A = 0) \geq \frac{1}{2n} \sum_{\ell \in A} \Big[\mathbb{P}(\sigma_{A \setminus \{\ell\}} = 0)\, \mathbb{E}_k\big[\mathbb{P}(\sigma'_{N^+(k)} = 0)\big] + \mathbb{P}(\sigma'_{A \setminus \{\ell\}} = 0)\, \mathbb{E}_k\big[\mathbb{P}(\sigma_{N^+(k)} = 0)\big]\Big].$$

By Step 1,

$$\mathbb{P}(\tau_A = 0) \geq \frac{\delta}{2n} \sum_{\ell \in A} \Big[\mathbb{P}(\sigma_{A \setminus \{\ell\}} = 0) + \mathbb{P}(\sigma'_{A \setminus \{\ell\}} = 0)\Big]. \tag{2.17}$$

We now combine (2.17) with the induction hypothesis. Fix $m < n$ and assume (2.16) holds at level m. Given $p_1, \dots, p_{2^{m+1}} \in \mathcal{P}_\alpha(\Omega_G)$, set $p = \Phi_m(p_1, \dots, p_{2^m}) \in \mathcal{P}_\alpha(\Omega_G)$ and $q = \Phi_m(p_{2^m+1}, \dots, p_{2^{m+1}}) \in \mathcal{P}_\alpha(\Omega_G)$, so that $\tau \sim \Phi_2(p, q) = \Phi_{m+1}(p_1, \dots, p_{2^{m+1}})$.

Applying (2.17) with $A = A_{m+1}$ and integrating, we obtain

$$\mathbb{E}_{A_{m+1}}\big[\mathbb{P}(\tau_{A_{m+1}} = 0)\big] \geq \frac{\delta}{2}\,\frac{m+1}{n}\,\mathbb{E}_{A_{m+1}}\left[\frac{1}{m+1}\sum_{\ell\in A_{m+1}}\big(\mathbb{P}(\sigma_{A_{m+1}\setminus\{\ell\}} = 0) + \mathbb{P}(\sigma'_{A_{m+1}\setminus\{\ell\}} = 0)\big)\right].$$

Removing a uniformly random element from a uniformly random subset of size $m+1$ yields a uniformly random subset of size m; therefore $A_{m+1} \setminus \{\ell\} \stackrel{d}{=} A_m$ when ℓ is uniform on A_{m+1}. Combining this with the induction hypothesis,

$$\mathbb{E}_{A_{m+1}}\big[\mathbb{P}(\tau_{A_{m+1}} = 0)\big] \geq \frac{\delta}{n}(m+1)\,\mathbb{E}_{A_m}\big[\mathbb{P}(\sigma_{A_m} = 0)\big] \geq \frac{\delta}{n}(m+1)\Big(\frac{\delta}{n}\Big)^m m!\,\frac{\Delta+\delta}{\Delta+1},$$

which is (2.16) at level $m+1$. This completes the induction.

Step 3. For all $\eta \in \Omega_G$ and all $p_1, \ldots, p_{2^{2n-1}} \in\in \mathcal{P}_\alpha(\Omega_G)$, if $\tau \sim \Phi_{2n-1}(p_1, \ldots, p_{2^{2n-1}})$, then

$$\mathbb{P}(\tau = \eta) > n!\left(\frac{\alpha\delta}{2n^2}\right)^n \frac{\Delta+\delta}{\Delta+1}.$$

We shall show that for every time $t \geq 1$, every $\eta \in \Omega_G$ and every vector of distributions $p_1, \cdots, p_{2^t} \in \mathcal{P}_\alpha(\Omega_G)$,

$$\begin{aligned}\mathbb{P}_{\tau\sim\Phi_t(p_1,\cdots,p_{2^t})}(\tau = S_u(\eta,1)) \geq \frac{\alpha}{2n}\mathbb{P}_{\tau\sim\Phi_t(p_1,\cdots,p_{2^{t-1}})}(\tau = S_u(\eta,1)), \qquad (2.18)\\ \forall u \in [n]:\ S_u(\eta,1) \in \Omega_G,\ \eta_u = 0.\end{aligned}$$

To see this, sample independently $\sigma \sim \Phi_{t-1}(p_1, \cdots, p_{2^{t-1}})$ and $\sigma' \sim \Phi_{t-1}(p_{2^{t-1}}, \cdots, p_{2^t})$, let $\ell, k \in [n]$ be two elements chosen independently and uniformly at random in $[n]$, and toss a fair coin $B \in \{0,1\}$. Then

$$\tau = \begin{cases} S_\ell(\sigma, \sigma'_k) & \text{if } B = 1 \\ S_k(\sigma', \sigma_\ell) & \text{if } B = 0 \end{cases}$$

has distribution $\Phi_t(p_1, \cdots, p_{2^t})$. Here, we used the convention that $S_u(\sigma, a) = \sigma$ when $S_u(\sigma, a) \notin \Omega_G$. The event $\{\tau = S_u(\eta, 1)\}$ contains $\{B = 1, \sigma = \eta, \sigma'_u = 1, \ell = u\}$, which gives the desired bound.

Iterating (2.18), the probability of reaching a given independent set η at time $t = 2n-1$, starting from the vector of distributions $p_1, \ldots, p_{2^{2n-1}}$, can be bounded below by

$$\mathbb{P}_{\tau\sim\Phi_{2n-1}(p_1,\ldots,p_{2^{2n-1}})}(\tau = \eta) \geq \left(\frac{\alpha}{2n}\right)^{|\eta|}\mathbb{P}_{\tau\sim\Phi_{2n-1-|\eta|}(p_1,\ldots,p_{2^{2n-1-|\eta|}})}(\tau = \mathbf{0}).$$

Indeed, removing one element at a time, each removal costs at most $\alpha/2n$, until all vertices of η have been removed after $|\eta|$ steps. On the other hand, since $2n - 1 - |\eta| \geq n - 1$, considering the collisions on the sub-trees and using the conservation law, one can find a

vector of distributions $q_1, \cdots, q_{2^{n-1}} \in \mathcal{P}_\alpha(\Omega_G)$ such that

$$\Phi_{2n-1-|\eta|}(p_1, \cdots p_{2^{2n-1-|\eta|}}) = \Phi_{n-1}(q_1, \cdots, q_{2^{n-1}}).$$

Then, using Step 2,

$$\begin{aligned}\mathbb{P}_{\tau\sim\Phi_{2n-1-|\eta|}(p_1,\ldots,p_{2^{2n-1-|\eta|}})}(\tau = \mathbf{0}) &= \mathbb{P}_{\tau\sim\Phi_{n-1}(q_1,\cdots,q_{2^{n-1}})}(\tau = \mathbf{0}) \\ &\geq \left(\frac{\delta}{n}\right)^n n! \, \frac{\Delta+\delta}{\Delta+1}.\end{aligned}$$

Combining the previous inequalities, we conclude the proof of Step 3:

$$\mathbb{P}_{\tau\sim\Phi_{n-1}(p_1,\cdots,p_{2^{n-1}})}(\tau = \eta) \geq \left(\frac{\alpha}{2n}\right)^{|\eta|} \left(\frac{\delta}{n}\right)^{n-1} n! \, \frac{\Delta+\delta}{\Delta+1} > n! \left(\frac{\alpha\delta}{2n^2}\right)^n \frac{\Delta+\delta}{\Delta+1}.$$

To prove the lemma, it remains to extend the above bound from time $2n-1$ to all later times. Let $t \geq 2n-1$ and set $s := t-(2n-1)$. Considering the collisions on the sub-trees of height $s > 0$ and using the conservation law, as in the proof of Step 3 above, one can find a vector of distributions $q_1, \cdots, q_{2^{2n-1}} \in \mathcal{P}_\alpha(\Omega_G)$ such that

$$\Phi_{2n-1}(q_1, \cdots, q_{2^{2n-1}}) = \Phi_t(p_1, \cdots, p_{2^t}).$$

We may then apply Step 3 to $\Phi_{2n-1}(q_1, \cdots, q_{2^{2n-1}})$, and therefore, for every $\eta \in \Omega_G$,

$$\mathbb{P}_{\tau\sim\Phi_t(p_1,\cdots,p_{2^t})}(\tau = \eta) = \mathbb{P}_{\tau\sim\Phi_{n-1}(q_1,\cdots,q_{2^{2n-1}})}(\tau = \eta) \geq n! \left(\frac{\alpha\delta}{2n^2}\right)^n \frac{\Delta+\delta}{\Delta+1}.$$

Thus the irreducibility condition holds with

$$\varepsilon := n! \left(\frac{\alpha\delta}{2n^2}\right)^n \frac{\Delta+\delta}{\Delta+1}.$$

This concludes the proof. □

2.3 Wasserstein Contraction for low fugacity

In this section we prove Theorem 2.3, which states that the solution p_t of the nonlinear equation (2.4) with density $\alpha(p_t) < \frac{1}{3(\Delta+1)}$ satisfies the exponential decay

$$W_1(p_t, \mu_{G,\lambda}) \leq e^{-\varepsilon t/n} \, W_1(p_0, \mu_{G,\lambda}) \tag{2.19}$$

in the 1-Wasserstein distance associated with the Hamming metric, where $\varepsilon = \varepsilon(\alpha) = 1 - 3(\Delta+1)\alpha$ and $\mu_{G,\lambda}$ is the unique hardcore measure with $\alpha(\mu_{G,\lambda}) = \alpha(p_0)$, whose existence is given by Lemma 2.12. An analogous result in the multi-component mean-field setting was established for the Ising model in Section 1.3. While the overall strategy is similar, the coupling used here is rather different.

The proof reduces to a contraction property of the collision operator $T(p,q) = p \circ q$ in Wasserstein distance, obtained via a coupling argument. The key estimate is the following.

Lemma 2.15. *Let $p_i, q_i \in \mathcal{P}(\Omega_G)$, $i \in \{1,2\}$, all have common density $\alpha < \frac{1}{3(\Delta+1)}$, and set $\varepsilon = 1 - 3(\Delta+1)\alpha$. Then*

$$W_1(T(p_1,q_1),\, T(p_2,q_2)) \le \frac{1}{4}\Big(1 - \frac{\varepsilon}{n}\Big)\Big(W_1(p_1,p_2) + W_1(p_1,q_2) + W_1(q_1,p_2) + W_1(q_1,q_2)\Big). \tag{2.20}$$

The fact that Lemma 2.15 implies (2.19), and hence Theorem 2.3, was already discussed in Subsection 1.3.1 for the Ising model. Once the contraction property (2.20) is established, the argument depends only on the graphical representation of the solution, which is independent of the underlying model; we therefore omit the details and refer the reader to that subsection.
The remainder of this section is devoted to the proof of Lemma 2.15. We begin with the preliminary tools needed for the coupling argument that lies at its core.

Given $\ell \in [n], \sigma, \tau \in \Omega_G$, recall that $(\ell,k)_{\sigma,\sigma'}$ is the event that the exchange (ℓ,k) is admissible for the pair of independent sets $\sigma, \sigma' \in \Omega_G$, and $\widehat{(\ell,k)}_{\sigma,\sigma'}$ denotes its complementary event. Define

$$\Sigma_p(\tau,\sigma,\ell) := \frac{1}{n}\sum_{k,\sigma'} p(\sigma')\Big(\mathbf{1}\big((\ell,k)_{\sigma,\sigma'}\big)\,\mathbf{1}\big(\tau = S_\ell(\sigma,\sigma'_k)\big) + \mathbf{1}\big(\widehat{(\ell,k)}_{\sigma,\sigma'}\big)\,\mathbf{1}(\tau=\sigma)\Big).$$

that is the probability, given σ, to see τ by updating σ at ℓ-coordinate. According to this definition, for any given pair of distribution $p, q \in \mathcal{P}(\Omega_G)$

$$T(p,q)(\tau) = \frac{1}{2n}\sum_{\ell\in[n],\sigma\in\Omega_G} p(\sigma)\Sigma_p(\tau,\sigma,\ell) + \frac{1}{2n}\sum_{\ell\in[n],\sigma\in\Omega_G} q(\sigma)\Sigma_p(\tau,\sigma,\ell). \tag{2.21}$$

where $T(p,q) := p \circ q$ is defined in (2.5). Now, we present a useful preliminary result.

Proposition 2.16. *Let $p, q \in \mathcal{P}(\Omega_G)$ be such that $\alpha(p) = \alpha(q)$. For any $\ell \in [n]$, $\sigma \in \Omega_G$, if $\sigma_\ell = 0$ then*

$$\Sigma_p(\tau,\sigma,\ell) = \Sigma_q(\tau,\sigma,\ell), \quad \forall \tau \in \Omega_G \tag{2.22}$$

Proof. Fix $\sigma \in \Omega_G$ and $\ell \in [n]$. If $\sigma_\ell = 0$, we distinguish two cases:

1. $N(\ell) \cap \sigma \neq \emptyset$, and $\sigma_\ell = 0$;
2. $N(\ell) \cap \sigma = \emptyset$ and $\sigma_\ell = 0$.

The proof is completed by showing that (2.22) holds in both cases. In the first case, we

necessarily have $\tau = \sigma$, even if the (ℓ, k)-exchange occurs. This implies

$$\begin{aligned}\Sigma_p(\tau, \sigma, \ell) &= \frac{1}{n} \sum_{\sigma', k} p(\sigma') \mathbf{1}(\tau = \sigma) \\ &= \mathbf{1}(\sigma = \tau) \\ &= \Sigma_q(\tau, \sigma, \ell).\end{aligned}$$

In the second case, the event $(\ell, k)_{\sigma,\sigma'}$ holds independently of σ' and $k \in [n]$. Therefore,

$$\begin{aligned}\Sigma_p(\tau, \sigma, \ell) &= \frac{1}{n} \sum_{k \in [n], \sigma'} p(\sigma' \mid \sigma'_k = 1)\, p_k(1)\, \mathbf{1}(\tau = S_\ell(\sigma, 1)) \\ &\quad + \frac{1}{n} \sum_{k \in [n], \sigma'} p(\sigma' \mid \sigma'_k = 0)\, p_k(0)\, \mathbf{1}(\tau = S_\ell(\sigma, 0)) \\ &= \frac{\mathbf{1}(\tau = S_\ell(\sigma, 1))}{n} \sum_{k \in [n]} p_k(1) \sum_{\sigma'} p(\sigma' \mid \sigma'_k = 1) \\ &\quad + \frac{\mathbf{1}(\tau = S_\ell(\sigma, 0))}{n} \sum_{k \in [n]} p_k(0) \sum_{\sigma'} p(\sigma' \mid \sigma'_k = 0) \\ &= \alpha(p)\, \mathbf{1}(\tau = S_\ell(\sigma, 1)) + (1 - \alpha(p))\, \mathbf{1}(\tau = S_\ell(\sigma, 0)) \\ &= \Sigma_q(\tau, \sigma, \ell),\end{aligned}$$

which concludes the proof of (2.22).

□

2.3.1 Proof of Lemma 2.15

By the definition of the collision operator (2.21), we have the convex decomposition

$$T(p, q) = \tfrac{1}{2}\widehat{T}(p, q) + \tfrac{1}{2}\widehat{T}(q, p),$$

where $\widehat{T}$ is the non-symmetric operator defined by

$$\widehat{T}(p, q)(\tau) = \frac{1}{n} \sum_{\ell \in [n]} \sum_{\sigma \in \Omega_G} p(\sigma)\, \Sigma_q(\tau, \sigma, \ell).$$

Let $p_i, q_i \in \mathcal{P}(\Omega_G)$, $i \in \{1, 2\}$, all have common density $\alpha < \frac{1}{3(\Delta+1)}$. By the convexity of the Wasserstein distance, to prove (2.20) it suffices to establish the following contraction property for $\widehat{T}$:

$$W_1(\widehat{T}(p_1, q_1),\, \widehat{T}(p_2, q_2)) \le \left[1 - \frac{1 - 2(\Delta+1)\alpha}{n}\right] W_1(p_1, p_2) + \frac{\alpha(\Delta+1)}{n}\, W_1(q_1, q_2). \quad (2.23)$$

We now turn to the proof of (2.23), which relies on an explicit coupling between $\widehat{T}(p_1, q_1)$ and $\widehat{T}(p_2, q_2)$. Recall that, for any pair of independent sets $\sigma, \sigma' \in \Omega_G$, the *discrepancy set*

is

$$D_{\sigma,\sigma'} = \{\ell \in [n] : \sigma_\ell \neq \sigma'_\ell\}.$$

We denote by $\pi^p_{\text{opt}} \in \Gamma(p_1, p_2)$ an optimal coupling realizing $W_1(p_1, p_2)$, i.e., $W_1(p_1, p_2) = \mathbb{E}_{(\sigma,\sigma')\sim\pi^p_{\text{opt}}}[|D_{\sigma,\sigma'}|]$, and analogously $\pi^q_{\text{opt}} \in \Gamma(q_1, q_2)$ for $W_1(q_1, q_2)$. The coupling between $\widehat{T}(p_1, q_1)$ and $\widehat{T}(p_2, q_2)$ is described in Algorithm 2.

Algorithm 2 Coupling between $\widehat{T}(p_1, q_1)$ and $\widehat{T}(p_2, q_2)$

Require: Distributions $p_i, q_i \in \mathcal{P}(\Omega_G)$ for $i \in \{1, 2\}$; optimal couplings $\pi^p_{\text{opt}} \in \Gamma(p_1, p_2)$ and $\pi^q_{\text{opt}} \in \Gamma(q_1, q_2)$ with respect to the expected Hamming distance.

Ensure: Coupled outputs $\tau^1 \sim \widehat{T}(p_1, q_1)$ and $\tau^2 \sim \widehat{T}(p_2, q_2)$.

Sample $(\sigma^{p_1}, \sigma^{p_2}) \sim \pi^p_{\text{opt}}$.
Sample $\ell, k \in [n]$ uniformly at random.
if $\sigma^{p_1}_\ell = 1,\ \sigma^{p_2}_\ell = 0$ **then**
    Sample $\sigma'^{q_1} \sim q_1$ and set $\sigma'^{q_2} = \sigma'^{q_1}$.
else if $\sigma^{p_2}_\ell = 0,\ \sigma^{p_1}_\ell = 1$ **or** $\sigma^{p_2}_\ell = \sigma^{p_1}_\ell = 0$ **then**
    Sample $\sigma'^{q_1} \sim q_1$ and set $\sigma'^{q_2} = \sigma'^{q_1}$.
else ▷ $\sigma^{p_1}_\ell = \sigma^{p_2}_\ell = 1$
    Sample $(\sigma'^{q_1}, \sigma'^{q_2}) \sim \pi^q_{\text{opt}}$.
end if
Sample $(\tau^1, \widehat{\tau}^1) \sim Q_{\ell,k}(\sigma^{p_1}, \sigma'^{q_1}; \cdot, \cdot)$ and output τ^1.
Sample $(\tau^2, \widehat{\tau}^2) \sim Q_{\ell,k}(\sigma^{p_2}, \sigma'^{q_2}; \cdot, \cdot)$ and output τ^2.
return (τ^1, τ^2).

By Proposition 2.16, this procedure defines a valid coupling between $\widehat{T}(p_1, q_1)$ and $\widehat{T}(p_2, q_2)$.

Let $\sigma^{p_1}, \sigma^{q_1}$ be the configurations producing τ^1, and $\sigma^{p_2}, \sigma^{q_2}$ those producing τ^2, according to the algorithm. The argument relies on controlling the evolution of discrepancies. Let ℓ, k be chosen as above.

For every realization of the random variables, we have

$$\begin{aligned} |D_{\tau^1,\tau^2}| - |D_{\sigma^{p_1},\sigma^{p_2}}| \leq -\mathbf{1}(\ell \in D_{\sigma^{p_1},\sigma^{p_2}}) \qquad (2.24)\\ (A) \quad + \mathbf{1}(\sigma^{p_1}_\ell = 1, \sigma^{p_2}_\ell = 0, \exists\, u \in N(k) : \sigma^{q_1}_k = 1)\\ (B) \quad + \mathbf{1}(\sigma^{p_1}_\ell = 0, \sigma^{p_2}_\ell = 1, \exists\, u \in N(k) : \sigma^{q_2}_k = 1)\\ (C) \quad + \mathbf{1}(\sigma^{q_1}_k = \sigma^{q_2}_k = 1,\ N(\ell) \cap D_{\sigma^{p_1},\sigma^{p_2}} \neq \emptyset)\\ (D) \quad + \mathbf{1}(\sigma^{p_1}_\ell = \sigma^{p_2}_\ell = 1,\ N^+(k) \cap D_{\sigma^{q_1},\sigma^{q_2}} \neq \emptyset). \end{aligned}$$

Proof of Inequality (2.24). We distinguish the three possible cases according to the number of discrepancies, namely $|D_{\tau^1,\tau^2}| - |D_{\sigma^{p_1},\sigma^{p_2}}| \in \{-1, 0, 1\}$. In the case $|D_{\tau^1,\tau^2}| - |D_{\sigma^{p_1},\sigma^{p_2}}| = -1$, inequality (2.24) is immediately satisfied.

$\boxed{|D_{\tau^1,\tau^2}| - |D_{\sigma^{p_1},\sigma^{p_2}}| = 0.}$ If $\ell \notin D_{\sigma^{p_1},\sigma^{p_2}}$, then (2.24) clearly holds. Assume now that $\ell \in |D_{\sigma^{p_1},\sigma^{p_2}}|$. Since the number of discrepancies is preserved, the discrepancy at site ℓ

must be maintained by the exchange. There are two possible cases: $\sigma_\ell^{p_1} = 1, \sigma_\ell^{p_2} = 0$ and $\sigma_\ell^{p_1} = 0, \sigma_\ell^{p_2} = 1$. We start with the first one.

By the coupling rule, we have $\sigma_k^{q_1} = \sigma_k^{q_2}$, hence either $\sigma_k^{q_1} = \sigma_k^{q_2} = 0$ or $\sigma_k^{q_1} = \sigma_k^{q_2} = 1$.

1. $\sigma_k^{q_1} = \sigma_k^{q_2} = 0$: in order to maintain the discrepancy, it must be that $\widehat{(\ell, k)}_{\sigma^{p_1}, \sigma^{p_2}}$ holds, that is, there exists $u \in N(k)$ such that $\sigma_u^{q_1} = 1$. In this case $A = 1$, and inequality (2.24) is satisfied.

2. $\sigma_k^{q_1} = \sigma_k^{q_2} = 1$: in order to maintain the discrepancy, $\widehat{(\ell, k)}_{\sigma^q, \sigma'^q}$ must hold. Therefore there exists $u \in N(\ell)$ such that $\sigma_u^{p_1} = 1$. Since $\sigma_\ell^{p_1} = 1$ and $\sigma_\ell^{p_2} = 0$, this implies that there is a discrepancy between σ^{p_1} and σ^{p_2} in the neighborhood $N(\ell)$. Hence $N(\ell) \cap D_{\sigma^{p_1}, \sigma^{p_2}} \neq \emptyset$ and $C = 1$.

The case $\sigma_\ell^{p_1} = 0, \sigma_\ell^{p_2} = 1$ is treated in the same way and leads to $B = 1$ when $\sigma_k^{q_1} = \sigma_k^{q_2} = 0$ and to $C = 1$ when $\sigma_k^{q_1} = \sigma_k^{q_2} = 1$.

$\boxed{|D_{\tau^1, \tau^2}| - |D_{\sigma^{p_1}, \sigma^{p_2}}| = 1.}$ In this case the discrepancy increases after the exchange. This implies that $\ell \notin D_{\sigma^{p_1}, \sigma^{p_2}}$ and $\ell \in D_{\tau^1, \tau^2}$. We distinguish two cases: $\sigma_\ell^{p_1} = \sigma_\ell^{p_2} = 0$ and $\sigma_\ell^{p_1} = \sigma_\ell^{p_2} = 1$.

1. $\sigma_\ell^{p_1} = \sigma_\ell^{p_2} = 0$: Here we know that $\sigma_\ell^{q_1} = \sigma_\ell^{q_2}$. To create a discrepancy, it must be that $\sigma_\ell^{q_1} = \sigma_\ell^{q_2} = 1$ and there must exist a discrepancy between σ^{p_1} and σ^{p_2} in the neighborhood $N(\ell)$. Therefore $C = 1$ and the inequality (2.24) holds.

2. $\sigma_\ell^{p_1} = \sigma_\ell^{p_2} = 1$: There are two possible ways to create a discrepancy. The first one is that $(\ell, k)_{\sigma^{p_1}, \sigma^{q_1}} \cap (\ell, k)_{\sigma^{p_2}, \sigma^{q_2}}$ holds but $\sigma_k^{q_1} \neq \sigma_k^{q_2}$, so that both exchanges are available but the updates differ. The second one is that the exchange is available in one of the two only and thus $N(k) \cap D_{\sigma^{q_1}, \sigma^{q_2}} \neq \emptyset$. In both situations, there must exist a discrepancy between σ^{q_1} and σ^{q_2} in the extended neighborhood $N^+(k) = \{k\} \cup N(k)$. Hence $D = 1$ and inequality (2.24) is satisfied.

Let $\alpha = \alpha(p_i) = \alpha(q_i)$ for $i \in \{1, 2\}$. We are going to estimate the expected value of terms A, B, C and D that appear in (2.24).

Estimation of the terms A and B. The expected value of A satisfies

$$\begin{aligned}
\mathbb{E}[A] &= \frac{1}{n^2} \sum_{\ell,k=1}^{n} \mathbb{P}(\sigma_\ell^{p_1} = 1, \sigma_\ell^{p_2} = 0, \exists\, u \in N(k) : \sigma_k^{q_1} = 1) \\
&= \frac{1}{n^2} \sum_{\ell,k=1}^{n} \mathbb{P}(\exists\, u \in N(k) : \sigma_k^{q_1} = 1 \mid \sigma_\ell^{p_1} = 1, \sigma_\ell^{p_2} = 0) \mathbb{P}(\sigma_\ell^{p_1} = 1, \sigma_\ell^{p_2} = 0) \\
&= \left[\frac{1}{n} \sum_{k=1}^{n} \mathbb{P}_{\sigma \sim q_1} \Big(\bigcup_{u \in N(k)} \{\sigma_u = 1\} \Big) \right] \left[\sum_{\ell=1}^{n} \frac{\mathbb{P}(\sigma_\ell^{p_1} = 1, \sigma_\ell^{p_2} = 0)}{n} \right] \\
&\leq \left[\frac{1}{n} \sum_{k=1}^{n} \sum_{u \in N(k)} \mathbb{P}_{\sigma \sim q_1}(\sigma_u = 1) \right] \left[\sum_{\ell=1}^{n} \frac{\mathbb{P}(\sigma_\ell^{p_1} = 1, \sigma_\ell^{p_2} = 0)}{n} \right] \\
&\leq \Delta\, \alpha \sum_{\ell=1}^{n} \frac{\mathbb{P}(\sigma_\ell^{p_1} = 1, \sigma_\ell^{p_2} = 0)}{n}
\end{aligned}$$

Similarly for B,

$$\mathbb{E}[B] \leq \Delta\, \alpha \sum_{\ell=1}^{n} \frac{\mathbb{P}(\sigma_\ell^{p_1} = 0, \sigma_\ell^{p_2} = 1)}{n}.$$

Summing $\mathbb{E}[A] + \mathbb{E}[B]$, since $(\sigma^{p_1}, \sigma^{p_2}) \sim \pi_{opt}^{p}$,

$$\begin{aligned}\mathbb{E}[A+B] &\leq \alpha\Delta \sum_{\ell=1}^{n} \frac{\mathbb{P}(\sigma_\ell^{p_1} = 1, \sigma_\ell^{p_2} = 0) + \mathbb{P}(\sigma_\ell^{p_1} = 0, \sigma_\ell^{p_2} = 1)}{n} \\ &= \frac{1}{n}\alpha\Delta W_1(p_1, p_2).\end{aligned}$$

Estimation of term C. Observe that the condition $N(\ell) \cap D_{\sigma^{p_1},\sigma^{p_2}} \neq \emptyset$ implies that $\sigma_\ell^{p_1} = \sigma_\ell^{p_2} = 1$ cannot occur. Therefore,

$$\mathbb{E}[C] = \frac{1}{n^2} \sum_{\ell,k=1} \mathbb{P}\Big(\sigma_k^{q_1} = \sigma_k^{q_2} = 1,\ N(\ell) \cap D_{\sigma^{p_1},\sigma^{p_2}} \neq \emptyset,\ 0 \in \{\sigma_\ell^{p_1}, \sigma_\ell^{p_2}\}\Big) \tag{2.25}$$

$$\begin{aligned}&= \frac{1}{n^2} \sum_{\ell,k=1} \mathbb{P}\Big(\sigma_k^{q_1} = \sigma_k^{q_2} = 1 \,\Big|\, \begin{matrix} N(\ell) \cap D_{\sigma^{p_1},\sigma^{p_2}} \neq \emptyset, \\ 0 \in \{\sigma_\ell^{p_1}, \sigma_\ell^{p_2}\}\end{matrix}\Big) \\ &\quad \times \mathbb{P}\Big(N(\ell) \cap D_{\sigma^{p_1},\sigma^{p_2}} \neq \emptyset,\ 0 \in \{\sigma_\ell^{p_1}, \sigma_\ell^{p_2}\}\Big)\end{aligned}$$

Letting $\mathbb{Q} = \mathbb{P}(\cdot \mid N(\ell) \cap D_{\sigma^{p_1},\sigma^{p_2}} \neq \emptyset,\ 0 \in \{\sigma_\ell^{p_1}, \sigma_\ell^{p_2}\})$,

$$\begin{aligned}\frac{1}{n}\sum_k \mathbb{Q}(\sigma_k^{q_1} = \sigma_k^{q_2} = 1) = \mathbf{1}(\sigma_\ell^{p_1} = 1, \sigma_\ell^{p_2} = 0)\frac{1}{n}\sum_k \mathbb{P}_{\sigma\sim q_1}(\sigma_k = 1) + \\ (\mathbf{1}(\sigma_k^{p_1} = 0,\ \sigma_\ell^{p_2} = 1) + \mathbf{1}(\sigma_\ell^{p_1} = \sigma_\ell^{p_2} = 0))\frac{1}{n}\sum_k \mathbb{P}_{\sigma\sim q_2}(\sigma_{q_2} = 1) = \alpha.\end{aligned} \tag{2.26}$$

Moreover,

$$\begin{aligned}\frac{1}{n}\sum_{\ell=1}^{n} \mathbb{P}(N(\ell) \cap D_{\sigma^{p_1},\sigma^{p_2}} \neq \emptyset,\ 0 \in \{\sigma_\ell^{p_1}, \sigma_\ell^{p_2}\}) &\leq \frac{1}{n}\sum_{\ell=1}^{n} \mathbb{P}(N(\ell) \cap D_{\sigma^{p_1},\sigma^{p_2}} \neq \emptyset) \\ &\leq \frac{1}{n}\sum_{\ell=1}^{n} \mathbb{P}(\ell \in \bigcup_{u\in D_{\sigma^{p_1},\sigma^{p_2}}} N(u)) \\ &\leq \frac{1}{n}\sum_{\ell=1}^{n}\sum_{u=1}^{n} \mathbb{P}(u \in D_{\sigma^{p_1},\sigma^{p_2}}, \ell \in N(u)) \\ &\leq \frac{\Delta}{n} W_1(p_1, p_2).\end{aligned} \tag{2.27}$$

Putting together (2.25),(2.26) and (2.27) we obtain

$$\mathbb{E}[C] \leq \frac{\alpha\Delta}{n} W_1(p_1, p_2).$$

Estimation of term D

$$\begin{aligned}
\mathbb{E}[D] &= \frac{1}{n}\sum_{\ell=1}^{n}\mathbb{P}(\sigma_\ell^{p_1}=\sigma_\ell^{p_2}=1)\left[\frac{1}{n}\sum_{k=1}^{n}\mathbb{P}(N^+(k)\cap D_{\sigma^{q_1},\sigma^{q_2}}\neq\emptyset|\sigma_\ell^{p_1}=\sigma_\ell^{p_2}=1)\right]\\
&\le \frac{1}{n}\sum_{\ell=1}^{n}\mathbb{P}(\sigma_\ell^{p_1}=1)\left[\frac{1}{n}\sum_{k=1}^{n}\sum_{u=1}^{n}\mathbb{P}(k\in N^+(u),\, u\in D_{\sigma^{q_1},\sigma^{q_2}}\mid \sigma_\ell^{p_1}=\sigma_\ell^{p_2}=1)\right]\\
&\le \frac{1}{n}\sum_{\ell=1}^{n}\mathbb{P}(\sigma_\ell^{p_1}=1)\left[\frac{\Delta+1}{n}W_1(q_1,q_2)\right]\\
&= \frac{(\Delta+1)\alpha}{n}W_1(q_1,q_2),
\end{aligned}$$

where we use that $\mathrm{Law}(\sigma^{q_1},\sigma^{q_2})|\sigma_\ell^{p_1}=\sigma_\ell^{p_2}=1\sim\pi_{opt}^{q}$

Conclusion. Observing that $\mathbb{E}[|D_{\sigma^{p_1},\sigma^{p_2}}|]=W_1(p_1,p_2)$ and $\mathbb{E}[\mathbf{1}(\ell\in D_{\sigma^{p_1},\sigma^{p_2}})]=\frac{1}{n}W_1(p_1,p_2)$, by previous estimations we have

$$\begin{aligned}
W_1(\widehat{T}(p_1,q_1),\widehat{T}(p_2,q_2)) &\le \mathbb{E}|D_{\tau^1,\tau^2}|\\
&\text{ineq. (2.24)} \le \mathbb{E}[|D_{\sigma^{p_1},\sigma^{p_2}}|]-\mathbb{E}[\mathbf{1}(\ell\in D_{\sigma^{p_1},\sigma^{p_2}})]+\mathbb{E}[A]+\mathbb{E}[B]+\mathbb{E}[C]+\mathbb{E}[D]\\
&\le \Big[1-\frac{1-2(\Delta+1)\alpha}{n}\Big]W_1(p_1,p_2)+\frac{\alpha(\Delta+1)}{n}W_1(q_1,q_2).
\end{aligned}$$

which concludes the proof of (2.23).

2.4 Proof of Theorem 2.6

Here we prove Theorem 2.6. Recall that the state space $\Omega_{N,\alpha}$, with $\alpha\in D_N$, can be identified with the collection of independent sets of size $Nn\alpha$ on the graph $\mathcal{G}_N$, which consists of N disjoint copies of G; under this identification, each $\sigma(i)$ is an independent set of the i-th copy of G.

Given $\sigma,\sigma'\in\Omega_{N,\alpha}$ we write $\sigma\sim\sigma'$ if $\sigma\neq\sigma'$ and there exist $(i,\ell),(j,k)$ with $i\neq j$ such that $(\ell,k)_{\sigma(i),\sigma(j)}$ is satisfied and $\sigma'=\sigma^{i,\ell;j,k}$. Define the Markov chain with state space $\Omega_{N,\alpha}$ and transition matrix $P_N(\sigma,\sigma')$ such that whenever $\sigma\neq\sigma'$,

$$P_N(\sigma,\sigma'):=\frac{1}{(Nn)^2}\,\mathbf{1}(\sigma\sim\sigma'). \tag{2.28}$$

In words, at each step we select uniformly at random two pairs (i,ℓ) and (j,k) in $[N]\times[n]$, and if $i\neq j$ and the exchange is admissible, we update $\sigma\to\sigma^{i,\ell;j,k}$. By symmetry, $\mu_{N,\alpha}$ is reversible for P_N for every $\alpha\in D_N$, and the associated Dirichlet form is

$$\bar{\mathcal{E}}_{N,\alpha}(F,G):=\mu_{N,\alpha}[F\cdot(I-P_N)G]=\frac{1}{2(Nn)^2}\sum_{\substack{\sigma,\sigma'\in\Omega_{N,\alpha}:\\ \sigma\sim\sigma'}}\mu_{N,\alpha}(\sigma)(F(\sigma)-F(\sigma'))(G(\sigma)-G(\sigma')).$$

Theorem 2.17 (Log-Sobolev inequality (LSI))**.** *Suppose that G has maximum degree $\Delta\ge 3$ and n vertices, and let $N\ge n^2$. For any $\gamma,\delta>0$, if $\alpha\in D_N$ satisfies $\alpha\in(\gamma,(1-\delta)\alpha_c(\Delta))$,*

then the following Log-Sobolev inequality holds:

$$\bar{\mathcal{E}}_{N,\alpha}(\sqrt{F},\sqrt{F}) \;\geq\; r\,\mathrm{Ent}_{\mu_{N,\alpha}}[F], \qquad \forall\, F:\Omega_{N,\alpha}\to\mathbb{R}_+, \tag{2.29}$$

with $r \gtrsim_{\gamma,\delta,\Delta} \frac{1}{Nn}$.

Here, the notation $\lesssim_{\gamma,\delta,\Delta}$ and $\gtrsim_{\gamma,\delta,\Delta}$ denotes inequalities up to multiplicative constants depending only on γ,δ, and Δ. In general, the Log-Sobolev constant is defined as the largest positive constant r for which (2.29) holds; see, for instance, the seminal work of [DS96] and the monograph [MT06].

Theorem 2.6 follows immediately from Theorem 2.17. Indeed, one first observes that $\mathcal{E}_{N,\alpha}(\sqrt{F},\sqrt{F}) \geq N\,\bar{\mathcal{E}}_{N,\alpha}(\sqrt{F},\sqrt{F})$, and then converts the LSI (2.29) into the MLSI (2.8) using the elementary inequality

$$(x-y)\log\Big(\frac{x}{y}\Big) \geq 4(\sqrt{x}-\sqrt{y})^2, \qquad x,y\geq 0,$$

which implies $\mathcal{E}_{N,\alpha}(F,\log F) \geq 4\,\mathcal{E}_{N,\alpha}(\sqrt{F},\sqrt{F})$. Combining the previous inequalities with Theorem 2.17, we obtain

$$\begin{aligned}\mathcal{E}_{N,\alpha}(F,\log F) &\geq 4\,\mathcal{E}_{N,\alpha}(\sqrt{F},\sqrt{F})\\ &\geq 4N\,\bar{\mathcal{E}}_{N,\alpha}(\sqrt{F},\sqrt{F})\\ &\gtrsim_{\gamma,\delta,\Delta} \frac{1}{n}\,\mathrm{Ent}_{\mu_{N,\alpha}}[F].\end{aligned}$$

The proof of Theorem 2.17 is based on two reductions. The first reduces the problem to the density regime $\alpha\in(\gamma, C/\Delta^2)$ for a universal constant C, via a stochastic localisation argument whose key ingredient is an entropy contraction property of the localisation scheme; see Section 2.4.1, which is based on the recent paper [JPV26]. The second applies the martingale method, introduced by H.T. Yau [Yau96] and developed by many authors afterwards. This iteratively reduces the number of components of $\mathcal{G}_N$, combining the analysis of the Log- Sobolev constant of the "induced chain" on each component with the highly nontrivial analysis introduced in [Jai+23] to handle the low density regime. In contrast with the analysis in [Jai+23], here we need to keep track of two growing parameters, namely N and n.

The extension to the full subcritical regime is carried out in Section 2.4.1.

2.4.1 Reduction to Small Densities

In this section, we show that it is sufficient to establish the logarithmic Sobolev inequality for P_N only in the small density regime $\alpha\in(\gamma,\frac{1}{1000\Delta^2})$ in order to recover the full subcritical range $\alpha\in(\gamma,(1-\delta)\alpha_c(\Delta))$. To carry out this density reduction, it is convenient to transfer the problem to a different dynamics, the down-up walk, for which the reduction argument takes a cleaner form. We therefore begin by establishing the equivalence between the

logarithmic Sobolev inequality for P_N and its counterpart for the down-up walk, both in the full regime and in the reduced one; these are Theorem 2.18 and Theorem 2.19 below. We then show that Theorem 2.19 implies Theorem 2.18 via a stochastic localization argument, whose two key ingredients are a entropy contraction property of the localization scheme and a convexity property of the Dirichlet form.

We represent independent sets as subsets $X \subseteq V$ rather than spin configurations, adopting the correspondence $\sigma \to X$ to avoid confusion between the two representations.

The down-up walk on $\Omega_{N,\alpha}$. It is convenient to transfer the proof to the down-up walk, in the specific formulation of [CLV21], which can be rephrased as Glauber dynamics up to a relabelling of independent sets as spin configurations, where the i-th spin encodes the position of the i-th element of the independent set. Given a current state X, a new state is generated by selecting a vertex $u \in X$ uniformly at random, and then choosing uniformly at random among all independent sets of the same size as X that contain $X \setminus \{u\}$. This down-up walk preserves the uniform distribution on $\Omega_{N,\alpha}$. Moreover, its transition probabilities are easily compared to those of the chain P_N defined in (2.28): both chains admit the same transitions, and for densities $\alpha \in (\gamma, (1-\delta)\alpha_c(\Delta))$ the respective transition probabilities differ by at most a multiplicative factor depending only on γ, δ, and Δ. Consequently, their logarithmic Sobolev constants are of the same order, and Theorem 2.17 is equivalent to the following.

Theorem 2.18. *Suppose $N \geq n^2$ and fix $\alpha \in (\gamma, (1-\delta)\alpha_c(\Delta))$. Then the down-up walk on independent sets of $\mathcal{G}_N$ of size $\alpha N n$ has logarithmic Sobolev constant $r = \Omega_{\gamma,\delta,\Delta}(1/Nn)$.*

The reduction statement consists in the sufficiency of the following theorem.

Theorem 2.19. *Let $\mathcal{G}$ be a disjoint union of $\Theta(N)$ graphs, each with at most n vertices and maximum degree Δ. The down-up walk on independent sets of $\mathcal{G}$ with density $\alpha \in (\gamma, \frac{1}{1000\Delta^2})$ has logarithmic Sobolev constant $r = \Omega_{\gamma,\Delta}(1/Nn)$.*

The logical structure is therefore: Theorem 2.19 $\Rightarrow$ Theorem 2.18 $\Leftrightarrow$ Theorem 2.17. We note that Theorem 2.19 allows the components of $\mathcal{G}$ to be distinct graphs, whereas Theorem 2.18 is stated for N identical copies of G. As it will be seen in a moment, the more general framework with possibly distinct components is necessary for the proof of the implication Theorem 2.19 $\Rightarrow$ Theorem 2.18.

Stochastic localization and density reduction. We sample $X \in \Omega_{N,\alpha}$, viewed as an independent set of $\mathcal{G}_N$, and draw a uniformly random ordering $(X_1, \dots, X_{Nn\alpha})$ of the elements of X. For each $t \in \{0, 1, \dots, Nn\alpha\}$, define the *pinned set* $Y_t := \{X_1, \dots, X_t\}$ and the *residual graph*

$$\mathcal{G}_{N,t} = \mathcal{G}_N[V \setminus (Y_t \cup N(Y_t))],$$

where $N(Y_t)$ denotes the neighborhood of Y_t in $\mathcal{G}_N$. Define the random measure

$$\mu_{N,\alpha,t} = \mathcal{P}(X \setminus Y_t \mid Y_t \subseteq X).$$

The random sequence $(\mu_{N,\alpha,t})_{t\in\{0,\dots,Nn\alpha\}}$ is a martingale and constitutes, in the framework of [CE22], a *localization scheme* for $\mu_{N,\alpha}$. One readily verifies that $\mu_{N,\alpha,t}$ is the uniform measure on independent sets of size $Nn\alpha - t$ in $\mathcal{G}_{N,t}$; we denote this state space by $\Omega_{N,\alpha,t}$, its vertex set by V_t, and the residual density by $\alpha_t = \frac{Nn\alpha - t}{|V_t|}$. The residual graph $\mathcal{G}_{N,t}$ is a disjoint union of $N(t)$ induced subgraphs $G_1^t, \dots, G_{N(t)}^t$, each with at most n vertices and maximum degree at most Δ.

Proposition 2.20 (Reduction to small densities)**.** *Let G be a graph with n vertices and maximum degree $\Delta \geq 3$, and let $N \geq n^2$. Fix $\alpha \in (\gamma, \alpha_c(\Delta))$ and define*

$$c_\Delta = \frac{1}{\Delta} \cdot \frac{0.26\Delta - 0.74}{1001\Delta + 1}.$$

Set $t^* = \lfloor Nn\alpha(1 - c_\Delta) \rfloor + 1$. Then

$$\alpha_{t^*} \leq \frac{1}{1000\Delta^2} \qquad \textit{and} \qquad N(t^*) \geq \frac{N}{2}.$$

Proof. Since each pinned vertex excludes at most Δ neighbors, the residual vertex set satisfies

$$|V_{t^*}| \geq Nn - t^*(\Delta + 1) \geq Nn\,[1 - \alpha(1 - c_\Delta)(\Delta + 1)] - (\Delta + 1),$$

which for Nn large gives $|V_{t^*}| \geq Nn\left[1 - \alpha(1 - c_\Delta)(\Delta + 1) - \frac{1}{Nn}\right]$.

For the density bound, using $Nn\alpha - t^* \leq Nn\alpha c_\Delta$,

$$\alpha_{t^*} = \frac{Nn\alpha - t^*}{|V_{t^*}|} \leq \frac{\alpha\, c_\Delta}{1 - \alpha(1 - c_\Delta)(\Delta + 1) - \frac{1}{Nn}} \leq \frac{1}{1000\Delta^2},$$

where the last inequality holds for all $\Delta \geq 3$ and $\alpha < \alpha_c(\Delta) \leq \frac{1}{e\Delta}$.

For the components bound, since each component has at most n vertices, $N(t^*) \cdot n \geq |V_{t^*}|$, and therefore

$$N(t^*) \geq \frac{|V_{t^*}|}{n} \geq N\left[1 - \alpha(1 - c_\Delta)(1 + \Delta) - \frac{1}{Nn}\right] \geq \frac{N}{2},$$

where the last inequality follows from $\alpha(1 - c_\Delta)(1 + \Delta) \leq \frac{1}{e} < \frac{1}{2}$ for all $\Delta \geq 3$. □

Proposition 2.20 guarantees that at time t^* the residual system falls within the small density regime of Theorem 2.19, while retaining at least $N/2$ components. It therefore remains to transfer the logarithmic Sobolev inequality from the reduced system back to the original measure $\mu_{N,\alpha}$. This is achieved via the following entropy contraction inequality, which appears as [JPV26, ineq. (6.3)] and is derived therein from the entropic independence property established in [JPV26, Proposition 6.5]: for $t = \Omega_\Delta(Nn\alpha)$,

$$\mathrm{Ent}_{\mu_{N,\alpha}}[F] \lesssim_{\gamma,\delta,\Delta} \mathbb{E}\Big[\mathrm{Ent}_{\mu_{N,\alpha,t}}[F]\Big]. \tag{2.30}$$

Applying (2.30) with $t = t^*$ and Theorem 2.19, we conclude: denoting by $\mathcal{E}_{N,\alpha,t}$ the

Dirichlet form of the down-up walk on $\Omega_{N,\alpha,t}$, for all $F\colon \Omega_{N,\alpha} \to \mathbb{R}_+$ we have

$$\begin{aligned}\mathrm{Ent}_{\mu_{N,\alpha}}[F] &\lesssim_{\gamma,\delta,\Delta} \mathbb{E}\Big[\mathrm{Ent}_{\mu_{N,\alpha,t^*}}[F]\Big] \\ &\lesssim_{\gamma,\delta,\Delta} Nn\,\mathbb{E}\Big[\mathcal{E}_{N,\alpha,t^*}(\sqrt{F},\sqrt{F})\Big] \\ &\le Nn\,\mathcal{E}_{N,\alpha,0}(\sqrt{F},\sqrt{F}),\end{aligned}$$

where the last inequality uses the convexity of the Dirichlet form along the localization scheme, established in [CE22, Proposition 2.21] for Glauber dynamics, which applies here via the relabelling of the down-up walk described above.

2.4.2 Martingale's Method in the Reduced Regime

For the sake of simplicity, we shall prove Theorem 2.19 in the simplified setting where one has N identical copies of G. However, a careful check of the proof shows that the general setting of Theorem 2.19 can be handled with very minor modifications.

Thus, we now focus on proving the following statement.

Theorem 2.21. *In the setting of Section 2.4, suppose G has maximum degree $\Delta \ge 3$ and n vertices, and let $N \ge n^2$. There exists an absolute constant $C \le 1/1000$, such that for any $\gamma > 0$, if $\alpha \in D_N$ satisfies $\alpha \in (\gamma, C/\Delta^2)$ for then for all $F : \Omega_{N,\alpha} \to \mathbb{R}_+$,*

$$\frac{1}{Nn}\,\mathrm{Ent}_{\mu_{N,\alpha}}[F] \lesssim_{\gamma,\Delta} \bar{\mathcal{E}}_{N,\alpha}(\sqrt{F},\sqrt{F}).$$

The proof of Theorem 2.21 is based on the martingale method, which reduces the logarithmic Sobolev inequality for a system of N components to the same inequality for systems of $N-1$ components, at the cost of a small additive error. The key quantity is the optimal constant $t_{N,\alpha}$ in the logarithmic Sobolev inequality, defined as follows.

Definition 2.22. For a given pair (N,α) with $\alpha \in D_N$, we define $t_{N,\alpha}$ as the smallest constant such that

$$\frac{1}{Nn}\,\mathrm{Ent}_{N,\alpha}[F] \le t_{N,\alpha}\,\bar{\mathcal{E}}_{N,\alpha}(\sqrt{F},\sqrt{F})$$

holds for all $F : \Omega_{N,\alpha} \to \mathbb{R}_+$, where

$$\bar{\mathcal{E}}_{N,\alpha}(\sqrt{F},\sqrt{F}) = \frac{1}{2(Nn)^2}\sum_{i,\ell}\sum_{j\neq i,k}\mu_{N,\alpha}\Big[\mathbf{1}(\ell,k)_{\sigma(i),\sigma(j)}\left(\nabla^{i,\ell;j,k}\sqrt{F}\right)^2\Big].$$

Theorem 2.19 is then equivalent to showing that $t_{N,\alpha} \lesssim_{\gamma,\Delta} 1$ for all $\alpha \in (\gamma, \frac{1}{1000\Delta^2})$ and $N \ge n^2$. The core of the argument is the following recursive inequality.

Lemma 2.23 (Recursive rule for $t_{N,\alpha}$)**.** *Suppose $\alpha \in D_N$ satisfies $\alpha \in (\gamma/2, 1/(241\Delta^2))$ for some $\gamma > 0$. Then*

$$t_{N,\alpha} \le \frac{N-2}{N-1}\max_{\sigma\sim\sigma'}\left\{\sum_{i\in[N]}\frac{\mathbf{1}(\sigma(i)=\sigma'(i))}{N-2}\,t_{N-1,\alpha^{\sigma(i)}}\right\} + \frac{C_{\Delta,\gamma}}{N}, \tag{2.31}$$

with $C_{\Delta,\gamma}$ a positive constant and $\alpha^\tau = \frac{N}{N-1}\alpha - \frac{|\tau|}{(N-1)n}$ for $\tau \in \Omega_G$.

Proof. The proof combines the elementary entropy decomposition

$$\mathrm{Ent}_{N,\alpha}[F] = \frac{1}{N}\sum_{i\in[N]} \mu_{N,\alpha}[\mathrm{Ent}_{N,\alpha}[F \mid \sigma(i)]] + \frac{1}{N}\sum_{i\in[N]} \mathrm{Ent}_{N,\alpha}[\mu_{N,\alpha}[F \mid \sigma(i)]]\,,$$

with the two lemmas below. Lemma 2.24 controls the first term and yields the main recursive contribution, while Lemma 2.25 controls the second term and produces the additive error $C_{\Delta,\gamma}/N$. □

Lemma 2.24. *For all $\alpha \in D_N$, we have*

$$\begin{aligned}&\frac{1}{N}\sum_{i\in[N]} \mathbb{E}_{\mu_{N,\alpha}}\left[\mathrm{Ent}_{\mu_{N,\alpha}}[F \mid \sigma(i)]\right] \\ &\qquad \le \frac{Nn(N-2)}{(N-1)} \max_{\sigma\sim\sigma'} \left\{ \sum_{i\in[N]} \frac{\mathbf{1}(\sigma(i)=\sigma'(i))}{N-2}\, t_{N-1,\alpha^{\sigma(i)}} \right\} \bar{\mathcal{E}}_{N,\alpha}(\sqrt{F},\sqrt{F}),\end{aligned}$$

for all $F : \Omega_{N,\alpha} \to \mathbb{R}_+$, with $\alpha^\tau = \frac{N}{N-1}\alpha - \frac{1}{(N-1)n}|\tau|$ for $\tau \in \Omega_G$.

Proof. Note that by definition, for all $i \in [N]$,

$$\mathrm{Ent}_{\mu_{N,\alpha}}[F \mid \sigma(i)=\tau] \le \frac{t_{N-1,\alpha^\tau}}{2(N-1)n} \sum_{\substack{\sigma\sim\sigma'\in\Omega_{N,\alpha}\\ \sigma(i)=\sigma'(i)=\tau}} \mu_{N,\alpha}(\sigma \mid \sigma(i)=\tau)\left(\sqrt{F(\sigma)}-\sqrt{F(\sigma')}\right)^2.$$

Therefore,

$$\begin{aligned}\frac{1}{N}\sum_{i\in[N]} \mu_{N,\alpha}[\mathrm{Ent}_{N,\alpha}[F \mid \sigma(i)]] &\le \frac{1}{N}\sum_{i\in[N]}\sum_{\tau\in\Omega_G} \mu_{N,\alpha,i}(\tau)\,\frac{t_{N-1,\alpha^\tau}}{2(N-1)n} \\ &\quad \times \sum_{\substack{\sigma\sim\sigma'\in\Omega_{N,\alpha}\\ \sigma(i)=\sigma'(i)=\tau}} \mu_{N,\alpha}(\sigma \mid \sigma(i)=\tau)\left(\sqrt{F(\sigma)}-\sqrt{F(\sigma')}\right)^2 \\ &= \frac{1}{(N-1)}\frac{1}{Nn} \sum_{\sigma\sim\sigma'\in\Omega_{N,\alpha}} \mu_{N,\alpha}(\sigma)\,\left(\sqrt{F(\sigma)}-\sqrt{F(\sigma')}\right)^2 \\ &\quad \times \sum_{i\in[N]} \mathbf{1}\{\sigma(i)=\sigma'(i)\}\, t_{N-1,\alpha^{\sigma(i)}} \\ &\le \frac{Nn(N-2)}{(N-1)} \max_{\sigma\sim\sigma'} \left\{ \sum_{i\in[N]} \frac{\mathbf{1}(\sigma(i)=\sigma'(i))}{N-2}\, t_{N-1,\alpha^{\sigma(i)}} \right\} \bar{\mathcal{E}}_{N,\alpha}(\sqrt{F},\sqrt{F}),\end{aligned}$$

which concludes the proof. □

Lemma 2.25. *Under the same hypotheses as Lemma 2.23,*

$$\frac{1}{N}\sum_{i\in[N]} \mathrm{Ent}_{\mu_{N,\alpha}}[\mathbb{E}[F \mid \sigma(i)]] \lesssim_{\gamma,\Delta} n\,\bar{\mathcal{E}}_{N,\alpha}(\sqrt{F},\sqrt{F}).$$

The proof of Lemma 2.25 is more involved and is deferred to Section 2.4.5, as it requires some additional tools originally introduced in [Jai+23]. The proof of Theorem 2.21 is then completed in Section 2.4.3, where the recursive rule is iterated via a probabilistic argument.

2.4.3 Proof of Theorem 2.21

The statement of Theorem 2.21 is equivalent to showing that for all $\alpha \in (\gamma, \frac{1}{1000\Delta^2})$ and $N \geq n^2$,

$$t_{N,\alpha} \lesssim_{\gamma,\Delta} 1.$$

The proof relies on the recursive rule established in Lemma 2.23 and follows a probabilistic argument from [Jai+23]. We begin by reformulating Lemma 2.23 in probabilistic terms and defining the density process that will be iterated.

Probabilistic reformulation and density process. Consider a pair $\sigma, \sigma' \in \Omega_{N,\alpha}$, $\sigma \sim \sigma'$, achieving the maximum in (2.31), and sample uniformly at random a component $\tau^1 = \sigma(i^*)$, where $i^* \in \{1, \dots, N\}$ is such that $\sigma(i^*) = \sigma'(i^*)$. Then

$$t_{N,\alpha} \leq \frac{N-2}{N-1}\,\mathbb{E}[t_{N-1,\alpha_1}] + \frac{C_{\Delta,\gamma}}{N}, \tag{2.32}$$

where $\alpha_1 = \alpha^{\tau^1} = \frac{N}{N-1}\alpha - \frac{|\tau^1|}{(N-1)n}$.

More generally, we define iteratively a process $(\alpha_s)_{s\in\mathbb{N}}$ on the densities, where each α_s is a density on the state space Ω_{N-s,α_s}, according to the following recursion:

1. $\alpha_0 = \alpha$;
2. given α_s at time $s \geq 0$: consider a pair $\sigma, \sigma' \in \Omega_{N-s,\alpha_s}$ maximizing (2.31) when N is replaced by $N - s$, sample uniformly at random $i^*_{s+1} \in \{1, \dots, N-s\}$ such that $\sigma(i^*_{s+1}) = \sigma'(i^*_{s+1})$, set $\tau^{s+1} = \sigma(i^*_{s+1})$, and define

$$\alpha_{s+1} = \alpha_s^{\tau^{s+1}} = \frac{N-s}{N-s-1}\alpha_s - \frac{|\tau^{s+1}|}{(N-s-1)n}. \tag{2.33}$$

We introduce the stopping time

$$S := \min\left\{ s > 0 : \alpha_s \notin \left(\gamma/2, \frac{1}{241\Delta^2}\right)\right\}, \tag{2.34}$$

and note that on the event $\{s \leq S\}$ the recursive estimate (2.32) can be applied at every step. Note in particular that $\alpha_0 \in (\gamma, \frac{1}{1000\Delta^2})$, which is strictly contained in the interval $(\gamma/2, \frac{1}{241\Delta^2})$ defining S.

Iteration of the recursive rule. The following proposition quantifies the effect of iterating (2.32) up to time s.

Proposition 2.26. *Let $s \le N-2$. Then*

$$\mathbb{E}[t_{N,\alpha} \mid s \le S] \le \frac{N-s-1}{N-1}\,\mathbb{E}[t_{N-s,\alpha_s} \mid s \le S] + \sum_{i=1}^{s} \frac{N-i}{N-1}\,\frac{C_{\Delta,\gamma}}{N-i+1}.$$

Proof. On the event $\{s \le S\}$, Lemma 2.23 guarantees that the recursive inequality can be applied at every step up to time s. Define

$$a_i = \frac{N-i-1}{N-1}, \qquad b_i = \frac{C_{\Delta,\gamma}}{N-i+1}.$$

The claim is equivalent to

$$\mathbb{E}[t_{N,\alpha} \mid s \le S] \le a_s\,\mathbb{E}[t_{N-s,\alpha_s} \mid s \le S] + \sum_{i=1}^{s} a_{i-1} b_i.$$

For $s = 1$ the statement follows from (2.32). Assuming the claim holds for some $i < s$, we obtain

$$\begin{aligned}
\mathbb{E}[t_{N,\alpha} \mid s \le S] &\le a_i\,\mathbb{E}[t_{N-i,\alpha_i} \mid s \le S] + \sum_{j=1}^{i} a_{j-1} b_j \\
&\le a_i \left[\frac{N-i-2}{N-i-1}\,\mathbb{E}[t_{N-i-1,\alpha_{i+1}} \mid s \le S] + b_{i+1}\right] + \sum_{j=1}^{i} a_{j-1} b_j \\
&= a_{i+1}\,\mathbb{E}[t_{N-i-1,\alpha_{i+1}} \mid s \le S] + \sum_{j=1}^{i+1} a_{j-1} b_j,
\end{aligned}$$

which concludes the induction. □

Concentration of the density process. The next proposition ensures that the stopping time S is large with high probability, so that the iteration in Proposition 2.26 can be carried out up to time $s = N - 16\log^2(N)$.

Proposition 2.27. *Let h be an integer with $h^2 \le N-2$ and $h \ge (6 \cdot 10^3 \Delta^2)^2$. Then*

$$\mathbb{P}(S \le N - h^2) \le \exp\left(-\frac{h}{2}\right).$$

Proof. Note that the claim is equivalent to

$$\mathbb{P}(S \le N-h^2) \le \mathbb{P}\left(\max_{i \le N-h^2} |\alpha_i - \alpha_0| \ge \gamma/2\right) \le \exp\left(-\frac{h}{2}\right), \tag{2.35}$$

since $\alpha_0 \in (\gamma, \frac{1}{1000\Delta^2})$ and, by definition of S in (2.34), the event $\{S \le N-h^2\}$ implies $\max_{i \le N-h^2} |\alpha_i - \alpha_0| \ge \gamma/2$. We establish the stronger bound

$$\mathbb{P}\left(\max_{i \le N-h} |\alpha_i - \alpha_0| \ge \frac{3}{\sqrt{h}}\right) \le \exp\left(-\frac{h}{2}\right), \tag{2.36}$$

which yields (2.35) provided $\sqrt{h} \geq 6/\gamma > 6000\Delta^2$.

Let $\mathcal{F}_s$ denote the σ-algebra generated by $\{\alpha_{s'}, \tau_{s'} : s' \leq s\}$, so that $\{\mathcal{F}_s\}_{s\geq 0}$ is a filtration. We introduce the auxiliary martingale $\{\bar{\alpha}_s\}_{s\geq 0}$ defined by

$$\bar{\alpha}_{s+1} = \bar{\alpha}_s + (\alpha_{s+1} - \alpha_s) - \mathbb{E}[\alpha_{s+1} - \alpha_s \mid \mathcal{F}_s],$$

and establish the following preliminary bounds. From (2.33),

$$\alpha_{s+1} - \alpha_s = \frac{\alpha_s}{N-s-1} - \frac{|\tau_s|}{(N-s-1)n}. \tag{2.37}$$

Using (2.37) and averaging over the uniform choice of i^*_{s+1},

$$\begin{aligned}
\mathbb{E}[\alpha_{s+1} - \alpha_s \mid \mathcal{F}_s] &= \frac{1}{N-s-2} \sum_{i:\,\sigma(i)=\sigma'(i)} \left(\frac{\alpha_s}{N-s-1} - \frac{|\sigma(i)|}{(N-s-1)n} \right) \qquad (2.38)\\
&= \frac{1}{(N-s-1)(N-s-2)} \left(\left(\frac{N-s-2}{N-s} - 1 \right) \sum_{i\in[N-s]} \frac{|\sigma(i)|}{n} + \sum_{i:\,\sigma(i)\neq\sigma'(i)} \frac{|\sigma(i)|}{n} \right)\\
&\leq \frac{1}{(N-s-2)(N-s-1)} \sum_{i:\,\sigma(i)\neq\sigma'(i)} \frac{|\sigma(i)|}{n} \leq \frac{2}{(N-s-2)(N-s-1)}.
\end{aligned}$$

Combining (2.37) and (2.38),

$$|\bar{\alpha}_{s+1} - \bar{\alpha}_s| \leq \frac{2}{N-s-1}.$$

For the quadratic variation, we compute

$$\mathbb{E}\Big[(\alpha_{s+1} - \alpha_s)^2 \mid \mathcal{F}_s\Big] = \frac{1}{N-s-2} \sum_{i:\,\sigma(i)=\sigma'(i)} \frac{[\alpha_s - |\sigma(i)|/n]^2}{(N-s-1)^2} \leq \frac{1}{(N-s-1)^2},$$

and therefore, using the identity $\frac{1}{n}\sum_i a_i^2 - \frac{1}{n^2}\sum_{i,j} a_i a_j = \frac{1}{2n^2}\sum_{i,j}(a_i - a_j)^2$ with $|a_i - a_j| \leq 1$,

$$\mathbb{E}\Big[(\bar{\alpha}_{s+1} - \bar{\alpha}_s)^2 \mid \mathcal{F}_s\Big] = \mathbb{E}\Big[(\alpha_{s+1} - \alpha_s)^2 \mid \mathcal{F}_s\Big] - \mathbb{E}[\alpha_{s+1} - \alpha_s \mid \mathcal{F}_s]^2 \leq \frac{1}{2(N-s-1)^2}.$$

The quadratic variation $\langle \bar{\alpha} \rangle_s = \sum_{i=1}^{s} \mathbb{E}[(\bar{\alpha}_i - \bar{\alpha}_{i-1})^2 \mid \mathcal{F}_{i-1}]$ therefore satisfies

$$\langle \bar{\alpha} \rangle_s \leq \sum_{i=1}^{s} \frac{1}{2(N-i)^2} \leq \frac{1}{2} \int_{N-s-1}^{N-1} \frac{dx}{x^2} = \frac{1}{2} \frac{s}{(N-1)(N-s-1)}.$$

Applying Freedman's martingale inequality [Fre75],

$$\mathbb{P}\left(\max_{i\leq s} |\bar{\alpha}_i - \bar{\alpha}_0| \geq x \right) \leq \exp\left[-\frac{(N-s-1)x^2}{\frac{s/2}{N-1} + \frac{4}{3}x} \right].$$

To translate this into a bound for α_s, we write $\bar{\alpha}_i - \bar{\alpha}_0 = \alpha_i - \alpha_0 - \sum_{j\leq i} \mathbb{E}[\alpha_j - \alpha_{j-1} \mid \mathcal{F}_{j-1}]$

and define

$$\bar{\Sigma}_s := \max_{i\le s}\left|\sum_{j\le i}\mathbb{E}[\alpha_j-\alpha_{j-1}\mid\mathcal{F}_{j-1}]\right|,$$

so that $|\bar{\alpha}_i-\bar{\alpha}_0|\ge|\alpha_i-\alpha_0|-\bar{\Sigma}_s$, and therefore

$$\mathbb{P}\Big(\max_{i\le s}|\alpha_i-\alpha_0|\ge x+\bar{\Sigma}_s\Big)\le\mathbb{P}\Big(\max_{i\le s}|\bar{\alpha}_i-\bar{\alpha}_0|\ge x\Big).$$

By (2.38),

$$\bar{\Sigma}_s\le 2\sum_{j=N-s-1}^{N-2}\frac{1}{j(j-1)}=2\left(\frac{1}{N-s-2}-\frac{1}{N-2}\right)=\frac{2s}{(N-s-2)(N-2)}.$$

Taking $s=N-h^2$ gives $\bar{\Sigma}_s\le\frac{2}{h^2-2}$. Choosing $x=\frac{1}{\sqrt{h}}$,

$$\mathbb{P}\Big(\max_{i\le s}|\alpha_i-\alpha_0|\ge\frac{3}{\sqrt{h}}\Big)\le\mathbb{P}\Big(\max_{i\le s}|\alpha_i-\alpha_0|\ge\frac{1}{\sqrt{h}}+\bar{\Sigma}_s\Big)\le\exp(-h),$$

which proves (2.36). Indeed, for $x=\frac{1}{\sqrt{h}}$ and $s=N-h^2$,

$$\exp\left[-\frac{(N-s-1)x^2}{\frac{s/2}{N-1}+\frac{4}{3}x}\right]=\exp\left[-\frac{\frac{h^2-1}{h}}{\frac{1}{2}\frac{N-h^2}{N-1}+\frac{4}{3\sqrt{h}}}\right]\le\exp\Big(-\frac{h}{2}\Big),$$

since $h^2-1\ge\frac{1}{2}h^2$ and $\frac{1}{2}\frac{N-h^2}{N-1}+\frac{4}{3\sqrt{h}}\le 1$ for $h\ge 9$. □

Conclusion of the proof of Theorem 2.21. We begin with the non-optimal bound

$$t_{N,\alpha}\lesssim_{\gamma,\Delta}Nn, \tag{2.39}$$

which follows from a spectral gap estimate of order $1/(Nn)$ for P_N when $\alpha\in(0,1/(241\Delta^2))$, combined with standard inequalities comparing spectral gap and logarithmic Sobolev constants [DS96]. The spectral gap bound can be obtained through path coupling; the coupling argument is the same as in Proposition A.1 below. From now on, we assume the validity of (2.39).

We now apply Proposition 2.26 up to time $s=N-16\log^2(N)$. By Proposition 2.27 with $h=4\log(N)$,

$$\mathbb{P}\Big(S\le N-16\log^2(N)\Big)\le\frac{1}{N^2}.$$

On the complementary event $\{S\ge N-16\log^2(N)\}$, the density $\alpha_{N-16\log^2(N)}$ lies in $(\gamma/2,1/(241\Delta^2))$, so by (2.39),

$$t_{16\log^2(N),\,\alpha_{N-16\log^2(N)}}\lesssim_{\gamma,\Delta}16n\log^2(N).$$

Applying Proposition 2.26 and decomposing over the two events,

$$\begin{aligned} t_{N,\alpha} &\le \frac{16\log^4(N)}{N-1}\,\mathbb{E}\Big[t_{16\log^2(N),\,\alpha_{N-16\log^2(N)}}\,\Big|\,S \ge N-16\log^2(N)\Big] + \\ &\quad + C_{\Delta,\gamma} \;+\; t_{N,\alpha}\,\mathbb{P}\Big(S \le N-16\log^2(N)\Big) \\ &\le \frac{16n\log^6(N)}{N-1} + C_{\Delta,\gamma} + \frac{n}{N} \lesssim_{\gamma,\Delta} 1, \end{aligned}$$

provided that $N \ge n^2$. This concludes the proof of Theorem 2.21. □

It remains to prove Lemma 2.25. The argument proceeds in two steps. First, in Subsection 2.4.4, we establish a log-Sobolev inequality for the chain induced on a single particle $\sigma(i)$, with the remaining particles frozen. Second, in Subsection 2.4.5, we bound the Dirichlet form of this induced chain by the appropriate terms in $\bar{\mathcal{E}}_{N,\alpha}$, which yields the proof of Lemma 2.25.

2.4.4 Induced Chain

Consider a density $\alpha \in \mathcal{D}_N$. The *induced chain* is the Markov chain on Ω_G obtained from P_N by tracking only the first particle $\sigma(1) \in \Omega_G$. Formally, its transition matrix P is given, for $\tau \ne \tau' \in \Omega_G$, by

$$P(\tau,\tau') := \frac{1}{n}\sum_{\ell\in[n]}\frac{1}{Nn}\sum_{j\in[N]\setminus\{1\}}\sum_{k\in[n]}\mathbb{E}_{\sigma\sim\mu_{N,\alpha}}\Big[\mathbf{1}\{(\ell,k)_{\sigma(1),\sigma(j)}\}\,\mathbf{1}\{\sigma^{1,\ell;\,j,k}(1)=\tau'\}\,\Big|\,\sigma(1)=\tau\Big]. \tag{2.40}$$

This chain describes exchanges between the first particle and the remaining $N-1$ particles, the latter sampled at equilibrium. In the language of statistical mechanics, the induced chain is the analogue of a *grand-canonical* dynamics: the first particle plays the role of a small subsystem in contact with a reservoir, the other $N-1$ particles, with which it exchanges 1's while preserving the total density of the joint system. From the Markov chain perspective, this is a *heat-bath* dynamics: each update resamples $\sigma_\ell(1)$ according to its equilibrium conditional law given the reservoir. Strictly speaking, this is not the Glauber dynamics of the hardcore model $\mu_{G,\lambda}$, but a closely related chain whose invariant measure is the marginal $P_1\mu_{N,\alpha_N}$, defined below in (2.41). As we shall show in Theorem **??**, this marginal converges to $\mu_{G,\lambda}$ as $N\to\infty$, so the induced chain can be viewed as a finite-N approximation of the Glauber dynamics for the hardcore model.

This approximation can be made quantitative. Using the local central limit theorem of Proposition 2.28 below, we compare the induced chain P with the Glauber dynamics of $\mu_{G,\lambda}$ in a quantitative sense, and transfer to P the log-Sobolev inequality known to hold for Glauber throughout the subcritical regime by the spectral-independence theory of [CLV21]. The log-Sobolev inequality for P_N is then obtained from the one for P via the martingale method.

We collect below the ingredients needed to state the log-Sobolev inequality for P: its invariant measure, the structure of its transitions, and the associated Dirichlet form. By

construction, a transition $\tau \to \tau'$ of P is admissible only if τ' differs from τ by the addition or removal of a single occupied site. Accordingly, for $\ell \in [n]$ we set

$$\tau^{\ell} = \begin{cases} \tau \cup \{\ell\}, & \text{if } \tau \cup \{\ell\} \in \Omega_G, \\ \tau, & \text{otherwise,} \end{cases} \qquad \tau^{\bar{\ell}} = \tau \setminus \{\ell\}.$$

By symmetry, P is reversible with respect to

$$P_1\mu_{N,\alpha}(\tau) := \mathbb{P}_{\sigma\sim\mu_{N,\alpha}}(\sigma(1) = \tau), \qquad \tau \in \Omega_G, \tag{2.41}$$

which is the first marginal of $\mu_{N,\alpha}$. The Dirichlet form of the reversible pair $(P, P_1\mu_{N,\alpha})$ is then

$$\mathcal{E}_P(f,g) = \sum_{\ell\in[n]} \mathbb{E}_{\tau\sim P_1\mu_{N,\alpha}}\Big[P(\tau,\tau^{\ell})\big(f(\tau) - f(\tau^{\ell})\big)\big(g(\tau) - g(\tau^{\ell})\big)\Big], \tag{2.42}$$

for $f, g\colon \Omega_G \to \mathbb{R}$.

The proof of the log-Sobolev inequality for P relies on a quantitative comparison between the canonical marginal $P_1\mu_{N,\alpha}$ and the grand-canonical hardcore measure $\mu_{G,\lambda}$ at the matching fugacity. This comparison is provided by the following local central limit theorem, whose proof we omit; it can be obtained by adapting the argument of [Jai+23, Corollary 13] to track the joint dependence on both N and n.

Proposition 2.28. *Let G be a graph on n vertices with maximum degree Δ, let N be the size of the particle system, let $\lambda \in (\gamma, (1-\delta)\lambda_c(\Delta))$ for some $\gamma, \delta > 0$, and let $\alpha \in D_N$. Denote by $\alpha(\mu_{G,\lambda})$ the density of $\mu_{G,\lambda}$ and suppose that*

$$Nn\,|\alpha - \alpha(\mu_{G,\lambda})|^2 = O_{\gamma,\delta,\Delta}(1).$$

Then

$$\left|\mu_{G,\lambda}^{\otimes N}(\Omega_{N,\alpha}) - \frac{1}{\sqrt{2\pi}\,v_N}\,\exp\Big(-\frac{N^2n^2|\alpha - \alpha(\mu_{G,\lambda})|^2}{2v_N^2}\Big)\right| = O_{\gamma,\delta,\Delta}\Big(\frac{1}{Nn}\Big),$$

where $v_N^2 = \mathrm{Var}_{\sigma\sim\mu_{G,\lambda}^{\otimes N}}(|\sigma|)$ satisfies $v_N^2 = \Theta_{\gamma,\delta,\Delta}(Nn)$.

We are now ready to state the log-Sobolev inequality for the induced chain.

Lemma 2.29. *If $\alpha \in (\gamma, (1-\delta)\alpha_c(\Delta))$ for some $\gamma, \delta > 0$ and $N \geq n$, then the matrix P defined in (2.40) satisfies a log-Sobolev inequality with respect to the Dirichlet form (2.42) with rate $r = \Omega_{\gamma,\delta,\Delta}(1/n)$.*

Proof. First, by the correspondence induced by the occupancy ratio $\alpha_G(\cdot)$, there exists $\lambda^* \in (c(\gamma), (1-c'(\delta))\lambda_c(\Delta))$ such that $\alpha_c(\lambda^*) = \alpha$, or equivalently $\alpha(\lambda^*) = \alpha$. Here $c(\gamma)$ and $c'(\delta)$ are positive constants.

We compare the dynamics generated by P with the Glauber dynamics P_{GD} of the hard-core model with fugacity λ^*, defined by

$$P_{GD}(\tau,\tau') = \frac{1}{n}\sum_{\ell\in[n]}\Big[\frac{\mu_{\lambda^*,G}(\tau^\ell)}{\mu_{\lambda^*,G}(\tau^\ell)+\mu_{\lambda^*,G}(\tau^{\bar\ell})}\mathbf{1}(\tau'=\tau^\ell) + \frac{\mu_{\lambda^*,G}(\tau^{\bar\ell})}{\mu_{\lambda^*,G}(\tau^\ell)+\mu_{\lambda^*,G}(\tau^{\bar\ell})}\mathbf{1}(\tau'=\tau^{\bar\ell})\Big].$$

In particular, $P_{GD}(\tau,\tau^\ell) = \frac{1}{n}\frac{\lambda^*}{1+\lambda^*}$ when $\tau^\ell \neq \tau$. The dynamics generated by P_{GD} satisfies a log-Sobolev inequality with rate $r = \Omega_{\gamma,\delta,\Delta}(1/n)$, as follows from Lemma 2.3 and Fact 3.5 in [CLV21]. For completeness, we briefly recall the main ingredients of the proof.

Lemma 2.3 in [CLV21] states that if a Gibbs measure is spectrally independent and marginally bounded, then it satisfies approximate tensorization of entropy. For the hard-core model, spectral independence in the subcritical regime is a classical result due to Weitz; see, for instance, [Wei06] for the seminal work and [CSV25] for more recent developments. Marginal boundedness is immediate, see e.g. Proposition 50 in [Ana+21], which ensures that for every subset $A \subset [n]$, for every σ_A such that $\sigma_A \cap N(\ell) = \emptyset$, for all $\lambda > 0$,

$$\frac{\lambda}{1+\lambda}(1+\lambda)^{-N(\ell)} \leq \mathbb{P}_{\sigma\sim\mu_{\lambda,G}}(\sigma_\ell = 1 \mid \sigma_A) \leq \frac{\lambda}{1+\lambda}.$$

In particular, marginal boundedness holds in the form

$$\min\Big\{\mathbb{P}_{\sigma\sim\mu_{\lambda^*,G}}(\sigma_\ell = 1 \mid \sigma_A), \mathbb{P}_{\sigma\sim\mu_{\lambda^*,G}}(\sigma_\ell = 0 \mid \sigma_A)\Big\} \geq b > 0, \tag{2.43}$$

for some $b = b(\gamma,\delta,\Delta) > 0$. By [CLV21, Lemma 2.3], this yields approximate tensorization of entropy,

$$\mathrm{Ent}_{\mu_{\lambda^*,G}}[f] \lesssim_{\gamma,\delta,\Delta} \sum_{\ell\in[n]} \mu_{\lambda^*,G}\Big[\mathrm{Ent}_{\mu_{\lambda^*,G}}\left[f \mid \{\sigma_k,\, k\neq\ell\}\right]\Big]. \tag{2.44}$$

Finally, using marginal boundedness (2.43) together with the classical result of Diaconis and Saloff-Coste (Appendix 1 in [DS96]), we obtain

$$\begin{aligned}\mathrm{Ent}_{\mu_{\lambda^*,G}}\left[f \mid \{\sigma_k,\, k\neq\ell\}\right] &\leq \frac{1-2b}{\log(1/b-1)} Var_{\mu_{\lambda^*,G}}[\sqrt{f} \mid \{\sigma_k,\, k\neq\ell\}] \\ &\lesssim_{\gamma,\delta,\Delta} Var_{\mu_{\lambda^*,G}}[\sqrt{f} \mid \{\sigma_k,\, k\neq\ell\}],\end{aligned}$$

which, combined with (2.44), yields

$$\mathrm{Ent}_{\mu_{\lambda^*,G}}[f] \lesssim_{\gamma,\delta,\Delta} \sum_{\ell\in[n]} \mu_{\lambda^*,G}\Big[Var_{\mu_{\lambda^*,G}}[\sqrt{f} \mid \{\sigma_k,\, k\neq\ell\}]\Big].$$

This corresponds to a log-Sobolev inequality for the Glauber dynamics with rate $r = \Omega_{\gamma,\delta,\Delta}(1/n)$.

Therefore, to conclude the proof of the lemma, it is sufficient to compare the two chains

P and P_{GD}. Specifically, we need that if $N \geq n$,

$$1)\ \frac{P_1\mu_{N,\alpha}}{\mu_{G,\lambda^*}} = \Theta_{\gamma,\delta,\Delta}(1), \qquad 2)\ \frac{P_{GD}(\tau,\tau^\ell)}{P(\tau,\tau^\ell)} = O_{\gamma,\delta,\Delta}(1),\ \forall \ell \in [n],\ \forall \tau \in \Omega_G :\ \tau \neq \tau^\ell. \tag{2.45}$$

Assuming (2.45), standard comparison arguments yield, for any $f : \Omega_G \to \mathbb{R}_+$,

$$\begin{aligned}
\mathrm{Ent}_{\mu_{N,\alpha,1}}[f] &\lesssim_{\gamma,\delta,\Delta} \mathrm{Ent}_{\mu_{\lambda^*,G}}[f] \\
&\lesssim_{\gamma,\delta,\Delta} n\mathcal{E}_{P_{GD}}(\sqrt{f},\sqrt{f}) \\
&= \sum_{\tau\in\Omega_G}\sum_{\ell\in[n]} \mu_{\lambda^*,G}(\tau)P_{GD}(\tau,\tau^\ell)\big(\sqrt{f(\tau)} - \sqrt{f(\tau^\ell)}\big)^2 \\
&\lesssim_{\gamma,\delta,\Delta} \sum_{\ell\in[n]} \mu_{N,\alpha,1}(\tau)P(\tau,\tau^\ell)\big(\sqrt{f(\tau)} - \sqrt{f(\tau^\ell)}\big)^2 = n\mathcal{E}_P(\sqrt{f},\sqrt{f}).
\end{aligned}$$

In the first row above, we used Holley-Strook variational representation for entropy:

$$Ent_\mu[F] = \inf_{a\in\mathbb{R}_+} \mu[F\log F - F\log a + a - F].$$

We begin with point 1) in (2.45). For any $\tau \in \Omega_G$, define

$$\alpha^\tau = \frac{N}{N-1}\alpha - \frac{1}{(N-1)n}|\tau|,$$

which is the density induced on the subsystem $\sigma(2),\dots,\sigma(N)$ when fixing $\sigma(1) = \tau$. By homogeneity, $\mu_{N,\alpha}$ can be written as $\mu^{\otimes N}_{\lambda^*,G}(\cdot \mid \Omega_{N,\alpha})$. Hence, for any $\tau \in \Omega_G$,

$$\begin{aligned}
P_1\mu_{N,\alpha}(\tau) &= \frac{\mathbb{P}_{\sigma\sim\mu^{\otimes N}_{\lambda^*,G}}(\sigma(1)=\tau,\ \sigma\in\Omega_{N,\alpha})}{\mathbb{P}_{\sigma\sim\mu^{\otimes N}_{\lambda^*,G}}(\sigma\in\Omega_{N,\alpha})} \\
&= \mu_{\lambda^*,G}(\tau)\frac{\mathbb{P}_{\sigma\sim\mu^{\otimes N-1}_{\lambda^*,G}}(\sigma\in\Omega_{N-1,\alpha^\tau})}{\mathbb{P}_{\sigma\sim\mu^{\otimes N}_{\lambda^*,G}}(\sigma\in\Omega_{N,\alpha})}.
\end{aligned}$$

By Proposition 2.28 below, when $N \geq n$ the ratio above is $\Theta_{\gamma,\delta,\Delta}(1)$, which proves point 1).

We now turn to point 2). First, we observe that

$$P_{GD}(\tau,\tau^\ell) = \frac{1}{n}\frac{\lambda^*}{1+\lambda^*} \leq \frac{1}{n}\frac{\lambda_c(\Delta)}{1+\lambda_c(\Delta)} = O_\Delta\left(\frac{1}{n}\right). \tag{2.46}$$

On the other hand,

$$P(\tau,\tau^\ell) = \frac{1}{n}\frac{N-1}{N}\alpha^\tau = \Omega_{\gamma,\delta,\Delta}(1/n), \tag{2.47}$$

where we assume N to be sufficiently large. Combining (2.46) and (2.47) concludes the proof of point 2), and hence the lemma. □

2.4.5 Proof of Lemma 2.25

The proof of Lemma 2.25 adapts the strategy of [Jai+23] to track the joint dependence on N and n. The idea is the following: Lemma 2.29 provides a log-Sobolev inequality for the induced chain governing a single particle, in terms of its Dirichlet form. To deduce a log-Sobolev inequality for P_N, one needs to compare the average over $i \in [N]$ of these single-particle Dirichlet forms with the global Dirichlet form of P_N. This comparison is carried out term by term: each contribution to the Dirichlet form of an induced chain is matched with a suitable combination of terms in the Dirichlet form of P_N.

The matching step is technically delicate. It relies on Stein's method for Markov chains [BN19; RR19], combined with a contractive coupling argument for the down-up walk in the small-density regime—similar in spirit to the path coupling technique used to bound the spectral gap of the down-up walk on independent sets of fixed size in [BD97]. The contractive coupling can be implemented only when $\alpha = O(1/\Delta^2)$, which is the source of the smallness assumption on α. We isolate this part of the argument in Lemma 2.35, deferred to the appendix of this chapter; see Proposition A.1.

The combination of Lemma 2.29 with the matching step just described leads to the second key ingredient, Lemma 2.30 below. Before stating it, we set up the relevant notation. For each $i \in [N]$ and each $F \colon \Omega_{N,\alpha} \to \mathbb{R}_+$, define

$$f_i(\tau) := \mathbb{E}_{\mu_{N,\alpha}}[F \mid \sigma(i) = \tau], \qquad \tau \in \Omega_G.$$

Lemma 2.30. *Under the hypotheses of Lemma 2.23, for every $\ell \in [n]$ and every $i \in [N]$,*

$$\left(\sqrt{f_i(\tau \setminus \{\ell\})} - \sqrt{f_i(\tau^\ell)}\right)^2 \lesssim_{\gamma,\Delta} \frac{1}{Nn} \sum_{(j,k)} \mathbb{E}_{\sigma\sim\mu_{N,\alpha}}\left[\mathbf{1}\{(\ell,k)_{\sigma(i),\sigma(j)}\} \left(\sqrt{F(\sigma)} - \sqrt{F(\sigma^{(i,\ell;j,k)})}\right)^2 \Big| \sigma(i) = \tau^\ell\right].$$

The remainder of this subsection is organized as follows. We first show how Lemma 2.30 implies Lemma 2.25. We then build the toolbox needed to prove Lemma 2.30: a change-of-measure inequality and Stein's method. Finally, we state the two key estimates and combine them to conclude the proof. By symmetry, it suffices to prove Lemma 2.30 for $i = 1$, which we assume henceforth.

Reduction to Lemma 2.30

From Lemma 2.29 we know that the induced chain on component $i \in [N]$ satisfies a log-Sobolev inequality with rate $\Omega_{\gamma,\Delta}(1/n)$. In particular,

$$\begin{aligned}\mathrm{Ent}(f_i) &\lesssim_{\gamma,\Delta} n \sum_{\tau\in\Omega_G} \mu_{N,\alpha}(\sigma(i)=\tau) \sum_{\ell\in[n]} P(\tau,\tau\setminus\{\ell\})\Big(\sqrt{f_i(\tau)}-\sqrt{f_i(\tau\setminus\{\ell\})}\Big)^2 \\ &\lesssim_{\gamma,\Delta} \sum_{\tau\in\Omega_G} \mu_{N,\alpha}(\sigma(i)=\tau) \sum_{\ell\in[n]} \Big(\sqrt{f_i(\tau)}-\sqrt{f_i(\tau\setminus\{\ell\})}\Big)^2,\end{aligned}$$

where we used $P(\tau,\tau\setminus\{\ell\}) = \Omega_{\gamma,\Delta}(1/n)$. The nonzero terms in the last sum must satisfy $\tau = \tau^\ell$. Combining this observation with Lemma 2.30 yields

$$\mathrm{Ent}(f_i) \lesssim_{\gamma,\Delta} \frac{1}{Nn} \sum_{\ell\in[n]} \sum_{(j,k)} \mathbb{E}_{\sigma\sim\mu_{N,\alpha}}\Big[\mathbf{1}\{(\ell,k)_{\sigma(i),\sigma(j)}\}\Big(\sqrt{F(\sigma)}-\sqrt{F(\sigma^{(i,\ell;j,k)})}\Big)^2\Big].$$

Summing over $i \in [N]$,

$$\frac{1}{N}\sum_{i\in[N]} \mathrm{Ent}(f_i) \lesssim_{\gamma,\Delta} n\,\bar{\mathcal{E}}_{N,\alpha}[\sqrt{F},\sqrt{F}].$$

This reduces the proof of Lemma 2.25 to that of Lemma 2.30. The remainder of this subsection is devoted to the latter. □

Toolbox: change of measure

Let $\tau \in \Omega_G$ be such that the site $\ell \in [n]$ is available for an update inserting a 1 at ℓ, so that $\tau^\ell = \tau \cup \{\ell\}$ and $\tau^{\bar\ell} = \tau\setminus\{\ell\}$. Given a density $\alpha \in D_N$, define

$$\nu_\ell := \mu_{N,\alpha}[\,\cdot \mid \sigma(1) = \tau^\ell], \quad \nu_{\bar\ell} := \mu_{N,\alpha}[\,\cdot \mid \sigma(1) = \tau^{\bar\ell}], \quad \nu := \mu_{N,\alpha}[\,\cdot \mid \sigma(1) \in \{\tau^\ell, \tau^{\bar\ell}\}]. \tag{2.48}$$

The first two are obtained from $\mu_{N,\alpha}$ by conditioning the first particle to equal τ^ℓ or $\tau^{\bar\ell}$, respectively; the third fixes $\sigma(1)$ to agree with τ on every site $k \neq \ell$, leaving the value at site ℓ unspecified. These three measures are linked by the identity

$$\nu = \theta\,\nu_\ell + (1-\theta)\,\nu_{\bar\ell}, \qquad \theta := \mu_{N,\alpha}(\sigma_\ell(1) = 1 \,|\, \sigma(1)_k = \tau_k,\ k \neq \ell). \tag{2.49}$$

The left-hand side of the inequality we wish to establish in Lemma 2.30 can be rewritten as $(\sqrt{\mathbb{E}_{\nu_\ell}F} - \sqrt{\mathbb{E}_{\nu_{\bar\ell}}F})^2$, for some $F\colon \Omega_{N,\alpha} \to \mathbb{R}_+$. It is convenient to replace $\nu_{\bar\ell}$ with ν, the latter leaving the state at site ℓ unspecified, leading to the following lemma.

Lemma 2.31 ([Jai+23], Lemma 36)**.** *Let $\alpha \in D_N$ satisfy $\alpha \in (\gamma, (1-\delta)\alpha_c(\Delta))$ and assume $N \geq n$. Then*

$$\Big(\sqrt{\mathbb{E}_{\nu_\ell}F} - \sqrt{\mathbb{E}_{\nu_{\bar\ell}}F}\Big)^2 \lesssim_{\gamma,\delta,\Delta} \Big(\sqrt{\mathbb{E}_{\nu_\ell}F} - \sqrt{\mathbb{E}_{\nu}F}\Big)^2. \tag{2.50}$$

The proof of Lemma 2.31 combines two elementary steps: a rationalization of both sides

of (2.50), and a control on the mixing weight θ in (2.49). The latter is provided by Proposition 2.32 below; granting it, the proof is straightforward.

Proof of Lemma 2.31. By Proposition 2.32, $\theta = \Theta_{\gamma,\delta,\Delta}(1)$, and the decomposition (2.49) gives $\mathrm{d}\nu_\ell/\mathrm{d}\nu = \theta$ and $\mathrm{d}\nu_{\bar\ell}/\mathrm{d}\nu = 1-\theta$. We rewrite the left-hand side of (2.50) as a ratio:

$$\Big(\sqrt{\mathbb{E}_{\nu_\ell} F} - \sqrt{\mathbb{E}_{\nu_{\bar\ell}} F}\Big)^2 = \frac{\big(\mathbb{E}_{\nu_\ell} F - \mathbb{E}_{\nu_{\bar\ell}} F\big)^2}{\big(\sqrt{\mathbb{E}_{\nu_\ell} F} + \sqrt{\mathbb{E}_{\nu_{\bar\ell}} F}\big)^2} \le \frac{\big(\mathbb{E}_{\nu_\ell} F - \mathbb{E}_{\nu_{\bar\ell}} F\big)^2}{\mathbb{E}_{\nu_\ell} F + \mathbb{E}_{\nu_{\bar\ell}} F}. \tag{2.51}$$

Similarly, viewing ν as the θ-convex combination in (2.49), we have $\mathbb{E}_{\nu_\ell} F - \mathbb{E}_\nu F = (1-\theta)\big(\mathbb{E}_{\nu_\ell} F - \mathbb{E}_{\nu_{\bar\ell}} F\big)$, hence

$$\Big(\sqrt{\mathbb{E}_{\nu_\ell} F} - \sqrt{\mathbb{E}_\nu F}\Big)^2 = \frac{\big(\mathbb{E}_{\nu_\ell} F - \mathbb{E}_\nu F\big)^2}{\big(\sqrt{\mathbb{E}_{\nu_\ell} F} + \sqrt{\mathbb{E}_\nu F}\big)^2} \ge \frac{(1-\theta)^2\big(\mathbb{E}_{\nu_\ell} F - \mathbb{E}_{\nu_{\bar\ell}} F\big)^2}{4\big(\mathbb{E}_{\nu_\ell} F + \mathbb{E}_{\nu_{\bar\ell}} F\big)}. \tag{2.52}$$

Combining (2.51) and (2.52) yields (2.50). □

It remains to prove that the mixing weight θ is bounded away from 0 and 1, uniformly in τ and ℓ. This relies on the local central limit theorem of Proposition 2.28.

Proposition 2.32. *Let $\alpha \in D_N$ satisfy $\alpha \in (\gamma, (1-\delta)\alpha_c(\Delta))$ and assume $N \ge n$. Then, for all $\tau \in \Omega_G$ with $\tau_{N(\ell)} = \emptyset$,*

$$\mu_{N,\alpha}\big(\sigma_\ell(1) = 1 \,\big|\, \sigma(1)_k = \tau_k,\ k \ne \ell\big) = \Omega_{\gamma,\delta,\Delta}(1),$$
$$\mu_{N,\alpha}\big(\sigma_\ell(1) = 0 \,\big|\, \sigma(1)_k = \tau_k,\ k \ne \ell\big) = \Omega_{\gamma,\delta,\Delta}(1).$$

In particular, $\mathrm{d}\nu_\ell/\mathrm{d}\nu$ and $\mathrm{d}\nu_{\bar\ell}/\mathrm{d}\nu$ are both $O_{\gamma,\delta,\Delta}(1)$.

Proof. By Lemma 2.12, there exists $\lambda > 0$ such that $\alpha(\mu_{G,\lambda}) = \alpha$. A direct computation gives

$$\mu_{N,\alpha}\big(\sigma_\ell(1) = 1 \,\big|\, \sigma(1)_k = \tau_k,\ k \ne \ell\big) = \frac{1}{1 + z(\ell, \tau, N, \alpha)},$$

where

$$z(\ell, \tau, N, \alpha) = \frac{\mu_{G,\lambda}(\tau^{\bar\ell})}{\mu_{G,\lambda}(\tau^\ell)} \cdot \frac{\mu_{G,\lambda}^{\otimes(N-1)}(\Omega_{N-1,\alpha^{\tau^{\bar\ell}}})}{\mu_{G,\lambda}^{\otimes(N-1)}(\Omega_{N-1,\alpha^{\tau^\ell}})}.$$

Since $N \ge n$, one has $(N-1)n\,|\alpha - \alpha^\tau|^2 \le 1$ for all $\tau \in \Omega_G$, so Proposition 2.28 yields

$$\frac{\mu_{G,\lambda}^{\otimes(N-1)}(\Omega_{N-1,\alpha^{\tau^{\bar\ell}}})}{\mu_{G,\lambda}^{\otimes(N-1)}(\Omega_{N-1,\alpha^{\tau^\ell}})} = \Theta_{\gamma,\delta,\Delta}(1).$$

Moreover, since $\tau_{N(\ell)} = \emptyset$, both τ^ℓ and $\tau^{\bar\ell}$ belong to Ω_G, and

$$\frac{\mu_{G,\lambda}(\tau^{\bar\ell})}{\mu_{G,\lambda}(\tau^\ell)} = \frac{\mu_{G,\lambda}(\sigma_\ell = 0 \mid \sigma_k = \tau_k,\ k \ne \ell)}{\mu_{G,\lambda}(\sigma_\ell = 1 \mid \sigma_k = \tau_k,\ k \ne \ell)} = \frac{1}{\lambda} = \Theta_{\gamma,\delta,\Delta}(1),$$

since $\lambda = \Theta_{\gamma,\delta,\Delta}(1)$. This concludes the proof. □

Toolbox: Stein's method

We now use Stein's method to estimate the right-hand side of (2.50). Let $P_{\nu,N}$ be the Markov chain on $\Omega_{N,\alpha}$ obtained from P_N by restricting the dynamics to the support of ν; that is, the transition matrix

$$P_{\nu,N}(\sigma,\sigma') := P_N(\sigma,\sigma')\,\mathbf{1}\{\sigma,\sigma' \in \mathrm{supp}(\nu)\}, \qquad \sigma \neq \sigma'.$$

In words, if the proposed update changes $\sigma(1)$ at a site different from ℓ, the update is rejected; in all other cases $P_{\nu,N}$ acts as P_N. We define $P_{\nu_\ell,N}$ analogously, replacing ν with ν_ℓ. The measures ν and ν_ℓ are reversible for $P_{\nu,N}$ and $P_{\nu_\ell,N}$, respectively.

Let $h\colon \Omega_{N,\alpha} \to \mathbb{R}$ be the solution of the Poisson equation

$$h - P_{\nu,N}\, h = F - \mathbb{E}_\nu[F],$$

which has the explicit form

$$h(\sigma) = \sum_{t=0}^{+\infty} \mathbb{E}\Big[F(\sigma_t) - \mathbb{E}_\nu[F] \,\Big|\, \sigma_0 = \sigma\Big],$$

where $(\sigma_t, t \in \mathbb{N})$ is the Markov chain with transition matrix $P_{\nu,N}$, and $\mathbb{E}$ denotes expectation with respect to its law. Using h, we can express

$$\begin{aligned} \mathbb{E}_{\nu_\ell}[F] - \mathbb{E}_\nu[F] &= \mathbb{E}_{\nu_\ell}[h - P_{\nu,N}\, h] \\ &= \mathbb{E}_{\nu_\ell}[P_{\nu_\ell,N}\, h - P_{\nu,N}\, h], \end{aligned} \tag{2.53}$$

where we used the invariance of ν_ℓ under $P_{\nu_\ell,N}$ to write $\mathbb{E}_{\nu_\ell}[P_{\nu_\ell,N}\, h] = \mathbb{E}_{\nu_\ell}[h]$. Combining (2.53) with (2.50), we obtain

$$\Big(\sqrt{\mathbb{E}_{\nu_\ell} F} - \sqrt{\mathbb{E}_{\nu_{\bar\ell}} F}\Big)^2 \lesssim_{\gamma,\delta,\Delta} \frac{\big|\mathbb{E}_{\nu_\ell}[P_{\nu_\ell,N}\, h - P_{\nu,N}\, h]\big|^2}{\mathbb{E}_\nu[F]}. \tag{2.54}$$

Two key estimates

Given $\eta_0, \eta_0' \in \Omega_{N,\alpha}$, denote by (η_t, η_t') the Markov chains started at (η_0, η_0') with transition matrix $P_{\nu,N}$.

Definition 2.33 (Distribution Γ)**.** Define a probability distribution Γ on pairs (η_0, η_0') by:

1) choose (j,k) uniformly at random among all pairs with $j \geq 2$ and $k \in [n]$;

2) sample $\eta_0 \sim \mu_{N,\alpha}$ conditioned on the event $\sigma(1) = \tau^\ell$ and $\sigma(j)_{k'} = 0$ for all $k' \in N^+(k)$;

3) set $\eta_0' := (\eta_0)^{1,\ell;\,j,k}$.

Equivalently,

$$\Gamma(\eta_0,\eta_0') = \frac{1}{(N-1)n}\sum_{j\geq 2}\sum_{k\in[n]} \mu_{N,\alpha}(\eta_0 \mid \sigma(1)=\tau^\ell,\, \sigma(j)_{N^+(k)}=0)\mathbf{1}\{\eta_0' = (\eta_0)^{1,\ell;\,j,k}\}.$$

For brevity, given a coupling of (η_t,η_t') with $(\eta_0,\eta_0')\sim\Gamma$ and $c>0$, set

$$H(c,\Gamma) := \sum_{t=0}^{\infty} e^{ct/(nN)}\,\mathbb{E}\Big[\big(\sqrt{F(\eta_t)}-\sqrt{F(\eta_t')}\big)^2\Big]. \tag{2.55}$$

We can now state the two key estimates. The first bounds the right-hand side of (2.54) in terms of $H(c,\Gamma)$ for *any* coupling; the second exhibits a specific coupling for which $H(c,\Gamma)$ is controlled by the Dirichlet-type quantity at $t=0$. Both proofs are deferred to the appendix of this chapter.

Lemma 2.34. *Suppose $\alpha\in D_N$ lies in the subcritical interval $(\gamma,(1-\delta)\alpha_c(\Delta))$ for some $\gamma,\delta>0$, and assume $N\geq 2n$. Then, for any coupling of (η_t,η_t') with $(\eta_0,\eta_0')\sim\Gamma$ and any $c>0$,*

$$\frac{\big|\mathbb{E}_{\nu_\ell}[P_{\nu_\ell,N}\,h - P_{\nu,N}\,h]\big|^2}{\mathbb{E}_\nu[F]} \lesssim_{\gamma,\delta,\Delta} \frac{1}{cNn}\,H(c,\Gamma). \tag{2.56}$$

Lemma 2.35. *Suppose $\alpha\in D_N$ satisfies $\alpha\in(0,1/(C\Delta^2))$ with $C\geq 241$. Then there exists a coupling of (η_t,η_t') with $(\eta_0,\eta_0')\sim\Gamma$ such that, for all $c\leq 1/1000$,*

$$H(c,\Gamma)\leq Nn\,\mathbb{E}_{(\eta_0,\eta_0')\sim\Gamma}\Big[\big(\sqrt{F(\eta_0)}-\sqrt{F(\eta_0')}\big)^2\Big]. \tag{2.57}$$

Proof of Lemma 2.30

Throughout this proof we apply, in sequence, the inequalities (2.54), (2.56), and (2.57). The restriction

$$\alpha\leq\frac{1}{241\,\Delta^2}$$

is required only for (2.57); this condition lies within the subcritical regime, since $1/(241\,\Delta^2)<\alpha_c(\Delta)$. Assume now

$$\gamma\leq\alpha\leq\frac{1}{241\,\Delta^2}\quad\text{for some }\gamma>0.$$

Under this assumption, all inequalities in this section depend only on γ and Δ, since the gap between $1/(241\,\Delta^2)$ and $\alpha_c(\Delta)$ is a function of Δ alone. Accordingly, we use the notation $\lesssim_{\gamma,\Delta}$.

Fix $\tau\in\Omega_G$ and let $\ell\in[n]$ be a site available for τ, i.e., $\tau\cup\{\ell\}\in\Omega_G$; otherwise, the statement of Lemma 2.30 is trivial. By definition of f_1 and of $\nu_\ell,\nu_{\bar\ell}$ in (2.48), we have

$f_1(\tau^\ell) = \mathbb{E}_{\nu_\ell} F$ and $f_1(\tau \setminus \{\ell\}) = \mathbb{E}_{\nu_{\bar{\ell}}} F$. Taking $c \le 1/1000$,

$$\begin{aligned}\Big(\sqrt{f_1(\tau \setminus \{\ell\})} - \sqrt{f_1(\tau^\ell)}\Big)^2 &= \Big(\sqrt{\mathbb{E}_{\nu_\ell} F} - \sqrt{\mathbb{E}_{\nu_{\bar{\ell}}} F}\Big)^2 \\ &\overset{(2.54)}{\lesssim_{\gamma,\Delta}} \frac{\big|\mathbb{E}_{\nu_\ell}[P_{\nu_\ell,N}\, h - P_{\nu,N}\, h]\big|^2}{\mathbb{E}_\nu[F]} \\ &\overset{(2.56)}{\lesssim_{\gamma,\Delta}} \frac{1}{cNn}\, H(c,\Gamma) \\ &\overset{(2.57)}{\le} \mathbb{E}_{(\eta_0,\eta_0')\sim\Gamma}\Big[\big(\sqrt{F(\eta_0)} - \sqrt{F(\eta_0')}\big)^2\Big].\end{aligned}$$

Expanding the last expectation,

$$\begin{aligned}&\mathbb{E}_{(\eta_0,\eta_0')\sim\Gamma}\Big[\big(\sqrt{F(\eta_0)} - \sqrt{F(\eta_0')}\big)^2\Big] \\ &\quad= \frac{1}{(N-1)n} \sum_\sigma \sum_{j>1,k} \mu_{N,\alpha}\big(\sigma \mid \sigma(1) = \tau^\ell,\, \sigma_{N^+(k)}(j) = 0\big)\Big(\sqrt{F(\sigma)} - \sqrt{F(\sigma^{1,\ell;\,j,k})}\Big)^2 \\ &\quad\lesssim_{\gamma,\Delta} \frac{1}{Nn} \sum_\sigma \sum_{j,k} \mu_{N,\alpha}\big(\sigma \mid \sigma(1) = \tau^\ell\big)\, \mathbf{1}\{(\ell,k)_{\sigma(1),\sigma(j)}\}\Big(\sqrt{F(\sigma)} - \sqrt{F(\sigma^{1,\ell;\,j,k})}\Big)^2,\end{aligned}$$

where we used the inequality

$$\mu_{N,\alpha}\big(\sigma \mid \sigma(1) = \tau^\ell,\, \sigma_{N^+(k)}(j) = 0\big) \lesssim_{\gamma,\Delta} \mu_{N,\alpha}\big(\sigma \mid \sigma(1) = \tau^\ell\big)\, \mathbf{1}\{(\ell,k)_{\sigma(1),\sigma(j)}\},$$

which follows from a single application of the local central limit theorem of Proposition 2.28; see, for instance, the analogous argument in (A.14). □

2.5 Completing Kac's Program

In the following we prove entropy decay for the nonlinear mean-field hard-core model, following essentially the same strategy as in Section 1.5.

Theorem 2.36 (Entropy decay)**.** *Let $G = (V, E)$ be a graph on n vertices with maximum degree $\Delta \ge 3$. For any $\gamma > 0$ and $\delta \in (0,1)$, and any initial state p_0 with density $\alpha(p_0) \in (\gamma, (1-\delta)\alpha_c(\Delta))$, the solution p_t of the mean-field dynamics (2.4) satisfies*

$$H(p_t \mid \mu_{G,\lambda}) \;\le\; e^{-r\,t/n}\, H(p_0 \mid \mu_{G,\lambda}), \qquad t \ge 0,$$

for some constant $r = r(\gamma, \delta, \Delta) > 0$ independent of n, where $\lambda = \lambda(p_0) < \lambda_c(\Delta)$ is the unique fugacity matching the density of p_0.

Theorem 2.36 is established by completing Kac's program: we deduce the entropy decay of the nonlinear mean-field dynamics from a modified log-Sobolev inequality (MLSI) for the mean-field particle system, which in turn follows from Theorem 2.6, and transfer it to the nonlinear process by taking the limit $N \to \infty$. We start by discussing the functional inequality associated to the entropy decay of the continuous time nonlinear dynamics.

Nonlinear Modified Log-Sobolev

For any $\lambda > 0$, taking $\mu = \mu_{G,\lambda}$, we consider the functional

$$\mathcal{D}_\mu(f,f) = \frac{1}{4n^2} \sum_{\sigma,\sigma',\tau,\tau'} \mu(\sigma)\mu(\sigma')Q(\sigma,\sigma';\tau,\tau') \left(f(\sigma)f(\sigma') - f(\tau)f(\tau')\right) \log\left(\frac{f(\sigma)f(\sigma')}{f(\tau)f(\tau')}\right),$$

for all $f : \Omega_G \to \mathbb{R}_+$. This is known as the entropy dissipation functional. Indeed, Lemma 2.9 shows that the solution p_t of (2.4), for any initial state $p_0 \in \mathcal{P}(\Omega_G)$, satisfies

$$\frac{d}{dt} H(p_t \mid \mu) = -\mathcal{D}_\mu(f_t, f_t), \tag{2.58}$$

where $f_t = p_t/\mu$ is the relative density of p_t with respect to μ.

Since $(x-y)\log(\frac{x}{y}) \geq 0$ for all $x, y \geq 0$, we have $\mathcal{D}_\mu(f,f) \geq 0$, for all f, so that $H(p_t \mid \mu)$ is monotone non increasing as a function of $t \geq 0$. This is the analogue of Boltzmann's H-theorem in our setting.

We reformulate Theorem 2.36 in terms of the following nonlinear modified log-Sobolev inequality. We use the notation

$$\mathrm{Ent}_\mu[f] = \mu\left[f \log\left(f / \log \mu[f]\right)\right],$$

for all $f : \Omega_G \to \mathbb{R}_+$. In particular, when $\mu[f] = 1$ the functional returns the relative entropy $\mathrm{Ent}_\mu[f] = H(f\mu \mid \mu)$.

Definition 2.37 (Nonlinear Modified Log Sobolev (NMLSI))**.** For all $\lambda > 0$, let $\mathcal{S}_\lambda$ be the set of functions $f : \Omega_G \to \mathbb{R}_+$ with $\mu_{G,\lambda}[f] = 1$ such that $f\mu_{G,\lambda} \in \mathcal{P}(\Omega_G)$ has the same density as $\mu_{G,\lambda}$,

$$\alpha\left(f\mu_{G,\lambda}\right) = \alpha\left(\mu_{G,\lambda}\right).$$

The nonlinear dynamics is said to satisfy the NMLSI with profile function $\mathcal{I} : (0,\infty) \to [0,\infty)$ if for all $\lambda > 0$,

$$\mathcal{D}_{\mu_{G,\lambda}}(f,f) \geq \mathcal{I}(\lambda)\, \mathrm{Ent}_{\mu_{G,\lambda}}[f], \quad \forall f \in \mathcal{S}_\lambda. \tag{2.59}$$

Since the nonlinear mean-field dynamics conserves the density, any solution p_t satisfies $\alpha(p_t) = \alpha(p_0)$ for all $t \geq 0$. As a consequence, for any initial datum $p_0 \in \mathcal{P}(\Omega_G)$, one has $f_t = p_t/\mu_{G,\lambda} \in \mathcal{S}_\lambda$ for all $t \geq 0$, where λ is the unique value of the fugacity matching the density $\alpha(p_0)$. Therefore, (2.58) implies that, if the NMLSI holds with profile $\mathcal{I}(\cdot)$, then every solution p_t satisfies

$$H(p_t \mid \mu_{G,\lambda}) \leq e^{-\mathcal{I}(\lambda)t} H(p_0 \mid \mu_{G,\lambda}),$$

where λ is determined by $\alpha(p_0) = \alpha(\mu_{G,\lambda})$; see also Lemma 1.35 for a derivation of this statement in a slightly different context. In particular, Theorem 2.36 is established by

proving that the nonlinear system satisfies the NMLSI with a profile $\mathcal{I}$ such that $\mathcal{I}(\lambda) \geq r/n$, with a constant $r = r(\gamma, \delta, \Delta) > 0$, for any $\lambda > 0$ such that $\alpha(\mu_{G,\lambda}) \in (\gamma, (1-\delta)\alpha_c(\Delta))$.

Proof of Theorem 2.36

Fix $\gamma > 0$, $\delta \in (0,1)$ and $\alpha \in (\gamma, (1-\delta)\alpha_c(\Delta))$, and let $\lambda > 0$ be such that $\alpha = \alpha(\mu_{G,\lambda})$. Consider the particle system defined by the chain P_N on the state space Ω_{N,α_N}, as in Definition 2.5, with density

$$\alpha_N = \frac{\lfloor Nn\alpha \rfloor}{Nn},$$

and note that $\alpha_N \in (\gamma/2, (1-\delta)\alpha_c(\Delta))$ for all N large enough. Recall the unique invariant measure μ_{N,α_N}, and observe that $\mu_{N,\alpha_N} = \mu_{G,\lambda}^{\otimes N}(\cdot \mid \Omega_{N,\alpha_N})$, that is the N-fold product of $\mu_{G,\lambda}$ restricted to Ω_{N,α_N}. Pick $f \in \mathcal{S}_\lambda$ with $f > 0$. By definition, $\nu = f\mu_{G,\lambda} \in \mathcal{P}_+(\Omega_G)$ with $\alpha(\nu) = \alpha$; define the function F_N on Ω_{N,α_N} by

$$F_N := \frac{\nu^{\otimes N}(\cdot \mid \Omega_{N,\alpha_N})}{\mu_{N,\alpha_N}}.$$

Let $N \geq n^2$, then the MLSI (2.8) implies

$$\mathcal{E}_{N,\alpha_N}(F_N, \log F_N) \geq \frac{r_0}{n} \operatorname{Ent}_{\mu_{N,\alpha_N}}[F_N], \tag{2.60}$$

with $r_0 = r_0(\gamma/2, \delta, \Delta) > 0$ independent of N, n.

From now on, we proceed as in Section 1.5. Since $\nu = f\mu_{G,\lambda} \in \mathcal{P}_+(\Omega_G)$ and $|\alpha_N - \alpha| \leq \frac{1}{N}$, by adapting the arguments in Proposition 1.42 one proves the so-called entropic chaos, that is

$$\lim_{N\to\infty} \frac{1}{N} \operatorname{Ent}_{\mu_{N,\alpha_N}^{\mathrm{mf}}}[F_N] = \operatorname{Ent}_{\mu_{G,\lambda}}[f].$$

Similarly, as in Proposition 1.43, one finds the so-called Fisher chaos, that is

$$\lim_{N\to\infty} \frac{1}{N} \mathcal{E}_{N,\alpha_N}(F_N, \log F_N) = 2\mathcal{D}_\mu(f, f).$$

Therefore, dividing by N and passing to the limit $N \to \infty$ in (2.60) yields (2.59) with

$$\mathcal{I}(\lambda) = \frac{r_0}{2n},$$

for every $f \in \mathcal{S}_\lambda$ with full support in Ω_G. It remains to extend the estimate to all $f \in \mathcal{S}_\lambda$. To this end, given $f \in \mathcal{S}_\lambda$, consider

$$f_\varepsilon(\sigma) = (1-\varepsilon)f(\sigma) + \varepsilon, \qquad \varepsilon \in (0,1),\ \sigma \in \Omega_G.$$

Note that $f_\varepsilon \in \mathcal{S}_\lambda$ for all $\varepsilon \in (0,1)$, since $\mu_{G,\lambda}[f_\varepsilon] = 1$ and $\alpha(f_\varepsilon\mu_{G,\lambda}) = (1-\varepsilon)\alpha(f\mu_{G,\lambda}) +$

$\varepsilon\,\alpha(\mu_{G,\lambda}) = \alpha(\mu_{G,\lambda})$. Since f has full support on Ω_G,

$$\mathcal{D}_{\mu_{G,\lambda}}(f_\varepsilon, f_\varepsilon) \geq \frac{r_0}{2n}\, \mathrm{Ent}_{\mu_{G,\lambda}}[f_\varepsilon]\,.$$

Letting $\varepsilon \to 0$, by continuity one has

$$\mathcal{D}_{\mu_{G,\lambda}}(f_\varepsilon, f_\varepsilon) \to \mathcal{D}_{\mu_{G,\lambda}}(f, f), \quad \mathrm{Ent}_{\mu_{G,\lambda}}[f_\varepsilon] \to \mathrm{Ent}_{\mu_{G,\lambda}}[f]\,,$$

which conclude the proof of the NMLSI with a profile $\mathcal{I}$ such that $\mathcal{I}(\lambda) \geq r/n$, with $r = r_0(\gamma/2, \delta, \Delta)/2 > 0$, for any $\lambda > 0$ such that $\alpha(\mu_{G,\lambda}) \in (\gamma, (1-\delta)\alpha_c(\Delta))$. □

Chapter 3

Algorithmic Applications to Hardcore Model

As discussed in the main introduction, the nonlinear exchange dynamics is the natural process for the sampling problem for the hard-core model μ_G^α with prescribed global density α or $\mu_G^{\boldsymbol{m}}$ with marginal vector $\boldsymbol{m}$, but due to its quadratic nature, and given our current understanding, it cannot be implemented directly at present. Following Kac [Kac56], we overcome this obstacle by introducing an auxiliary *particle system*: a linear Markov chain that approximates the nonlinear evolution and is built on the same exchange mechanism, mean-field for the prescribed density problem and single-site for the prescribed marginals one.

The idea is to consider a finite population of N particles $\sigma^{(1)}, \dots, \sigma^{(N)} \in \Omega_G$, undergoing pairwise collisions

$$(\sigma^{(i)}, \sigma^{(j)}) \mapsto (\tau^{(i)}, \tau^{(j)}),$$

where the pair on the right is obtained by exchanging two bits $\sigma_v^{(i)}$ and $\sigma_u^{(j)}$, the sites u, v being chosen so as to match the nonlinear collision kernel. This defined a linear Markov chain implementable provided N being not too large. On the other hand, if N is sufficiently large, the particle system approximates the nonlinear evolution; since the latter converges to the desired Gibbs measure, after enough steps the law of a single particle should be close to the target measure. Particle systems of this kind already appeared in the previous chapters, in the analysis of the continuous-time nonlinear process. Here, for algorithmic purposes, we work instead with a discrete-time version, and with a different time scaling: whereas the process considered earlier performs on average N updates per unit of time, the one considered here performs a single collision per step.

The algorithms we present follow this general template:

(1) Initialize N particles so that the relevant conserved quantity (density α or marginal vector $\boldsymbol{m}$) has the desired value.

(2) Run the corresponding particle-system chain for t steps, with t large enough.

(3) Output one particle.

The analysis proceeds in two steps. First, we show that for N large enough the one-particle marginal of the stationary distribution of the particle system is close to the target hard-core measure. Second, we bound the mixing time of the chain, which determines how large t must be. Together, these give an output particle approximately distributed according to the target measure in polynomial time.

3.0.1 Plan of the chapter

The two constrained sampling problems require separate analyses, and the chapter is organized accordingly.

In Section 3.1 we introduce the two particle systems, associated with the mean-field and the single-site kernel respectively, and we study their convergence to equilibrium: Theorem 3.2 bounds the mixing time of the mean-field chain, and Theorem 3.5 that of the single-site chain. We then turn the particle systems into algorithms, treating the two sampling problems in separate sections. In Section 3.2.1 we address the prescribed density problem: Algorithm 4 takes a target density α as input and samples from μ_G^α, its correctness and running time being established in Theorem 3.9. In Section 3.2.2 we complement this with a hardness result, showing that the regime in which the algorithm operates is optimal. Finally, in Section 3.2.3 we address the prescribed marginals problem: Algorithm 6 takes a target marginal vector $\boldsymbol{m}$ and samples from $\mu_G^{\boldsymbol{m}}$, with correctness and running time established in Theorem 3.18.

3.1 The particle system

We consider a system of $N \in \mathbb{N}$ *particles* $\boldsymbol{\sigma} = \{\sigma^{(i)}, i = 1, \dots, N\}$, where $\sigma^{(i)} = (\sigma_u^{(i)}, u \in V) \in \Omega_G$, for each $i \in [N]$. Equivalently, $\boldsymbol{\sigma} \in \Omega_G^N$ can be seen as an independent set on the graph $G_N := G^{\sqcup N}$, the disjoint union of N copies $G(i)$ of G. The vertex set of G_N is $V_N = \{(i, v) : i \in [N], v \in V\}$. For $i \in [N]$ and a probability measure ν on Ω_G^N, let $P_i\nu$ denote the i-th one-particle marginal of ν, that is,

$$(P_i\nu)(\eta) := \nu\big(\{\boldsymbol{\sigma} \in \Omega_G^N : \sigma^{(i)} = \eta\}\big), \qquad \eta \in \Omega_G.$$

The evolution of the particle system is a Markov chain with state space Ω_G^N where the occupation variables $\sigma_v^{(i)}$ and $\sigma_w^{(j)}$ are exchanged between vertices $(i, v), (j, w) \in V_N$, with the constraint that the resulting configuration is again a valid independent set of G_N.

To define the Markov chain, we introduce some notation. Given vertices $(i, v), (j, w) \in$

V_N, let $\widehat{\boldsymbol{\sigma}}$ be the configuration obtained from $\boldsymbol{\sigma}$ by the swap of $\sigma_v^{(i)}$ and $\sigma_w^{(j)}$, namely

$$\widehat{\sigma}_x^{(k)} = \begin{cases} \sigma_w^{(j)}, & (k,x) = (i,v), \\ \sigma_v^{(i)}, & (k,x) = (j,w), \\ \sigma_x^{(k)}, & \text{otherwise}, \end{cases}$$

for all $(k,x) \in V_N$, and define

$$\boldsymbol{\sigma}^{i,v;j,w} := \begin{cases} \widehat{\boldsymbol{\sigma}}, & \widehat{\boldsymbol{\sigma}} \in \Omega_G^N, \\ \boldsymbol{\sigma}, & \widehat{\boldsymbol{\sigma}} \notin \Omega_G^N, \end{cases}$$

with the convention $\boldsymbol{\sigma}^{i,v;j,w} = \boldsymbol{\sigma}$ when $(i,v) = (j,w)$. For $i \neq j$ the move exchanges a spin of particle i with a spin of particle j, whereas for $i = j$ it is an internal swap within the single particle $\sigma^{(i)}$; in both cases the update is accepted only if the resulting configuration remains a valid independent set, i.e. $\widehat{\boldsymbol{\sigma}} \in \Omega_G^N$. The corresponding jump is

$$\boldsymbol{\sigma} \to \boldsymbol{\sigma}^{i,v;j,w}.$$

We consider two different evolutions: the mean-field dynamics, where u, v are arbitrary sites, and the single-site dynamics, where $u = v$.

Mean-field particle system

Here both pairs $i, j \in [N]$ and $v, w \in [n]$ are picked uniformly at random. Thus, the stochastic matrix of the Markov chain is given by

$$P_N^{\mathrm{mf}}(\boldsymbol{\sigma}, \boldsymbol{\eta}) := \frac{1}{N^2 n^2} \sum_{i,j \in [N]} \sum_{v,w \in V} \mathbf{1}\left\{\boldsymbol{\sigma}^{i,v;j,w} = \boldsymbol{\eta}\right\}.$$

The dynamics conserves the total occupation number

$$M(\boldsymbol{\sigma}) = \sum_{i=1}^{N} |\sigma^{(i)}|, \qquad |\sigma^{(i)}| = \sum_{v \in V} \sigma_v^{(i)}.$$

Hence, in order to obtain an irreducible chain, it is necessary to fix the number of occupied sites. Equivalently, we restrict to a state space with a fixed density.

Definition 3.1. Given a density $\alpha \in \left(0, \frac{1}{\Delta+1}\right]$,

$$\Omega_{N,\alpha}^{\mathrm{mf}} = \left\{\boldsymbol{\sigma} \in \Omega_G^N : M(\boldsymbol{\sigma}) = \lfloor Nn\alpha \rfloor\right\}.$$

We note that the restriction $\alpha \in (0, \frac{1}{\Delta+1}]$ is optimal, in the sense that it is the largest range ensuring that $\Omega_{N,\alpha}^{\mathrm{mf}}$ is nonempty for every graph of maximum degree Δ and every $N \geq 1$. For $\alpha \leq \frac{1}{\Delta+1}$, the initialization procedure in Section 3.2.1 constructs a configuration

$\boldsymbol{\sigma} \in \Omega^{\mathrm{mf}}_{N,\alpha}$; see Algorithm 4. Conversely, if $\alpha > \frac{1}{\Delta+1}$, then taking $G = K_{\Delta+1}$ makes $\Omega^{\mathrm{mf}}_{N,\alpha}$ empty for all sufficiently large N.

We define $\mu^{\mathrm{mf}}_{N,\alpha}$ as the uniform distribution over $\Omega^{\mathrm{mf}}_{N,\alpha}$,

$$\mu^{\mathrm{mf}}_{N,\alpha} = \mathrm{Unif}(\Omega^{\mathrm{mf}}_{N,\alpha})\,.$$

By symmetry of the stochastic matrix P^{mf}_N, the pair $(P^{\mathrm{mf}}_N, \mu^{\mathrm{mf}}_{N,\alpha})$ is reversible, for any $\alpha \in (0, \frac{1}{\Delta+1}]$. When the density is sufficiently small, we have the following quantitative convergence statement, expressed in terms of entropy contraction. Recall that $\alpha_c(\Delta)$ is the critical density introduced in (2.6).

Theorem 3.2. *Let $G = (V, E)$ be a graph on n vertices with maximum degree $\Delta \geq 3$. For any $\gamma > 0$ and $\delta \in (0, 1)$, consider a density $\alpha \in (\gamma, (1-\delta)\alpha_c(\Delta))$. Then, there exist constants $r := r(\gamma, \delta, \Delta) > 0$ and $N_0 := N_0(\gamma, \delta, \Delta)$ such that for any probability measure $\nu \in \mathcal{P}(\Omega^{\mathrm{mf}}_{N,\alpha})$ and any $N \geq N_0$, it holds that*

$$H(\nu P^{\mathrm{mf}}_N \mid \mu^{\mathrm{mf}}_{N,\alpha}) \leq \left(1 - \frac{r}{Nn}\right) H(\nu \mid \mu^{\mathrm{mf}}_{N,\alpha}). \tag{3.1}$$

In Section 3.2.1, we use this entropy contraction result to prove the $\widetilde{O}(Nn)$ mixing-time bound for our mean-field particle-system sampler.

Proof. The entropy contraction (3.1) can be derived from a logarithmic Sobolev inequality (LSI) for the pair $(P^{\mathrm{mf}}_N, \mu^{\mathrm{mf}}_{N,\alpha})$ established recently in [Jai+23; JPV26].

The LSI is formulated as follows. The Dirichlet form associated to the reversible pair $(P^{\mathrm{mf}}_N, \mu^{\mathrm{mf}}_{N,\alpha})$ is given by

$$\mathcal{E}_{N,\alpha}(F, G) = \frac{1}{2N^2n^2} \sum_{i,j=1}^{N} \sum_{v,w \in V} \mu^{\mathrm{mf}}_{N,\alpha}\Big[\nabla^{i,v;j,w} F \, \nabla^{i,v;j,w} G\Big]\,,$$

for all $F, G : \Omega^{\mathrm{mf}}_{N,\alpha} \to \mathbb{R}$, where the gradient operator $\nabla^{i,v;j,w}$ is defined by

$$\nabla^{i,v;j,w} F(\boldsymbol{\sigma}) := F(\boldsymbol{\sigma}^{i,v;j,w}) - F(\boldsymbol{\sigma}).$$

The LSI for the pair $(P^{\mathrm{mf}}_N, \mu^{\mathrm{mf}}_{N,\alpha})$ with constant $\varrho \geq 0$ is the inequality

$$\mathcal{E}_{N,\alpha}(\sqrt{F}, \sqrt{F}) \geq \varrho \, \mathrm{Ent}_{\mu^{\mathrm{mf}}_{N,\alpha}}[F], \tag{3.2}$$

for all $F : \Omega^{\mathrm{mf}}_{N,\alpha} \to \mathbb{R}_{\geq 0}$, where we use the notation $\mathrm{Ent}_{\mu^{\mathrm{mf}}_{N,\alpha}}[F] = \mu^{\mathrm{mf}}_{N,\alpha}[F \log\left(F/\mu^{\mathrm{mf}}_{N,\alpha}[F]\right)]$.

Set

$$\beta := \frac{\lfloor \alpha N n \rfloor}{Nn}.$$

Since $\alpha \in (\gamma, (1-\delta)\alpha_c(\Delta))$, by taking $N \geq N_0(\gamma, \delta, \Delta)$ we may ensure that

$$\beta \in \left(\frac{\gamma}{2}, (1-\delta)\alpha_c(\Delta)\right). \tag{3.3}$$

The inequality (3.2) is established with

$$\varrho = \frac{r_0}{Nn},$$

with a constant $r_0 = r_0(\gamma, \delta, \Delta) > 0$, by applying [Jai+23, Theorem 2] and [JPV26, Equation 6.3] in the density window (3.3), together with a simple comparison of the Dirichlet form $\mathcal{E}_{N,\alpha}(\cdot,\cdot)$ with the Dirichlet form of the down-up walk defined in [Jai+23].

Once the statement (3.2) is available, the desired contraction (3.1), with $r = r_0$, follows from [Cap23, Lemma 2.5] and [Mic97, Proposition 6]. We note that the argument in [Cap23, Lemma 2.5] requires the positive semidefiniteness of the stochastic matrix P_N^{mf}.

The latter can be checked by verifying the stronger property that the holding probability satisfies

$$P_N^{\mathrm{mf}}(\boldsymbol{\sigma}, \boldsymbol{\sigma}) \geq \frac{1}{2}, \qquad \forall\, \boldsymbol{\sigma} \in \Omega_{N,\alpha}^{\mathrm{mf}}.$$

To see this, note that $\beta \in (0,1)$. Since exactly $M(\boldsymbol{\sigma}) = \beta Nn$ sites are occupied and $Nn - M(\boldsymbol{\sigma}) = (1-\beta)Nn$ are empty, the probability that both selected sites $(i,v), (j,w)$ are occupied equals β^2, and the probability that both are empty equals $(1-\beta)^2$. Hence

$$P_N^{\mathrm{mf}}(\boldsymbol{\sigma}, \boldsymbol{\sigma}) \geq (1-\beta)^2 + \beta^2 = 1 - 2\beta(1-\beta) \geq \frac{1}{2},$$

where the last inequality uses $\beta(1-\beta) \leq \frac{1}{4}$, $\beta \in [0,1]$. □

Single-site particle system

Here, for $N \geq 2$, we pick $i, j \in [N]$ uniformly at random with the constraint $i \neq j$ and then we pick $v = w \in V$ uniformly at random. We use the notation, $\boldsymbol{\sigma}^{i,j;v} := \boldsymbol{\sigma}^{i,v;j,v}$, so that the moves of the chain are of the type

$$\boldsymbol{\sigma} \to \boldsymbol{\sigma}^{i,j;v}. \tag{3.4}$$

The stochastic matrix of the Markov chain is given by

$$P_N^{\mathrm{ss}}(\boldsymbol{\sigma}, \boldsymbol{\eta}) := \frac{1}{N(N-1)n} \sum_{\substack{i,j \in [N] \\ i \neq j}} \sum_{v \in V} \mathbf{1}\left\{\boldsymbol{\sigma}^{i,j;v} = \boldsymbol{\eta}\right\}. \tag{3.5}$$

As in the mean-field case, the dynamics is conservative. However, here there is a conserved quantity for every $v \in V$. For every $\boldsymbol{\sigma} \in \Omega_G^N$ and every $v \in V$, let $R_v(\boldsymbol{\sigma}) = \sum_{i \in [N]} \sigma_v^{(i)}$ be the number of ones in $\boldsymbol{\sigma}$ along the stripe $\{(i,v) : i \in [N]\}$. The vector

$$R(\boldsymbol{\sigma}) := (R_v(\boldsymbol{\sigma}))_{v \in V} \tag{3.6}$$

is conserved by the dynamics defined in (3.5), since all allowed exchanges are of the type (3.4). To obtain an irreducible chain one must therefore restrict the state space to a fixed value of (3.6).

Definition 3.3. For $\vec{r} = (r_v)_{v \in V} \in \{0, \dots, N\}^V$, define

$$\Omega^{\mathrm{ss}}_{N,\vec{r}} := \left\{ \boldsymbol{\sigma} \in \Omega_G^N : R_v(\boldsymbol{\sigma}) = r_v \text{ for every } v \in V \right\}.$$

We note that $\Omega^{\mathrm{ss}}_{N,\vec{r}}$ is nonempty whenever

$$(\Delta + 1) \max_{v \in V} r_v \leq N.$$

Under this condition, in Section 3.2.3, we provide an explicit algorithm for constructing an initial configuration $\boldsymbol{\sigma} \in \Omega^{\mathrm{ss}}_{N,\vec{r}}$ (see Algorithm 5).

By symmetry of the chain P^{ss}_N, whenever $\Omega^{\mathrm{ss}}_{N,\vec{r}} \neq \emptyset$, the uniform measure

$$\mu^{\mathrm{ss}}_{N,\vec{r}} := \mathrm{Unif}(\Omega^{\mathrm{ss}}_{N,\vec{r}})$$

is reversible with respect to P^{ss}_N. We now prove a rapid mixing bound for the chain P^{ss}_N by comparing it with Glauber dynamics for proper colorings of an auxiliary graph.

For a graph H and an integer $q \geq 1$, let $\mathcal{C}_q(H)$ denote the set of proper q-colorings of H. The Glauber dynamics on $\mathcal{C}_q(H)$ is the Markov chain which is defined as follows: at each step, it chooses a vertex $x \in V(H)$ uniformly at random and a color $j \in [q]$ uniformly at random, and recolors x with j if the resulting coloring is proper; otherwise it stays put. This chain is reversible with respect to the uniform distribution on $\mathcal{C}_q(H)$.

We will use the following hypothesis about mixing of colorings. It is widely conjectured that Glauber dynamics for proper q-colorings of a graph H mixes rapidly whenever $q \geq \Delta(H) + 2$. In this paper, we assume the following form of this conjecture: there exist absolute constants $c_{\mathrm{col}} > 1$ and $\Delta_0 > 0$ such that, for every graph H of maximum degree $\Delta(H) \geq \Delta_0$, the Glauber dynamics on $\mathcal{C}_q(H)$ mixes in

$$O\left(|V(H)| \log \frac{|V(H)|}{\zeta} \right)$$

steps within total variation distance ζ of the uniform distribution on $\mathcal{C}_q(H)$ whenever $q \geq c_{\mathrm{col}} \Delta(H)$. Note that we assume the constant hidden in big-O is absolute (determined by c_{col} and Δ_0) and independent of $\Delta(H)$ or q. The classical coupling argument of Jerrum proves such a bound for any fixed $c_{\mathrm{col}} > 2$ and $\Delta_0 = 0$; see [Jer95].

Let $\gamma \in (0, 1)$. Let $\vec{r} \in \mathbb{Z}^V_{\geq 0}$ satisfy the following assumption

$$\gamma \leq \frac{r_v}{N} \leq \frac{1}{c_{\mathrm{col}}(\Delta + 1)} \qquad \forall v \in V. \tag{3.7}$$

Note that by the discussion above, $\Omega^{\mathrm{ss}}_{N,\vec{r}} \neq \emptyset$. We associate to $\vec{r}$ an auxiliary graph $H_{\vec{r}}$ as follows. For each $v \in V$, let

$$C_v := (v, a) : a \in [r_v],$$

and define $H_{\vec{r}}$ to be the graph with vertex set

$$V(H_{\vec{r}}) = \bigcup_{v \in V} C_v,$$

where each C_v induces a clique, and for every edge $\{u, v\} \in E$, all vertices of C_u are adjacent to all vertices of C_v.

Let P_{col} denote the Glauber dynamics on $\mathcal{C}_N(H_{\vec{r}})$. We compare P_N^{ss} with the Glauber dynamics for proper N-colorings of $H_{\vec{r}}$.

For a coloring $c \in \mathcal{C}_N(H_{\vec{r}})$, define

$$\Psi_{\vec{r}}(c) := (\sigma^{(1)}, \dots, \sigma^{(N)}),$$

where

$$\sigma^{(i)} := v \in V : \text{there exists } a \in [r_v] \text{ such that } c(v, a) = i.$$

Lemma 3.4. *Let $H_{\vec{r}}$ and $\Psi_{\vec{r}}$ be the graph and the map defined above.*

First, for every $c \in \mathcal{C}_N(H_{\vec{r}})$, one has

$$\Psi_{\vec{r}}(c) \in \Omega^{\mathrm{ss}}_{N,\vec{r}}.$$

Second, for every $\boldsymbol{\sigma} \in \Omega^{\mathrm{ss}}_{N,\vec{r}}$,

$$\Psi_{\vec{r}}^{-1}(\boldsymbol{\sigma}) = \prod_{v \in V} r_v!.$$

Consequently, if c is sampled uniformly from $\mathcal{C}_N(H_{\vec{r}})$, then $\Psi_{\vec{r}}(c)$ is distributed according to $\mu^{\mathrm{ss}}_{N,\vec{r}}$.

Proof. Fix $c \in \mathcal{C}_N(H_{\vec{r}})$, and let

$$\Psi_{\vec{r}}(c) = (\sigma^{(1)}, \dots, \sigma^{(N)}).$$

Since C_v is a clique of size r_v, the coloring c uses r_v distinct colors on C_v. Therefore v belongs to exactly r_v of the sets $\sigma^{(1)}, \dots, \sigma^{(N)}$, and hence

$$\sum_{i=1}^{N} \sigma^{(i)} = \vec{r}.$$

It remains to show that each $\sigma^{(i)}$ is an independent set. Suppose that $u, v \in \sigma^{(i)}$ for some $i \in [N]$. Then there exist $a \in [r_u]$ and $b \in [r_v]$ such that $c(u, a) = c(v, b) = i$. Since c is proper, (u, a) and (v, b) are not adjacent in $H_{\vec{r}}$. In particular, one cannot have $\{u, v\} \in E$. Thus $\sigma^{(i)}$ is an independent set, and therefore $\Psi_{\vec{r}}(c) \in \Omega^{\mathrm{ss}}_{N,\vec{r}}$.

Now fix $\boldsymbol{\sigma} = (\sigma^{(1)}, \dots, \sigma^{(N)}) \in \Omega^{\mathrm{ss}}_{N,\vec{r}}$. For each $v \in V$, let

$$I_v := i \in [N] : v \in \sigma^{(i)}.$$

Note that $I_v = r_v$ for every $v \in V$. Moreover, if $\{u, v\} \in E$, then $I_u \cap I_v = \emptyset$, since each $\sigma^{(i)}$ is an independent set.

A coloring c satisfies $\Psi_{\vec{r}}(c) = \boldsymbol{\sigma}$ if and only if, for every $v \in V$, the colors used on C_v are exactly the elements of I_v. Since C_v is a clique, this is equivalent to choosing a bijection from C_v onto I_v. The number of such choices is $r_v!$, independently for each $v \in V$. Hence

$$\Psi_{\vec{r}}^{-1}(\boldsymbol{\sigma}) = \prod_{v \in V} r_v!.$$

The final claim follows immediately, since every element of $\Omega^{\mathrm{ss}}_{N,\vec{r}}$ has the same number of preimages under $\Psi_{\vec{r}}$. □

The following theorem proves the chain P^{ss}_N mixes rapidly by comparing it with P_{col}.

Theorem 3.5. *Let $\gamma \in (0,1)$, and let $c_{\mathrm{col}} > 1$ and $\Delta_0 > 0$ be the coloring-mixing constants from Section 3.1. Let $G = (V, E)$ be a graph on n vertices with maximum degree $\Delta \geq \Delta_0$. Then there exists $C_0 = C_0(\Delta, c_{\mathrm{col}})$ such that, for every $N \geq C_0$, if $\vec{r} \in \mathbb{Z}^V_{\geq 0}$ satisfies (3.7), then the particle system chain P^{ss}_N mixes within total variation distance ζ of $\mu^{\mathrm{ss}}_{N,\vec{r}}$ in*

$$O_{\gamma,\Delta,c_{\mathrm{col}}}\left(Nn \log \frac{Nn}{\zeta}\right)$$

steps.

Proof. For every $(v, a) \in V(H_{\vec{r}})$,

$$\deg_{H_{\vec{r}}}(v, a) = (r_v - 1) + \sum_{u \in v} r_u \leq \sum_{u \in v} r_u.$$

Now we have

$$\deg_{H_{\vec{r}}}(v, a) \leq (\Delta + 1) \max_{u \in V} r_u \leq \frac{N}{c_{\mathrm{col}}},$$

where the last inequality follows from (3.7). Hence

$$N \geq c_{\mathrm{col}} \Delta(H_{\vec{r}}).$$

Also, $\Delta(H_{\vec{r}}) \geq \Delta \geq \Delta_0$. Therefore, by Section 3.1, P_{col} mixes within total variation distance ζ of the uniform distribution on $\mathcal{C}_N(H_{\vec{r}})$ in

$$O_{c_{\mathrm{col}},\Delta_0}\left(|V(H_{\vec{r}})| \log \frac{|V(H_{\vec{r}})|}{\zeta}\right)$$

steps.

We now compare P^{ss}_N with P_{col}. Let

$$m_{\vec{r}} := |V(H_{\vec{r}})| = \sum_{v \in V} r_v.$$

Define an auxiliary coloring chain $\widehat{P_{\mathrm{col}}}$ on $\mathcal{C}_N(H_{\vec{r}})$ as follows. At each step, choose $v \in V$ uniformly at random and choose an ordered pair of distinct colors $i, j \in [N]$ uniformly at random. If exactly one of the colors i, j appears on the clique C_v, recolor the unique vertex of C_v carrying the present color with the absent color, provided the resulting coloring is proper; otherwise stay put. If either both colors appear on C_v or neither appears on C_v, stay put.

We claim that the projection of $\widehat{P_{\mathrm{col}}}$ under $\Psi_{\vec{r}}$ is exactly P_N^{ss}. Suppose that $\Psi_{\vec{r}}(c) = \boldsymbol{\sigma}$. For each $v \in V$, the colors used on C_v are exactly the indices $k \in [N]$ such that $v \in \sigma^{(k)}$. Thus, choosing v, i, j in the coloring chain corresponds to choosing the same v, i, j in the particle system chain. If exactly one of i, j appears on C_v, then exactly one of the particles $\sigma^{(i)}$ and $\sigma^{(j)}$ contains v. Recoloring on C_v replaces the present color by the absent color, which is the same as swapping the site v between $\sigma^{(i)}$ and $\sigma^{(j)}$. The recoloring is proper if and only if the resulting two particle configurations are independent sets. Therefore the projected update is exactly the update of P_N^{ss}.

It remains to compare $\widehat{P_{\mathrm{col}}}$ with P_{col}. Consider a nontrivial recoloring move $c \to c'$ in which a vertex of some clique C_v with current color i is recolored with color j. Under the standard Glauber chain,

$$P_{\mathrm{col}}(c, c') = \frac{1}{m_{\vec{r}} N}.$$

Under $\widehat{P_{\mathrm{col}}}$, the same move is proposed by the two ordered choices (i, j) and (j, i), and hence

$$\widehat{P_{\mathrm{col}}}(c, c') = \frac{2}{nN(N-1)}.$$

Define $\beta := \frac{2m_{\vec{r}}}{n(N-1)}$. By (3.7), we have

$$\frac{2m_{\vec{r}}}{n(N-1)} \le \frac{2 \max_{v \in V} r_v}{N-1} \le \frac{N}{c_{\mathrm{col}}(\Delta+1)(N-1)}.$$

Thus if we choose $C_0 = C_0(\Delta, c_{\mathrm{col}})$ sufficiently large, for $N \ge C_0$,

$$\beta \le 1.$$

Now, note that for off-diagonal transition,

$$\widehat{P_{\mathrm{col}}}(c, c') = \beta \, P_{\mathrm{col}}(c, c').$$

Thus

$$\widehat{P_{\mathrm{col}}} = (1-\beta)I + \beta P_{\mathrm{col}}.$$

The lower bound $\frac{r_v}{N} \ge \gamma$ also gives

$$m_{\vec{r}} = \sum_{v \in V} r_v \ge Nn\gamma.$$

Thus $\beta = \Theta_{\gamma,\Delta,c_{\mathrm{col}}}(1)$. Consequently, $\widehat{P_{\mathrm{col}}}$ is a constant-factor lazy version of P_{col}, and hence it mixes within total variation distance ζ of the uniform distribution on $\mathcal{C}_N(H_{\vec{r}})$ in

$$O_{\gamma,\Delta,c_{\mathrm{col}}}\left(m_{\vec{r}} \log \frac{m_{\vec{r}}}{\zeta}\right)$$

steps.

Since P_N^{ss} is the projection of $\widehat{P_{\mathrm{col}}}$ under $\Psi_{\vec{r}}$, total variation distance cannot increase under this projection. Moreover, by Theorem 3.4, the uniform distribution on $\mathcal{C}_N(H_{\vec{r}})$ is mapped by $\Psi_{\vec{r}}$ to $\mu_{N,\vec{r}}^{\mathrm{ss}}$. Therefore P_N^{ss} has stationary distribution $\mu_{N,\vec{r}}^{\mathrm{ss}}$ and mixes within total variation distance ζ of $\mu_{N,\vec{r}}^{\mathrm{ss}}$ in

$$O_{\gamma,\Delta,c_{\mathrm{col}}}\left(m_{\vec{r}} \log \frac{m_{\vec{r}}}{\zeta}\right)$$

steps. Finally, since

$$m_{\vec{r}} = |V(H_{\vec{r}})| = \sum_{v \in V} r_v = O_{\gamma,\Delta,c_{\mathrm{col}}}(Nn),$$

the claimed bound follows. □

In Section 3.2.3, we apply Theorem 3.5 to analyze the mixing time of our single-site particle-system sampler.

3.2 Algorithmic Applications

We now turn to the algorithmic problem of sampling from the hard-core model when the desired conserved quantities are prescribed in advance. There are two natural versions of this problem. In the mean-field setting, a target density α is given, and we want to sample from the unique uniform-fugacity Gibbs measure with density α. In the single-site setting, a target marginal vector $(m_v)_{v \in V}$ is given, and we want to sample from the unique hard-core Gibbs measure whose vertex marginals are given by $(m_v)_{v \in V}$.

Remark 3.6. Drawing one sample from a distribution with prescribed vertex marginals can be done efficiently. Given a graph $G = (V, E)$ and a marginal vector $\boldsymbol{m} = (m_v)_{v \in V}$ satisfying

$$\max_{v \in V} m_v \leq \frac{1}{\Delta + 1},$$

the following algorithm outputs an independent set $\sigma \in \Omega_G$ with

$$\Pr[\sigma_v = 1] = m_v \qquad \forall v \in V.$$

Each color class is an independent set, so the output belongs to Ω_G. Moreover, for every $v \in V$,

$$\Pr[\sigma_v = 1] = \frac{1}{\Delta + 1} \cdot (\Delta + 1) m_v = m_v.$$

Algorithm 3 Initialization with prescribed marginals

Require: Graph $G = (V, E)$ with maximum degree Δ, and marginals $\boldsymbol{m} = (m_v)_{v \in V}$ satisfying $\max_{v \in V} m_v \leq 1/(\Delta + 1)$.

Ensure: A random independent set $\sigma \in \Omega_G$ with $\Pr[\sigma_v = 1] = m_v$ for all $v \in V$.

1: Find a proper coloring $c : V \to [\Delta + 1]$.
2: Choose a color $i \in [\Delta + 1]$ uniformly at random.
3: Initialize $\sigma_v \leftarrow 0$ for all $v \in V$.
4: **for all** $v \in c^{-1}(i)$ **do**
5: Set $\sigma_v \leftarrow 1$ with probability $(\Delta + 1)m_v$, independently.
6: **end for**
7: **return** σ.

For prescribed density α, the same algorithm applied to the constant marginal vector $m_v = \alpha$ gives an initial distribution with density α, provided $\alpha \leq 1/(\Delta + 1)$.

In order to sample from the maximum entropy distributions with the prescribed marginals or density, one possible approach is to use the nonlinear dynamics directly. First, we choose an initial distribution with the prescribed conserved quantity as in Theorem 3.6: density α in the mean-field case, or vertex marginals $(m_v)_{v \in V}$ in the single-site case. Then we run the corresponding nonlinear dynamics until it is close to its limiting Gibbs measure. The issue is that this dynamics is not easy to implement efficiently. For example, one can use the derivation tree construction of [CS24b]: to sample from the distribution obtained after T nonlinear steps, draw 2^T independent samples from an initial distribution $\mu_0 \in \mathcal{P}(\Omega_G)$ with the prescribed conserved quantity, place them at the leaves of a binary tree of height T, and apply the exchange kernel level by level from the leaves to the root. The configuration at the root has the law obtained after T steps of the nonlinear dynamics. However, this method already requires 2^T initial samples, and is therefore not efficient.

For this reason, we do not use the nonlinear dynamics itself as the algorithm. Instead, we use the particle system introduced in Section 3.1. In the prescribed-density case, we work on the fixed total-occupation space $\Omega^{\mathrm{mf}}_{N,\alpha}$ and use the mean-field particle chain P^{mf}_N. In the prescribed-marginal case, we work on the fixed vertex-count space $\Omega^{\mathrm{ss}}_{N,\vec{r}}$ and use the single-site particle chain P^{ss}_N. In both cases, we run the corresponding particle-system chain close to stationarity and output one particle.

3.2.1 Sampling with prescribed density

We now consider the mean-field inverse sampling problem, where the prescribed conserved quantity is the density α. Fix $\gamma, \delta \in (0, 1)$. Throughout this subsection we assume that

$$\gamma \leq \alpha \leq (1 - \delta)\alpha_c(\Delta). \tag{3.8}$$

One can show that there exists a unique fugacity $\lambda_\alpha > 0$ for which the density of the

uniform fugacity hard-core Gibbs measure μ_{G,λ_α} is exactly α. Further, λ_α lies in the regime

$$\gamma \le \lambda_\alpha \le (1-\delta)\lambda_c(\Delta).$$

See the beginning of Appendix B.1.2 for the detailed proof.
The goal is to approximately sample from μ_{G,λ_α}. For simplicity, throughout this subsection we write

$$\mu_G^\alpha := \mu_{G,\lambda_\alpha}.$$

For sufficiently large N, we write

$$M = Nn\alpha,$$

and we consider the space $\Omega_{N,\alpha}^{\mathrm{mf}}$ with uniform measure $\mu_{N,\alpha}^{\mathrm{mf}}$. We first observe that, if $\boldsymbol{\sigma}$ is a uniform draw from $\mu_{N,\alpha}^{\mathrm{mf}}$, then the law of $\sigma^{(1)}$ is a good approximation for μ_G^α. After that, we show that the mean-field particle system chain P_N^{mf} mixes rapidly to $\mu_{N,\alpha}^{\mathrm{mf}}$.

The next lemma proves the first step.

Lemma 3.7. *Assume $\Delta \ge 3$, and α satisfies (3.8). Then there exist constants $C = C(\gamma,\delta,\Delta) > 0$ and $N_0 = N_0(\gamma,\delta,\Delta)$ such that, for every $N \ge N_0$,*

$$\left\| P_1 \mu_{N,\alpha}^{\mathrm{mf}} - \mu_G^\alpha \right\|_{\mathrm{TV}} \le \frac{C}{N}.$$

Proof. We write $\mu := \mu_G^\alpha$. By the bounds on λ_α stated at the beginning of this subsection,

$$\gamma \le \lambda_\alpha \le (1-\delta)\lambda_c(\Delta).$$

All constants below depend only on γ, δ, and Δ. Note that

$$\mu_{N,\alpha}^{\mathrm{mf}} = \mu^{\otimes N}(\,\cdot \mid \Omega_{N,\alpha}^{\mathrm{mf}}).$$

By duality of total variation distance, it suffices to show that for every $f : \Omega_G \to \mathbb{R}$,

$$\left| \mathbb{E}_{\boldsymbol{\sigma}\sim\mu_{N,\alpha}^{\mathrm{mf}}}\left[f(\sigma^{(1)})\right] - \mathbb{E}_{\sigma\sim\mu}[f(\sigma)] \right| \le \frac{C}{N}\|f\|_\infty.$$

Let

$$X_N := M(\boldsymbol{\sigma}) = \sum_{i=1}^N |\sigma^{(i)}| \qquad \text{when } \boldsymbol{\sigma} \sim \mu^{\otimes N}.$$

Set

$$a_N := \mathbb{E}_{\boldsymbol{\sigma}\sim\mu^{\otimes N}}[X_N] = Nn\alpha, \qquad v_N^2 := \mathrm{Var}_{\boldsymbol{\sigma}\sim\mu^{\otimes N}}(X_N).$$

Since $\Omega_{N,\alpha}^{\mathrm{mf}} = \{\boldsymbol{\sigma} : M(\boldsymbol{\sigma}) = M\} = \{X_N = M\}$, we have

$$\mathbb{E}_{\boldsymbol{\sigma}\sim\mu_{N,\alpha}^{\mathrm{mf}}}\left[f(\sigma^{(1)})\right] - \mathbb{E}_{\sigma\sim\mu}[f(\sigma)] = \frac{\mathcal{N}(f)}{\mu^{\otimes N}(X_N = M)}, \tag{3.9}$$

where

$$\mathcal{N}(f) := \mathbb{E}_{\boldsymbol{\sigma}\sim\mu^{\otimes N}}\Big[f(\sigma^{(1)})\mathbf{1}\{X_N = M\}\Big] - \mathbb{E}_{\sigma\sim\mu}[f(\sigma)]\,\mu^{\otimes N}(X_N = M).$$

We estimate the denominator and the numerator in (3.9) separately.

Denominator. Let

$$Y := |\sigma| \qquad \text{when } \sigma \sim \mu,$$

and set

$$v_1^2 := \mathrm{Var}_{\sigma\sim\mu}(Y).$$

Since

$$(1-\delta)\lambda_c(\Delta) = \lambda_c(\Delta) - \delta\lambda_c(\Delta),$$

the variance estimate [Jai+22, Lemma 3.2] applies with additive gap $\delta\lambda_c(\Delta)$ from the threshold. Since also $\lambda_\alpha \geq \gamma$, there exist constants $c_0, C_0 > 0$, depending only on γ, δ, and Δ, such that

$$c_0 n \leq v_1^2 \leq C_0 n.$$

Equivalently, since $v_N^2 = N v_1^2$,

$$c_0 N n \leq v_N^2 \leq C_0 N n.$$

Recall that G_N is the disjoint union of N copies of G. Thus G_N has Nn vertices and maximum degree at most Δ. Moreover, the hard-core Gibbs measure on G_N with fugacity λ_α is exactly $\mu^{\otimes N}$, and X_N is the size of the independent set sampled from this measure. Therefore, by the local central limit theorem [Jai+23, Corollary 13], applied to G_N, we have

$$\left|\mu^{\otimes N}(X_N = M) - \frac{1}{\sqrt{2\pi}\,v_N}\exp\left(-\frac{(M-a_N)^2}{2v_N^2}\right)\right| \leq C_1(Nn)^{-3/2},$$

for some constant $C_1 > 0$.

Since $M = Nn\alpha$ and $a_N = Nn\alpha$, we have

$$|M - a_N| \leq 1.$$

We first lower bound the Gaussian leading term. Using $v_N^2 \leq C_0 N n$, we get

$$\frac{1}{v_N} \geq \frac{1}{\sqrt{C_0 N n}}.$$

Using $v_N^2 \geq c_0 N n$ and $|M - a_N| \leq 1$, we also get

$$\frac{(M-a_N)^2}{2v_N^2} \leq \frac{1}{2c_0 N n}.$$

Hence, for all sufficiently large N,

$$\exp\left(-\frac{(M-a_N)^2}{2v_N^2}\right) \geq \frac{1}{2}.$$

Therefore,

$$\frac{1}{\sqrt{2\pi}\,v_N}\exp\left(-\frac{(M-a_N)^2}{2v_N^2}\right) \geq \frac{1}{2\sqrt{2\pi C_0}}\frac{1}{\sqrt{Nn}}.$$

On the other hand,

$$C_1(Nn)^{-3/2} = \frac{C_1}{Nn}\frac{1}{\sqrt{Nn}},$$

and hence, for all sufficiently large N,

$$C_1(Nn)^{-3/2} \leq \frac{1}{4\sqrt{2\pi C_0}}\frac{1}{\sqrt{Nn}}.$$

Combining this with the local central limit theorem gives

$$\begin{aligned}\mu^{\otimes N}(X_N = M) &\geq \frac{1}{\sqrt{2\pi}\,v_N}\exp\left(-\frac{(M-a_N)^2}{2v_N^2}\right) - C_1(Nn)^{-3/2} \\ &\geq \frac{1}{2\sqrt{2\pi C_0}}\frac{1}{\sqrt{Nn}} - \frac{1}{4\sqrt{2\pi C_0}}\frac{1}{\sqrt{Nn}} \\ &= \frac{1}{4\sqrt{2\pi C_0}}\frac{1}{\sqrt{Nn}}.\end{aligned}$$

Thus, after setting

$$c_1 := \frac{1}{4\sqrt{2\pi C_0}},$$

we obtain

$$\mu^{\otimes N}(X_N = M) \geq \frac{c_1}{\sqrt{Nn}}. \tag{3.10}$$

Numerator. Define the normalized characteristic function

$$\psi(t) := \mathbb{E}_{\sigma\sim\mu}\left[\exp\left(\frac{it(Y-n\alpha)}{v_1}\right)\right].$$

Using again the additive gap $\delta\lambda_c(\Delta)$ from the threshold, the characteristic function estimate [Jai+22, Lemma 3.5], applied with $\lambda = \lambda_\alpha$, gives a constant $c > 0$, depending only on δ and Δ, such that

$$|\psi(t)| \leq \exp\left(-\frac{c\lambda_\alpha n t^2}{v_1^2}\right) \qquad \forall\, t \in [-\pi v_1, \pi v_1].$$

Since $\lambda_\alpha \geq \gamma$, setting $c_2 := c\gamma$ gives

$$|\psi(t)| \leq \exp\left(-\frac{c_2 n t^2}{v_1^2}\right) \qquad \forall\, t \in [-\pi v_1, \pi v_1].$$

We shall estimate the numerator $\mathcal{N}(f)$ with f replaced by a recentered function $\tilde{f}$, and then control the recentering term separately.

To this end, note that for any function $h : \Omega_G \to \mathbb{R}$, Fourier inversion gives

$$\mathcal{N}(h) = \frac{1}{2\pi v_N} \int_{-\pi v_N}^{\pi v_N} e^{-it(a_N - M)/v_N} \psi(t/\sqrt{N})^{N-1} \operatorname{Cov}_\mu \left(h, \exp\left(\frac{it(Y - n\alpha)}{v_N} \right) \right) dt. \quad (3.11)$$

For $t \in [-\pi v_N, \pi v_N]$, we have $t/\sqrt{N} \in [-\pi v_1, \pi v_1]$. Hence, for all sufficiently large N,

$$|\psi(t/\sqrt{N})|^{N-1} \le \exp\left(-\frac{c_2 n t^2}{2v_1^2} \right).$$

We now recenter f by removing its component in the direction of the conserved quantity. Set

$$\beta_f := \frac{\operatorname{Cov}_\mu(f, Y)}{v_1^2}, \qquad g(\sigma) := \beta_f \left(|\sigma| - n\alpha \right), \qquad \widetilde{f} := f - g.$$

Then

$$\operatorname{Cov}_\mu(\widetilde{f}, Y) = 0, \qquad \mathbb{E}_\mu[g] = 0.$$

Moreover, by Cauchy–Schwarz,

$$|\beta_f| \le \frac{\sqrt{\operatorname{Var}_\mu(f)}}{v_1} \le \frac{\|f\|_\infty}{v_1}.$$

Therefore

$$\operatorname{Var}_\mu(g) = \beta_f^2 v_1^2 \le \|f\|_\infty^2,$$

and hence

$$\operatorname{Var}_\mu(\widetilde{f}) \le 2 \operatorname{Var}_\mu(f) + 2 \operatorname{Var}_\mu(g) \le C \|f\|_\infty^2.$$

By the cumulant bound of [Jai+23, Lemma 11], for every fixed $d \ge 1$,

$$|\kappa_d(Y)| \le O_{\Delta,\delta,d}(n), \quad \kappa_d(Y) = \frac{d^d}{dt^d} \log \mathbb{E}_\mu \left[e^{tY} \right] \Big|_{t=0}.$$

In particular, taking $d = 4$, there exists a constant C such that

$$|\kappa_4(Y)| \le Cn. \quad (3.12)$$

Recall also that $v_1^2 = \operatorname{Var}_\mu(Y) \le C_0 n$ for some constant C_0. Now, note that

$$\mathbb{E}_\mu \left[(Y - n\alpha)^4 \right] = \kappa_4(Y) + 3 \operatorname{Var}_\mu(Y)^2 \le |\kappa_4(Y)| + 3 \operatorname{Var}_\mu(Y)^2.$$

Thus, by applying (3.12) and the variance bound, there exists a constant C_4 such that

$$\mathbb{E}_\mu \left[(Y - n\alpha)^4 \right] \le C_4 n^2. \quad (3.13)$$

Since $\mathrm{Cov}_\mu(\widetilde{f}, Y) = 0$, we have

$$\mathrm{Cov}_\mu\left(\widetilde{f}, \exp\left(\frac{it(Y-n\alpha)}{v_N}\right)\right) = \mathrm{Cov}_\mu\left(\widetilde{f}, \exp\left(\frac{it(Y-n\alpha)}{v_N}\right) - 1 - \frac{it(Y-n\alpha)}{v_N}\right).$$

Using

$$|e^{i\theta} - 1 - i\theta| \le \frac{\theta^2}{2},$$

Cauchy–Schwarz, and (3.13), we get

$$\begin{aligned}
&\left|\mathrm{Cov}_\mu\left(\widetilde{f}, \exp\left(\frac{it(Y-n\alpha)}{v_N}\right)\right)\right| \\
&\quad\le \sqrt{\mathrm{Var}_\mu(\widetilde{f})}\left(\mathbb{E}_\mu\left[\left|\exp\left(\frac{it(Y-n\alpha)}{v_N}\right) - 1 - \frac{it(Y-n\alpha)}{v_N}\right|^2\right]\right)^{1/2} \\
&\quad\le C\|f\|_\infty \frac{t^2}{v_N^2}\left(\mathbb{E}_\mu\left[(Y-n\alpha)^4\right]\right)^{1/2} \\
&\quad\le C\|f\|_\infty \frac{t^2}{N}.
\end{aligned} \tag{3.14}$$

In the last step we used $v_N^2 = Nv_1^2$ and $v_1^2 \ge c_0 n$. The value of the constant C may change from line to line but it remains independent of N, n.

Combining (3.11) and (3.14), we obtain

$$\begin{aligned}
|\mathcal{N}(\widetilde{f})| &\le \frac{1}{2\pi v_N}\int_{-\pi v_N}^{\pi v_N} \exp\left(-\frac{c_2 n t^2}{2v_1^2}\right) C\|f\|_\infty \frac{t^2}{N}\, dt \\
&\le \frac{C\|f\|_\infty}{Nv_N} \\
&\le \frac{C\|f\|_\infty}{N\sqrt{Nn}}.
\end{aligned} \tag{3.15}$$

Combining (3.9), (3.10), and (3.15), gives

$$\left|\mathbb{E}_{\boldsymbol{\sigma}\sim\mu_{N,\alpha}^{\mathrm{mf}}}\left[\widetilde{f}(\sigma^{(1)})\right] - \mathbb{E}_{\sigma\sim\mu}\left[\widetilde{f}(\sigma)\right]\right| \le \frac{C\|f\|_\infty}{N}.$$

It remains to estimate the linear correction g. Since $\mathbb{E}_\mu[g] = 0$ and, by symmetry under $\mu_{N,\alpha}^{\mathrm{mf}}$,

$$\mathbb{E}_{\boldsymbol{\sigma}\sim\mu_{N,\alpha}^{\mathrm{mf}}}\left[|\sigma^{(1)}|\right] = \frac{M}{N},$$

we have

$$\begin{aligned}
\left|\mathbb{E}_{\boldsymbol{\sigma}\sim\mu_{N,\alpha}^{\mathrm{mf}}}\left[g(\sigma^{(1)})\right] - \mathbb{E}_{\sigma\sim\mu}[g(\sigma)]\right| &= |\beta_f|\frac{|M - Nn\alpha|}{N} \\
&\le \frac{\|f\|_\infty}{v_1}\frac{1}{N} \\
&\le \frac{C\|f\|_\infty}{N\sqrt{n}}.
\end{aligned}$$

Here we used $|M - Nn\alpha| \le 1$ and $v_1^2 \ge c_0 n$.

Since $f = \widetilde{f} + g$, we conclude that

$$\left|\mathbb{E}_{\boldsymbol{\sigma}\sim\mu_{N,\alpha}^{\mathrm{mf}}}\left[f(\sigma^{(1)})\right] - \mathbb{E}_{\sigma\sim\mu}\left[f(\sigma)\right]\right| \le \frac{C\|f\|_\infty}{N}.$$

Taking the supremum over $\|f\|_\infty \le 1$ gives

$$\left\|P_1\mu_{N,\alpha}^{\mathrm{mf}} - \mu_G^{\alpha}\right\|_{\mathrm{TV}} \le \frac{C}{N}.$$

□

It remains to show that we can efficiently sample from $\mu_{N,\alpha}^{\mathrm{mf}}$ by running the chain P_N^{mf}. Recall that the pair $(P_N^{\mathrm{mf}}, \mu_{N,\alpha}^{\mathrm{mf}})$ is reversible. The following lemma is a direct implication of Theorem 3.2.

Lemma 3.8. *Assume $\Delta \ge 3$, and suppose α satisfies (3.8). Then, there exists a constant $N_0 := N_0(\gamma, \delta, \Delta)$ such that for every $\zeta \in (0,1)$ and every $N \ge N_0$, the chain P_N^{mf} mixes within total variation distance ζ of $\mu_{N,\alpha}^{\mathrm{mf}}$ in*

$$O_{\gamma,\delta,\Delta}\left(Nn \log \frac{Nn}{\zeta}\right)$$

steps.

Proof. By Theorem 3.2, applied with parameters $\gamma/2$ and $\delta/2$, there exist constants $r = r(\gamma, \delta, \Delta) > 0$ and $N_0 = N_0(\gamma, \delta, \Delta)$ such that, for every $N \ge N_0$ and every probability measure ν on $\Omega_{N,\alpha}^{\mathrm{mf}}$,

$$H(\nu P_N^{\mathrm{mf}} \mid \mu_{N,\alpha}^{\mathrm{mf}}) \le \left(1 - \frac{r}{Nn}\right) H(\nu \mid \mu_{N,\alpha}^{\mathrm{mf}}).$$

Iterating the entropy contraction gives

$$H(\nu (P_N^{\mathrm{mf}})^t \mid \mu_{N,\alpha}^{\mathrm{mf}}) \le \exp\left(-\frac{rt}{Nn}\right) H(\nu \mid \mu_{N,\alpha}^{\mathrm{mf}}).$$

Since $\mu_{N,\alpha}^{\mathrm{mf}}$ is uniform on $\Omega_{N,\alpha}^{\mathrm{mf}}$, for every initial distribution ν we have

$$H(\nu \mid \mu_{N,\alpha}^{\mathrm{mf}}) \le \log|\Omega_{N,\alpha}^{\mathrm{mf}}| \le Nn \log 2.$$

Thus

$$H(\nu (P_N^{\mathrm{mf}})^t \mid \mu_{N,\alpha}^{\mathrm{mf}}) \le Nn \log 2 \exp\left(-\frac{rt}{Nn}\right).$$

By Pinsker's inequality,

$$\left\|\nu (P_N^{\mathrm{mf}})^t - \mu_{N,\alpha}^{\mathrm{mf}}\right\|_{\mathrm{TV}} \le \sqrt{\frac{Nn \log 2}{2} \exp\left(-\frac{rt}{Nn}\right)}.$$

Choosing the implicit constant sufficiently large in

$$t = O_{\gamma,\delta,\Delta}\left(Nn \log \frac{Nn}{\zeta}\right)$$

makes the right-hand side at most ζ. This proves the claimed mixing bound. □

We now describe the resulting algorithm. By Theorem 3.7, if $\boldsymbol{\sigma}$ is sampled from $\mu^{\mathrm{mf}}_{N,\alpha}$, then the first particle $\sigma^{(1)}$ has law close to μ^{α}_{G}. Thus it remains to approximately sample from $\mu^{\mathrm{mf}}_{N,\alpha}$. For this, we run the chain P^{mf}_{N}, which mixes rapidly to $\mu^{\mathrm{mf}}_{N,\alpha}$ by Theorem 3.8, and then output the first particle.

We first explain how to initialize the chain. Since

$$\alpha \le (1-\delta)\alpha_c(\Delta) < \frac{1}{\Delta+1},$$

we have

$$M = Nn\alpha < \frac{Nn}{\Delta+1}.$$

Find a proper $(\Delta+1)$-coloring of G by the greedy algorithm, and use the same coloring on each of the N copies of G in G_N. Then some color class of G_N has size at least $Nn/(\Delta+1)$. Thus, choosing any M vertices of this color class gives an independent set of G_N of size M, or equivalently an initial configuration in $\Omega^{\mathrm{mf}}_{N,\alpha}$.

Algorithm 4 Prescribed-Density Particle System Sampler

Require: Graph $G = (V, E)$ with maximum degree Δ, target density α satisfying the assumptions above, and accuracy $\varepsilon \in (0, 1)$.

Ensure: A random independent set whose law is within total variation distance ε from μ^{α}_{G}.

1: Choose $N = O_{\gamma,\delta,\Delta}(\frac{1}{\varepsilon})$ sufficiently large.
2: Set $M := Nn\alpha$.
3: Find a proper $(\Delta+1)$-coloring of G by the greedy algorithm, and use the same coloring on each of the N copies of G in G_N.
4: Choose a color class of G_N of size at least $Nn/(\Delta+1)$, and choose any M vertices from this color class.
5: Let $\boldsymbol{\sigma}_0 \in \Omega^{\mathrm{mf}}_{N,\alpha}$ be the corresponding initial particle configuration.
6: Run the particle system chain P^{mf}_{N} from $\boldsymbol{\sigma}_0$ for

$$O_{\gamma,\delta,\Delta}\left(Nn \log \frac{Nn}{\varepsilon}\right)$$

steps.
7: **return** $\sigma^{(1)}$, where $\boldsymbol{\sigma} = (\sigma^{(1)}, \dots, \sigma^{(N)})$ is the resulting particle configuration.

Theorem 3.9. *Let $\gamma, \delta \in (0, 1)$. Let $G = (V, E)$ be a graph on n vertices with maximum degree $\Delta \ge 3$, and let the target density α satisfy*

$$\gamma \le \alpha \le (1-\delta)\alpha_c(\Delta).$$

Then, for every $\varepsilon \in (0,1)$, Algorithm 4 outputs a random independent set whose law is within total variation distance ε from μ_G^α. Moreover, the running time is

$$\widetilde{O}_{\gamma,\delta,\Delta}\left(\frac{n}{\varepsilon}\right).$$

Proof. Let C and N_1 be the constants from Theorem 3.7, and let N_2 be the constant from Theorem 3.8. We choose N sufficiently large so that

$$N \geq \max\left\{N_1, N_2, \frac{2C}{\varepsilon}\right\}.$$

In particular,

$$N = O_{\gamma,\delta,\Delta}(\frac{1}{\varepsilon}).$$

Let $\widehat{\mu}$ be the law of the particle configuration after running P_N^{mf} for the number of steps specified in Algorithm 4. By Theorem 3.8, choosing the implicit constant in the running time sufficiently large gives

$$\left\|\widehat{\mu} - \mu_{N,\alpha}^{\mathrm{mf}}\right\|_{\mathrm{TV}} \leq \frac{\varepsilon}{2}.$$

Let ν denote the law of the output $\sigma^{(1)}$. Since total variation distance cannot increase under projection to the first particle,

$$\left\|\nu - P_1\mu_{N,\alpha}^{\mathrm{mf}}\right\|_{\mathrm{TV}} = \left\|P_1\widehat{\mu} - P_1\mu_{N,\alpha}^{\mathrm{mf}}\right\|_{\mathrm{TV}} \leq \frac{\varepsilon}{2}.$$

On the other hand, by Theorem 3.7,

$$\left\|P_1\mu_{N,\alpha}^{\mathrm{mf}} - \mu_G^\alpha\right\|_{\mathrm{TV}} \leq \frac{C}{N} \leq \frac{\varepsilon}{2}.$$

Therefore,

$$\|\nu - \mu_G^\alpha\|_{\mathrm{TV}} \leq \left\|\nu - P_1\mu_{N,\alpha}^{\mathrm{mf}}\right\|_{\mathrm{TV}} + \left\|P_1\mu_{N,\alpha}^{\mathrm{mf}} - \mu_G^\alpha\right\|_{\mathrm{TV}} \leq \varepsilon.$$

It remains to bound the running time. The initialization takes $O_\Delta(Nn)$ time. The chain is run for

$$O_{\gamma,\delta,\Delta}\left(Nn\log\frac{Nn}{\varepsilon}\right)$$

steps, and each step can be implemented in $O_\Delta(1)$ time. Since $N = O_{\gamma,\delta,\Delta}(\frac{1}{\varepsilon})$, the total running time is

$$O_{\gamma,\delta,\Delta}\left(\frac{n}{\varepsilon}\log\frac{n}{\varepsilon}\right) = \widetilde{O}_{\gamma,\delta,\Delta}\left(\frac{n}{\varepsilon}\right).$$

This completes the proof. □

Remark 3.10. The same approach can be used to generate k independent samples from μ_G^α. Let G' be the disjoint union of k copies of G. The Gibbs measure on G' with density α is

$$\mu_{G'}^\alpha = (\mu_G^\alpha)^{\otimes k}.$$

We can apply Algorithm 4 to G', which has kn vertices and the same maximum degree as G, and produce a sample whose distribution is within total variation distance ε from $(\mu_G^\alpha)^{\otimes k}$ in time

$$\widetilde{O}_{\gamma,\delta,\Delta}\left(\frac{kn}{\varepsilon}\right).$$

For comparison, there is also a classical approach to address the same algorithmic problem: first estimate the fugacity λ_α and then sample from μ_{G,λ_α} using existing samplers for the hard-core model. We present this approach in Appendix B.1.2. As shown there, the bisection approach requires time

$$\widetilde{O}_{\gamma,\delta,\Delta}\left(\frac{n}{\varepsilon^2}\right)$$

for one sample ($k = 1$). Applied to G', for all $k \geq 1$, it requires time

$$\widetilde{O}_{\gamma,\delta,\Delta}\left(\frac{kn}{\varepsilon^2}\right)$$

to approximate $(\mu_G^\alpha)^{\otimes k}$. Thus, the particle-system approach improves the dependence on ε for both one sample and k independent samples.

3.2.2 Hardness of prescribed-density sampling

To complete our understanding of prescribed-density sampling problem, we show that the threshold $\alpha_c(\Delta)$ is also a computational barrier for this problem.

Throughout this section we use the notion of *polynomial time approximate sampler*: a polynomial time approximate sampler for μ_G^α is an algorithm that, given a graph G and an accuracy parameter ε, in time $\mathrm{poly}(|G|, 1/\varepsilon)$ outputs a random independent set whose law is within ε total variation distance of μ_G^α.

The following is the hardness result we prove in this section.

Theorem 3.11. *Fix $\Delta \geq 3$ and $\alpha \in (\alpha_c(\Delta), \frac{1}{\Delta+1})$. Unless* NP $=$ RP*, there is no polynomial time approximate sampler for μ_G^α on graphs with maximum degree Δ.*

This hardness result should be compared with the known hardness of approximately counting independent sets of size αn above $\alpha_c(\Delta)$ [DP23]. Our argument is a direct adaptation of that of [DP23]. The main difference is that our target is μ_G^α, not a uniform distribution on fixed-size independent sets. For a suitable choice of λ, we reduce the hard problem of approximately sampling from $\mu_{G,\lambda}$ on graphs with maximum degree Δ to our problem [SS12]. As in [DP23], we add many copies of cliques $K_{\Delta+1}$ to G in order to bring down the density of the resulting graph H close to α. Then, by a variance comparison we show that approximately sampling from μ_H^α gives an approximate sample from $\mu_{G,\lambda}$.

Lemma 3.12. *Let G be a graph on n vertices and $\hat{\alpha} \geq \alpha > 0$. Assume there exist constants c, C such that for every $\beta \in [\alpha, \hat{\alpha}]$,*

$$cn \leq \mathrm{Var}_{\sigma \sim \mu_G^\beta}(|\sigma|) \leq Cn.$$

Then, we have

$$\|\mu_G^{\alpha} - \mu_G^{\hat{\alpha}}\|_{\mathrm{TV}} \leq \frac{\sqrt{Cn}}{2c}|\hat{\alpha} - \alpha|.$$

We provide the proof of Theorem 3.12 in Appendix B.1.3.

Lemma 3.13. *Fix $\Delta \geq 3$ and $\alpha \in (\alpha_c(\Delta), \frac{1}{\Delta+1})$, and let*

$$\lambda := \frac{\alpha}{1 - \alpha(\Delta + 1)}.$$

If there is a polynomial time approximate sampler for μ_G^{α} on graphs with maximum degree Δ, then there is a polynomial time approximate sampler for $\mu_{G,\lambda}$ on graphs with maximum degree Δ.

Proof. Assume $\mathcal{O}_\alpha$ is an oracle that, given a graph H with maximum degree Δ and $\varepsilon > 0$, in time $\mathrm{poly}(|H|, 1/\varepsilon)$ outputs a random independent set whose law is within ε total variation distance of μ_H^{α}.

Let $G = (V, E)$ be a graph with maximum degree Δ, and write $n = |V|$. We construct a graph H by adding M disjoint copies of complete graphs on $\Delta + 1$ vertices to G, i.e.,

$$H := G \sqcup MK_{\Delta+1},$$

where M will be chosen later as a polynomial in n and $1/\varepsilon$. Note that the graph H also has maximum degree Δ.

Note also that $\alpha(\mu_{K_{\Delta+1},\lambda}) = \alpha$. Let

$$\hat{\alpha} := \alpha(\mu_{H,\lambda}).$$

Since H is a disjoint union, and the lowest density of $\mu_{G',\lambda}$ occurs when $G' = K_{\Delta+1}$, we conclude $\hat{\alpha} \geq \alpha$. Also, we have

$$\hat{\alpha} = \frac{n}{|V(H)|}\alpha(\mu_{G,\lambda}) + \frac{M(\Delta+1)}{|V(H)|}\alpha.$$

Therefore,

$$\hat{\alpha} \leq \alpha + \frac{n}{|V(H)|}.$$

We next verify that the variance assumption in Theorem 3.12 holds for all $\beta \in [\alpha, \hat{\alpha}]$. Let λ_β be the fugacity such that

$$\mu_H^{\beta} = \mu_{H,\lambda_\beta}.$$

Since H is a disjoint union,

$$\mathrm{Var}_{\sigma\sim\mu_H^{\beta}}(|\sigma|) = \mathrm{Var}_{\sigma\sim\mu_{G,\lambda_\beta}}(|\sigma|) + M\,\mathrm{Var}_{\sigma\sim\mu_{K_{\Delta+1},\lambda_\beta}}(|\sigma|).$$

We first prove the upper bound. Since $0 \le |\sigma| \le n$ on G,

$$\mathrm{Var}_{\sigma\sim\mu_{G,\lambda_\beta}}(|\sigma|) \le \frac{n^2}{4}.$$

Also, on $K_{\Delta+1}$, the size of an independent set is either 0 or 1, so

$$\mathrm{Var}_{\sigma\sim\mu_{K_{\Delta+1},\lambda_\beta}}(|\sigma|) \le \frac{1}{4}.$$

Therefore,

$$\mathrm{Var}_{\sigma\sim\mu_H^\beta}(|\sigma|) \le \frac{n^2}{4} + \frac{M}{4}.$$

If $M \ge n^2$, then

$$\mathrm{Var}_{\sigma\sim\mu_H^\beta}(|\sigma|) \le \frac{M}{2} \le \frac{|V(H)|}{2}.$$

This gives the desired upper bound with $C = 1/2$.

We now prove the lower bound. Since

$$\beta = \frac{n}{|V(H)|}\alpha(\mu_{G,\lambda_\beta}) + \frac{M(\Delta+1)}{|V(H)|}\alpha(\mu_{K_{\Delta+1},\lambda_\beta}),$$

with a similar argument as above, we have

$$\alpha(\mu_{K_{\Delta+1},\lambda_\beta}) \le \beta \le \alpha(\mu_{K_{\Delta+1},\lambda_\beta}) + \frac{n}{|V(H)|}.$$

Moreover, since β lies between α and $\hat{\alpha}$, and $|\hat{\alpha} - \alpha| \le \frac{n}{|V(H)|}$, we get

$$|\beta - \alpha| \le \frac{n}{|V(H)|}.$$

Hence

$$|\alpha(\mu_{K_{\Delta+1},\lambda_\beta}) - \alpha| \le \frac{2n}{|V(H)|}.$$

Since $\alpha \in (\alpha_c(\Delta), \frac{1}{\Delta+1})$ is fixed, by taking $M \ge A_{\alpha,\Delta} n$ for a sufficiently large constant $A_{\alpha,\Delta}$, we can ensure that

$$\alpha(\mu_{K_{\Delta+1},\lambda_\beta}) \in \left[\frac{\alpha}{2}, \frac{\alpha + \frac{1}{\Delta+1}}{2}\right].$$

On $K_{\Delta+1}$, the random variable $|\sigma|$ is Bernoulli with parameter

$$x_\beta = \Pr_{\sigma\sim\mu_{K_{\Delta+1},\lambda_\beta}}[|\sigma| = 1] = (\Delta+1)\alpha(\mu_{K_{\Delta+1},\lambda_\beta}).$$

The previous display implies that x_β is bounded away from both 0 and 1 by constants

depending only on α and Δ. Therefore, there exists a constant $b = b(\alpha, \Delta) > 0$ such that

$$\mathrm{Var}_{\sigma \sim \mu_{K_{\Delta+1}, \lambda_\beta}}(|\sigma|) = x_\beta(1 - x_\beta) \geq b.$$

Consequently,

$$\mathrm{Var}_{\sigma \sim \mu_H^\beta}(|\sigma|) \geq Mb.$$

If $M \geq n$, then

$$|V(H)| = n + M(\Delta + 1) \leq M(\Delta + 2),$$

and hence

$$\mathrm{Var}_{\sigma \sim \mu_H^\beta}(|\sigma|) \geq \frac{b}{\Delta + 2}|V(H)|.$$

Thus the variance lower bound in Theorem 3.12 holds with

$$c = \frac{b}{\Delta + 2}.$$

Combining the upper and lower bounds, Theorem 3.12 gives

$$\|\mu_H^\alpha - \mu_H^{\hat{\alpha}}\|_{\mathrm{TV}} \leq \frac{\sqrt{C|V(H)|}}{2c}|\hat{\alpha} - \alpha| \leq \frac{\sqrt{C|V(H)|}}{2c} \cdot \frac{n}{|V(H)|} = \frac{\sqrt{C}}{2c}\frac{n}{\sqrt{|V(H)|}}.$$

But by definition of $\hat{\alpha}$, we have

$$\mu_H^{\hat{\alpha}} = \mu_{H,\lambda}.$$

Choose

$$M = \left\lceil A_{\alpha,\Delta}\frac{n^2}{\varepsilon^2} \right\rceil$$

for a sufficiently large constant $A_{\alpha,\Delta}$. Then

$$\|\mu_H^\alpha - \mu_{H,\lambda}\|_{\mathrm{TV}} \leq \frac{\varepsilon}{2}.$$

Now run the oracle $\mathcal{O}_\alpha$ on H with accuracy $\varepsilon/2$, and let ν be the law of its output. Then

$$\|\nu - \mu_H^\alpha\|_{\mathrm{TV}} \leq \frac{\varepsilon}{2}.$$

We output the restriction of this independent set to the original vertex set $V(G)$. Let P_G denote this projection. Since total variation distance cannot increase under projection,

$$\|P_G\nu - P_G\mu_H^\alpha\|_{\mathrm{TV}} \leq \frac{\varepsilon}{2}.$$

Also,

$$\|P_G\mu_H^\alpha - P_G\mu_{H,\lambda}\|_{\mathrm{TV}} \leq \|\mu_H^\alpha - \mu_{H,\lambda}\|_{\mathrm{TV}} \leq \frac{\varepsilon}{2}.$$

Finally, because $H = G \sqcup MK_{\Delta+1}$ is a disjoint union,

$$P_G\mu_{H,\lambda} = \mu_{G,\lambda}.$$

Therefore,

$$\|P_G\nu - \mu_{G,\lambda}\|_{\mathrm{TV}} \le \varepsilon.$$

Note that by the choice of M, $|H|$ is polynomial in $|G|$, and the oracle runs in time $\mathrm{poly}(|H|, 1/\varepsilon)$. Thus the whole procedure runs in time $\mathrm{poly}(|G|, 1/\varepsilon)$. □

Proof of Theorem 3.11. Set

$$\lambda := \frac{\alpha}{1 - \alpha(\Delta + 1)}.$$

Since

$$\alpha_c(\Delta) = \frac{\lambda_c(\Delta)}{1 + (\Delta + 1)\lambda_c(\Delta)},$$

and the map $x \mapsto \frac{x}{1+(\Delta+1)x}$ is strictly increasing on $(0, \infty)$, the assumption $\alpha > \alpha_c(\Delta)$ implies that

$$\lambda > \lambda_c(\Delta).$$

Assume for contradiction that there is a polynomial time approximate sampler for μ_G^α on graphs with maximum degree Δ. By Theorem 3.13, this gives a polynomial time approximate sampler for $\mu_{G,\lambda}$ on graphs with maximum degree Δ. But hard-core approximate sampling at fugacity $\lambda > \lambda_c(\Delta)$ on Δ-regular graphs is known to be hard [SS12]. Thus, unless NP = RP, there is no polynomial approximate sampler for μ_G^α on graphs with maximum degree Δ. □

3.2.3 Sampling with prescribed vertex marginals

We now focus on the single-site inverse sampling problem. Let $\boldsymbol{m} = (m_v)_{v \in V}$ be a target marginal vector, and let $\boldsymbol{\lambda_m} \in (0, \infty)^V$ be the unique fugacity vector such that [1]

$$\Pr_{\sigma \sim \mu_{G,\boldsymbol{\lambda_m}}}[\sigma_v = 1] = m_v \qquad \forall v \in V.$$

For brevity, throughout this subsection we write

$$\mu_G^{\boldsymbol{m}} := \mu_{G,\boldsymbol{\lambda_m}}.$$

The particle system approach is as follows. We choose N sufficiently large, set

$$r_v = N m_v \qquad \forall v \in V,$$

and consider the space $\Omega^{\mathrm{ss}}_{N,\vec{r}}$ with uniform measure $\mu^{\mathrm{ss}}_{N,\vec{r}}$. We then run the single-site particle chain P^{ss}_N close to its stationary distribution $\mu^{\mathrm{ss}}_{N,\vec{r}}$ and output the first particle. Thus the law of the output is approximately $P_1\mu^{\mathrm{ss}}_{N,\vec{r}}$.

The analysis has two parts. First, we prove that $P_1\mu^{\mathrm{ss}}_{N,\vec{r}}$ is close to the target measure $\mu_G^{\boldsymbol{m}}$. Second, we show how to sample efficiently from $\mu^{\mathrm{ss}}_{N,\vec{r}}$ by running the particle chain

[1] See Appendix B.1 for a discussion of existence and uniqueness of the fugacity vector realizing a prescribed marginal vector.

P_N^{ss}.

To prove the first step, we use a relative entropy argument, along with concentration inequalities. One could also proceed as in the mean-field case, by using a direct estimation via the local central limit theorem (Theorem 3.7). In this case, however, one has to deal with an n-dimensional event and the Fourier transform argument used there appears to lead to a substantially worse dependence of N on n.[2]

The next lemmas provide the proof of the first step.

Lemma 3.14. *With the above notation, assume $\Omega^{\mathrm{ss}}_{N,\vec{r}} \neq \emptyset$. Then we have*

$$D_{\mathrm{KL}}\Big(P_1\mu^{\mathrm{ss}}_{N,\vec{r}}\,\Big\|\,\mu_G^{\boldsymbol{m}}\Big) \leq \frac{1}{N}\log\frac{1}{(\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{r}})}.$$

Consequently,

$$\Big\|P_1\mu^{\mathrm{ss}}_{N,\vec{r}} - \mu_G^{\boldsymbol{m}}\Big\|_{\mathrm{TV}} \leq \sqrt{\frac{1}{2N}\log\frac{1}{(\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{r}})}}.$$

Proof. First, note that for $\boldsymbol{\sigma} \in \Omega^{\mathrm{ss}}_{N,\vec{r}}$ and $\boldsymbol{\lambda} = \boldsymbol{\lambda}_{\boldsymbol{m}}$, we have

$$\mu_{G,\boldsymbol{\lambda}}^{\otimes N}(\boldsymbol{\sigma}) \propto \prod_{i\in[N]}\mu_{G,\boldsymbol{\lambda}}(\sigma^{(i)}) \propto \prod_{i\in[N]}\prod_{v\in V}\lambda_v^{\sigma_v^{(i)}} = \prod_{v\in V}\lambda_v^{\sum_{i\in[N]}\sigma_v^{(i)}} = \prod_{v\in V}\lambda_v^{r_v}.$$

Hence, all elements of $\Omega^{\mathrm{ss}}_{N,\vec{r}}$ have the same measure under $(\mu_G^{\boldsymbol{m}})^{\otimes N}$, which means that

$$\mu^{\mathrm{ss}}_{N,\vec{r}} = (\mu_G^{\boldsymbol{m}})^{\otimes N}\Big(\,\cdot\,\Big|\,\Omega^{\mathrm{ss}}_{N,\vec{r}}\Big).$$

We now have

$$\mu^{\mathrm{ss}}_{N,\vec{r}}(\boldsymbol{\sigma}) = \frac{(\mu_G^{\boldsymbol{m}})^{\otimes N}(\boldsymbol{\sigma})}{(\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{r}})}\,\mathbf{1}(\boldsymbol{\sigma}\in\Omega^{\mathrm{ss}}_{N,\vec{r}}), \qquad \boldsymbol{\sigma}\in\Omega_G^N.$$

Hence, for every $\boldsymbol{\sigma}\in\Omega^{\mathrm{ss}}_{N,\vec{r}}$,

$$\frac{\mu^{\mathrm{ss}}_{N,\vec{r}}(\boldsymbol{\sigma})}{(\mu_G^{\boldsymbol{m}})^{\otimes N}(\boldsymbol{\sigma})} = \frac{1}{(\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{r}})}.$$

Therefore,

$$D_{\mathrm{KL}}(\mu^{\mathrm{ss}}_{N,\vec{r}}\,\|\,(\mu_G^{\boldsymbol{m}})^{\otimes N}) = \sum_{\boldsymbol{\sigma}\in\Omega^{\mathrm{ss}}_{N,\vec{r}}}\mu^{\mathrm{ss}}_{N,\vec{r}}(\boldsymbol{\sigma})\log\frac{1}{(\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{r}})} = \log\frac{1}{(\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{r}})}. \tag{3.16}$$

Now, by the chain rule for relative entropy,

$$D_{\mathrm{KL}}(\mu^{\mathrm{ss}}_{N,\vec{r}}\,\|\,(\mu_G^{\boldsymbol{m}})^{\otimes N}) = \sum_{i=1}^{N}\mathbb{E}_{\boldsymbol{\sigma}\sim\mu^{\mathrm{ss}}_{N,\vec{r}}}\Big[D_{\mathrm{KL}}\big(\mu^{\mathrm{ss}}_{N,\vec{r}}(\sigma^{(i)}\mid\sigma^{(1)},\dots,\sigma^{(i-1)})\,\|\,\mu_G^{\boldsymbol{m}}\big)\Big].$$

[2] On the other hand, we note that a relative entropy argument, along with concentration inequalities, could also be used in the mean-field case, but it gives a worse dependence on N than the one-dimensional local central limit theorem.

By convexity of relative entropy in the first argument,

$$\mathbb{E}_{\boldsymbol{\sigma}\sim\mu^{\mathrm{ss}}_{N,\vec{r}}}\Big[D_{\mathrm{KL}}(\mu^{\mathrm{ss}}_{N,\vec{r}}(\sigma^{(i)} \mid \sigma^{(1)},\dots,\sigma^{(i-1)}) \,\|\, \mu_G^{\boldsymbol{m}})\Big] \geq D_{\mathrm{KL}}(P_i\mu^{\mathrm{ss}}_{N,\vec{r}} \,\|\, \mu_G^{\boldsymbol{m}}).$$

Thus,

$$D_{\mathrm{KL}}(\mu^{\mathrm{ss}}_{N,\vec{r}} \,\|\, (\mu_G^{\boldsymbol{m}})^{\otimes N}) \geq \sum_{i=1}^{N} D_{\mathrm{KL}}(P_i\mu^{\mathrm{ss}}_{N,\vec{r}} \,\|\, \mu_G^{\boldsymbol{m}}) = N \cdot D_{\mathrm{KL}}(P_1\mu^{\mathrm{ss}}_{N,\vec{r}} \,\|\, \mu_G^{\boldsymbol{m}}).$$

Combining this with (3.16) gives

$$D_{\mathrm{KL}}(P_1\mu^{\mathrm{ss}}_{N,\vec{r}} \,\|\, \mu_G^{\boldsymbol{m}}) \leq \frac{1}{N}\log\frac{1}{(\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{r}})}.$$

Finally, by Pinsker's inequality,

$$\|P_1\mu^{\mathrm{ss}}_{N,\vec{r}} - \mu_G^{\boldsymbol{m}}\|_{\mathrm{TV}} \leq \sqrt{\frac{1}{2}D_{\mathrm{KL}}(P_1\mu^{\mathrm{ss}}_{N,\vec{r}} \,\|\, \mu_G^{\boldsymbol{m}})},$$

and therefore

$$\|P_1\mu^{\mathrm{ss}}_{N,\vec{r}} - \mu_G^{\boldsymbol{m}}\|_{\mathrm{TV}} \leq \sqrt{\frac{1}{2N}\log\frac{1}{(\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{r}})}}.$$

□

By Theorem 3.14, to find an upper bound on the total variation distance, it suffices to provide a lower bound on the probability $(\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{r}})$. The next two lemmas establish such a lower bound for the case where $\vec{r}$ is an integer vector such that $\|\vec{r} - N\boldsymbol{m}\|_\infty$ is $O(\sqrt{N\log n})$.

Lemma 3.15. *Assume that*

$$\gamma \leq m_u \leq \frac{1-\delta}{\Delta+1} \qquad \forall u \in V$$

for some constants $\gamma > 0$ and $\delta \in (0,1)$. Then there exists $C = C(\gamma,\delta,\Delta)$ such that for all $N \geq C\log(4n)$, the following hold.

First, for every $\vec{s} \in \mathbb{Z}^V_{\geq 0}$ satisfying

$$\|\vec{s} - N\boldsymbol{m}\|_\infty \leq \sqrt{\frac{N}{2}\log(4n)},$$

one has $\Omega^{\mathrm{ss}}_{N,\vec{s}} \neq \emptyset$.

Second, let $\vec{r}, \vec{r}' \in \mathbb{Z}^V_{\geq 0}$ satisfy

$$\|\vec{r} - N\boldsymbol{m}\|_\infty \leq \sqrt{\frac{N}{2}\log(4n)}, \qquad \|\vec{r}' - N\boldsymbol{m}\|_\infty \leq \sqrt{\frac{N}{2}\log(4n)},$$

and suppose that for some $v \in V$,

$$r'_v = r_v + 1, \qquad r'_u = r_u \quad \textit{for all } u \in V \setminus \{v\}.$$

Then

$$\frac{\gamma}{2} \le \frac{\Omega^{\mathrm{ss}}_{N,\vec{r}}}{\Omega^{\mathrm{ss}}_{N,\vec{r}'}} \le \frac{2}{\delta}.$$

Proof. Let $C = C(\gamma, \delta, \Delta)$ be sufficiently large so that for every $N \ge C \log(4n)$ one has

$$2(\Delta+1)\sqrt{\frac{N}{2}\log(4n)} \le \delta N \qquad \text{and} \qquad \sqrt{\frac{N}{2}\log(4n)} \le \frac{\gamma}{2}N.$$

We now prove the first statement. Fix $\vec{s} \in \mathbb{Z}^V_{\ge 0}$ satisfying $\|\vec{s} - N\boldsymbol{m}\|_\infty \le \sqrt{\frac{N}{2}\log(4n)}$. We show that $\Omega^{\mathrm{ss}}_{N,\vec{s}} \neq \emptyset$. Indeed, since G has maximum degree at most Δ, it admits a proper $(\Delta+1)$-coloring. Let $V = V_1 \cup \cdots \cup V_{\Delta+1}$ be the corresponding partition of V into independent sets. For each $j \in [\Delta+1]$, define

$$M_j := \max_{u \in V_j} s_u,$$

with the convention that $M_j = 0$ if $V_j = \emptyset$. Since $\|\vec{s} - N\boldsymbol{m}\|_\infty \le \sqrt{\frac{N}{2}\log(4n)}$, we have for every $u \in V$,

$$s_u \le N m_u + \sqrt{\frac{N}{2}\log(4n)} \le \frac{1-\delta}{\Delta+1}N + \sqrt{\frac{N}{2}\log(4n)}.$$

Hence

$$M_j \le \frac{1-\delta}{\Delta+1}N + \sqrt{\frac{N}{2}\log(4n)} \qquad \text{for all } j \in [\Delta+1].$$

Therefore,

$$\sum_{j=1}^{\Delta+1} M_j \le (1-\delta)N + (\Delta+1)\sqrt{\frac{N}{2}\log(4n)} \le N.$$

Now, partition $[N]$ into pairwise disjoint sets $I_1, \ldots, I_{\Delta+1}, I_0$ such that $I_j = M_j$ for each $j \in [\Delta+1]$. For each $j \in [\Delta+1]$ and each $u \in V_j$, since $s_u \le M_j$, we may choose a set

$$J_u \subseteq I_j \qquad \text{with} \qquad J_u = s_u.$$

Now, for each $t \in I_j$, define

$$\sigma^{(t)} := u \in V_j : t \in J_u,$$

and for each $t \in I_0$, set $\sigma^{(t)} := \emptyset$.

Since each V_j is an independent set, every $\sigma^{(t)}$ is an independent set. Moreover, for each $u \in V_j$, the vertex u belongs to exactly the sets $\sigma^{(t)}$ with $t \in J_u$, and hence appears

exactly s_u times among $\sigma^{(1)}, \dots, \sigma^{(N)}$. Therefore

$$\sum_{t=1}^{N} \sigma^{(t)} = \vec{s},$$

which shows that $(\sigma^{(1)}, \dots, \sigma^{(N)}) \in \Omega^{\mathrm{ss}}_{N,\vec{s}}$. Hence $\Omega^{\mathrm{ss}}_{N,\vec{s}} \neq \emptyset$.

Note that the first part of the lemma shows that the ratio we want to bound in the second part is well-defined, and thus now we can prove the second statement.

Let

$$L := \Omega^{\mathrm{ss}}_{N,\vec{r}}, \qquad R := \Omega^{\mathrm{ss}}_{N,\vec{r}'},$$

and define a bipartite graph $H = (L \cup R, F)$ as follows. For $\boldsymbol{\sigma} = (\sigma^{(1)}, \dots, \sigma^{(N)}) \in L$ and $\boldsymbol{\tau} = (\tau^{(1)}, \dots, \tau^{(N)}) \in R$, we place an edge between $\boldsymbol{\sigma}$ and $\boldsymbol{\tau}$ if there exists $j \in [N]$ such that

$$\tau^{(j)} = \sigma^{(j)} \cup \{v\}, \qquad \tau^{(k)} = \sigma^{(k)} \quad \text{for all } k \neq j,$$

and $\sigma^{(j)} \cup \{v\}$ is an independent set. Equivalently, $\boldsymbol{\sigma}$ is obtained from $\boldsymbol{\tau}$ by deleting v from exactly one coordinate j for which $v \in \tau^{(j)}$. Clearly, every vertex of H has degree at most N.

We first bound the degree of vertices in R from below. Fix $\boldsymbol{\tau} \in R$. Then $\deg_H(\boldsymbol{\tau}) = r'_v$, since each coordinate $i \in [N]$ with $v \in \tau^{(i)}$ yields one neighbor in L, obtained by deleting v from $\tau^{(i)}$. Therefore,

$$\deg_H(\boldsymbol{\tau}) = r'_v \geq Nm_v - \sqrt{\frac{N}{2}\log(4n)} \geq \gamma N - \sqrt{\frac{N}{2}\log(4n)} \geq \frac{\gamma}{2}N.$$

Next, we bound the degree of vertices in L from below. Fix $\boldsymbol{\sigma} \in L$, and let

$$M := \sum_{i \in [N]} \sigma_v^{(i)}.$$

Let

$$B := i \in [N] : \sigma_v^{(i)} \neq \emptyset.$$

Note that if $i \notin B$, then v can be added to $\sigma^{(i)}$ while preserving independence. Thus

$$\deg_H(\boldsymbol{\sigma}) = N - B.$$

Since each $i \in B$ contributes at least 1 to the sum defining M, we have $B \leq M$, and therefore

$$\deg_H(\boldsymbol{\sigma}) \geq N - M.$$

On the other hand, since $\boldsymbol{\sigma} \in \Omega^{\mathrm{ss}}_{N,\vec{r}}$,

$$M = \sum_{u \in v} r_u.$$

Using $v \le \Delta + 1$ and the bound $r_u \le N m_u + \sqrt{\frac{N}{2}\log(4n)}$ for every $u \in V$, we obtain

$$M \le (\Delta+1)\left(\max_{u\in V} m_u \cdot N + \sqrt{\frac{N}{2}\log(4n)}\right) \le (\Delta+1)\left(\frac{1-\delta}{\Delta+1}N + \sqrt{\frac{N}{2}\log(4n)}\right).$$

Hence

$$M \le (1-\delta)N + (\Delta+1)\sqrt{\frac{N}{2}\log(4n)},$$

and so

$$\deg_H(\boldsymbol{\sigma}) \ge N - M \ge \delta N - (\Delta+1)\sqrt{\frac{N}{2}\log(4n)} \ge \frac{\delta}{2}N.$$

We now double-count the edges of H. Since every vertex has degree at most N, while every vertex in L has degree at least $\frac{\delta}{2}N$ and every vertex in R has degree at least $\frac{\gamma}{2}N$, we have

$$\frac{\delta}{2}N\,L \le F \le N\,R,$$

and

$$\frac{\gamma}{2}N\,R \le F \le N\,L.$$

Therefore,

$$\frac{L}{R} \le \frac{2}{\delta} \qquad \text{and} \qquad \frac{L}{R} \ge \frac{\gamma}{2}.$$

Since $L = \Omega^{\mathrm{ss}}_{N,\vec{r}}$ and $R = \Omega^{\mathrm{ss}}_{N,\vec{r}'}$, this gives

$$\frac{\gamma}{2} \le \frac{\Omega^{\mathrm{ss}}_{N,\vec{r}}}{\Omega^{\mathrm{ss}}_{N,\vec{r}'}} \le \frac{2}{\delta},$$

as claimed. □

Lemma 3.16. *Assume that*

$$\gamma \le m_u \le \frac{1-\delta}{\Delta+1} \qquad \forall u \in V$$

for some constants $\gamma > 0$ and $\delta \in (0,1)$. Then there exists $C = C(\gamma,\delta,\Delta)$ such that for all $N \ge C\log(4n)$, if we define

$$L := \sqrt{\frac{N}{2}\log(4n)} \qquad \textit{and} \qquad \mathcal{B} := \vec{r} \in \mathbb{Z}^V : \|\vec{r} - N\boldsymbol{m}\|_\infty \le L,$$

then

$$\min_{\vec{r}\in\mathcal{B}} (\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{r}}) \ge \frac{1}{2}(2L+1)^{-n}\left(\frac{\delta\gamma}{2}\right)^{2nL}.$$

Proof. Let $C = C(\gamma,\delta,\Delta)$ be as in Theorem 3.15, and fix $N \ge C\log(4n)$. For $\vec{r} \in \mathcal{B}$, we first show that

$$\sum_{\vec{r}\in\mathcal{B}} (\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{r}}) \ge \frac{1}{2}.$$

Let $\boldsymbol{\sigma} = (\sigma^{(1)}, \dots, \sigma^{(N)}) \sim (\mu_G^{\boldsymbol{m}})^{\otimes N}$, and set $\vec{R} := \sum_{i=1}^N \sigma^{(i)}$. Fix $v \in V$. Then $R_v = \sum_{i=1}^N \sigma_v^{(i)}$, where $\sigma_v^{(1)}, \dots, \sigma_v^{(N)}$ are independent $\{0,1\}$-valued random variables with $\mathbb{E}[\sigma_v^{(i)}] = m_v$. Hence, by Hoeffding's inequality,

$$(\mu_G^{\boldsymbol{m}})^{\otimes N}(|R_v - Nm_v| > L) \le 2\exp\left(-\frac{2L^2}{N}\right) = 2e^{-\log(4n)} = \frac{1}{2n}.$$

Applying the union bound over all $v \in V$, we obtain

$$(\mu_G^{\boldsymbol{m}})^{\otimes N}\left(\|\vec{R} - N\boldsymbol{m}\|_\infty > L\right) \le \frac{1}{2}.$$

Therefore,

$$\sum_{\vec{r}\in\mathcal{B}} (\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{r}}) = (\mu_G^{\boldsymbol{m}})^{\otimes N}\left(\|\vec{R} - N\boldsymbol{m}\|_\infty \le L\right) \ge \frac{1}{2}.$$

We next compare the probabilities of neighboring vectors. Let $\vec{r}, \vec{r}' \in \mathcal{B}$ satisfy

$$r'_v = r_v + 1, \qquad r'_u = r_u \quad \text{for all } u \in V \setminus \{v\},$$

for some $v \in V$. Let $\boldsymbol{\lambda} = \boldsymbol{\lambda}_{\boldsymbol{m}}$. Since all elements of $\Omega^{\mathrm{ss}}_{N,\vec{r}}$ have the same measure under $(\mu_G^{\boldsymbol{m}})^{\otimes N}$, we have

$$\frac{(\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{r}})}{(\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{r}'})} = \frac{|\Omega^{\mathrm{ss}}_{N,\vec{r}}|}{|\Omega^{\mathrm{ss}}_{N,\vec{r}'}|} \prod_{u\in V} \lambda_u^{r_u - r'_u} = \frac{|\Omega^{\mathrm{ss}}_{N,\vec{r}}|}{|\Omega^{\mathrm{ss}}_{N,\vec{r}'}|} \cdot \frac{1}{\lambda_v}.$$

By Theorem 3.15,

$$\frac{\gamma}{2} \le \frac{|\Omega^{\mathrm{ss}}_{N,\vec{r}}|}{|\Omega^{\mathrm{ss}}_{N,\vec{r}'}|} \le \frac{2}{\delta}.$$

We now bound the fugacities. For each $v \in V$, one has

$$m_v = \lambda_v \, \mu_G^{\boldsymbol{m}}(\sigma_u = 0 \text{ for all } u \in v)\,.$$

Hence $\lambda_v \ge m_v \ge \gamma$. Also, by the union bound,

$$\mu_G^{\boldsymbol{m}}(\sigma_u = 0 \text{ for all } u \in v) \ge 1 - \sum_{u\in v} m_u \ge 1 - (\Delta+1)\cdot\frac{1-\delta}{\Delta+1} = \delta,$$

and therefore

$$\lambda_v = \frac{m_v}{\mu_G^{\boldsymbol{m}}(\sigma_u = 0 \text{ for all } u \in v)} \le \frac{1}{\delta}.$$

Thus $\gamma \le \lambda_v \le \frac{1}{\delta}$, which implies

$$\frac{\delta\gamma}{2} \le \frac{(\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{r}})}{(\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{r}'})} \le \frac{2}{\delta\gamma}.$$

Now let $\vec{r}_* \in \mathcal{B}$ be such that

$$(\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{r}_*}) = \min_{\vec{r}\in\mathcal{B}} (\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{r}}).$$

Fix any $\vec{s} \in \mathcal{B}$. There exists a path $\vec{r}_* = \vec{r}^{(0)}, \vec{r}^{(1)}, \dots, \vec{r}^{(m)} = \vec{s}$ in $\mathcal{B}$ such that consecutive points differ in exactly one coordinate by ± 1, and $m = \|\vec{s} - \vec{r}_*\|_1$. Applying the previous bound along this path gives

$$(\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{s}}) \le \left(\frac{2}{\delta\gamma}\right)^m (\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{r}_*}) = \left(\frac{2}{\delta\gamma}\right)^{\|\vec{s}-\vec{r}_*\|_1} (\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{r}_*}).$$

Since $\vec{s}, \vec{r}_* \in \mathcal{B}$,

$$\|\vec{s} - \vec{r}_*\|_1 \le \sum_{u\in V} (|s_u - Nm_u| + |r_{*,u} - Nm_u|) \le 2nL.$$

Hence

$$(\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{s}}) \le \left(\frac{2}{\delta\gamma}\right)^{2nL} (\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{r}_*}).$$

Also, for each $u \in V$, the coordinate r_u ranges over an interval of length $2L$, and hence has at most $2L+1$ possible integer values. Therefore $|\mathcal{B}| \le (2L+1)^n$. Summing over $\vec{s} \in \mathcal{B}$, we obtain

$$\frac{1}{2} \le \sum_{\vec{s}\in\mathcal{B}} (\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{s}}) \le |\mathcal{B}| \left(\frac{2}{\delta\gamma}\right)^{2nL} (\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{r}_*}) \le (2L+1)^n \left(\frac{2}{\delta\gamma}\right)^{2nL} (\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{r}_*}).$$

Rearranging yields

$$(\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{r}_*}) \ge \frac{1}{2}(2L+1)^{-n} \left(\frac{\delta\gamma}{2}\right)^{2nL}.$$

we conclude that

$$\min_{\vec{r}\in\mathcal{B}} (\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{r}}) = (\mu_G^{\boldsymbol{m}})^{\otimes N}(\Omega^{\mathrm{ss}}_{N,\vec{r}_*}) \ge \frac{1}{2}(2L+1)^{-n} \left(\frac{\delta\gamma}{2}\right)^{2nL}.$$

This proves the lemma. □

Corollary 3.17. *By Theorems 3.14 to 3.16, for every $\vec{r} \in \mathcal{B}$,*

$$\|P_1\mu^{\mathrm{ss}}_{N,\vec{r}} - \mu_G^{\boldsymbol{m}}\|_{\mathrm{TV}} \le \sqrt{\frac{1}{2N}\left(\log 2 + n\log(2L+1) + 2nL\log\frac{2}{\delta\gamma}\right)},$$

where

$$L = \sqrt{\frac{N}{2}\log(4n)}.$$

In particular, there exists $C = C(\gamma, \delta, \Delta)$ such that if

$$N \ge C\frac{n^2\log(4n)}{\varepsilon^4},$$

then

$$\|P_1\mu^{\mathrm{ss}}_{N,\vec{r}} - \mu^{\boldsymbol{m}}_G\|_{\mathrm{TV}} \le \varepsilon \qquad \forall \vec{r} \in \mathcal{B}.$$

Moreover, one has $\vec{r} = N\boldsymbol{m} \in \mathcal{B}$*, and hence*

$$\|P_1\mu^{\mathrm{ss}}_{N,\vec{r}} - \mu^{\boldsymbol{m}}_G\|_{\mathrm{TV}} \le \varepsilon.$$

We now consider sampling from $\mu^{\mathrm{ss}}_{N,\vec{r}}$ by running the particle-system chain P^{ss}_N. Let $\gamma \in (0,1)$, and let $c_{\mathrm{col}} > 1$ and Δ_0 be the coloring-mixing constants in Section 3.1. Assume that

$$\Delta \ge \Delta_0, \qquad \gamma \le m_v \le \frac{1}{c_{\mathrm{col}}(\Delta+1)} \qquad \forall v \in V. \tag{3.17}$$

Note that

$$\frac{r_v}{N} = \frac{Nm_v}{N} \le \frac{1}{c_{\mathrm{col}}(\Delta+1)}.$$

Also, if $N \ge \frac{2}{\gamma}$

$$\frac{r_v}{N} = \frac{Nm_v}{N} \ge m_v - \frac{1}{N} \ge \frac{\gamma}{2}.$$

Thus, for $N \ge \frac{2}{\gamma}$, $\vec{r}$ satisfies the assumption in (3.7) with parameter $\frac{\gamma}{2}$. Thus, applying Theorem 3.5 with parameter $\frac{\gamma}{2}$, we conclude that for every $N \ge N_0(\Delta, c_{\mathrm{col}})$, P^{ss}_N mixes within total variation distance ζ of $\mu^{\mathrm{ss}}_{N,\vec{r}}$ in

$$O_{\gamma,\Delta,c_{\mathrm{col}}}\left(Nn \log \frac{Nn}{\zeta}\right)$$

steps.

We are now ready to state the resulting algorithm. First, we give a deterministic procedure for constructing an initial configuration in $\Omega^{\mathrm{ss}}_{N,\vec{r}}$. This will be used as the starting point for the particle-system chain in the final sampling algorithm. Algorithm 5 constructs an element of $\Omega^{\mathrm{ss}}_{N,\vec{r}}$ whenever

$$(\Delta+1)\max_{v\in V} r_v \le N.$$

In our application, this condition holds for $r_v = Nm_v$ by (3.17). Note that this is exactly the constructive procedure used in the proof of Theorem 3.15 to show that $\Omega^{\mathrm{ss}}_{N,\vec{r}} \neq \emptyset$.

Theorem 3.18. *Let* $\gamma \in (0,1)$*, and let* $c_{\mathrm{col}} > 1$ *and* Δ_0 *be the coloring-mixing constants in Section 3.1. Let* $G = (V,E)$ *be a graph on* n *vertices with maximum degree* $\Delta \ge \Delta_0$*, and let* $\boldsymbol{m}$ *be a marginal vector satisfying*

$$\gamma \le m_v \le \frac{1}{c_{\mathrm{col}}(\Delta+1)} \qquad \forall v \in V.$$

Then, for every $\varepsilon \in (0,1)$*, Algorithm 6 outputs a random independent set whose law is*

Algorithm 5 Constructing an Initial Particle Configuration

Require: Graph $G = (V, E)$ with maximum degree Δ, an integer vector $\vec{r} \in \mathbb{Z}_{\geq 0}^{V}$ satisfying $\max_{v \in V} r_v \leq \frac{N}{\Delta+1}$, and an integer N.

Ensure: A configuration $\boldsymbol{\sigma}_0 \in \Omega_{N,\vec{r}}^{\mathrm{ss}}$.

1: Find a proper $(\Delta+1)$-coloring of G, and let

$$V = V_1 \cup \cdots \cup V_{\Delta+1}$$

be the corresponding partition into independent sets.

2: Set

$$M_k := \max_{v \in V_k} r_v \quad \text{with } M_k = 0 \text{ if } V_k = \emptyset \qquad \forall k \in [\Delta+1]$$

3: Partition $[N]$ into pairwise disjoint sets

$$I_1, \ldots, I_{\Delta+1}, I_0$$

such that $I_k = M_k$ for every $k \in [\Delta+1]$.

4: For every $k \in [\Delta+1]$ and every $v \in V_k$, choose

$$J_v \subseteq I_k \qquad \text{with} \qquad J_v = r_v.$$

5: Define

$$\sigma^{(t)} := v \in V : t \in J_v \qquad \text{for every } t \in [N].$$

6: **return** $\boldsymbol{\sigma}_0 = (\sigma^{(1)}, \ldots, \sigma^{(N)})$.

within total variation distance ε from $\mu_G^{\boldsymbol{m}}$. Moreover, the running time is

$$\widetilde{O}_{\gamma,\Delta,c_{\mathrm{col}}}\left(\frac{n^3}{\varepsilon^4}\right).$$

Proof. By choosing the constant $C = C(\gamma, \Delta, c_{\mathrm{col}})$ sufficiently large, Theorem 3.17, applied with $\delta = 1 - \frac{1}{c_{\mathrm{col}}}$ and accuracy $\varepsilon/2$, gives

$$\left\| P_1 \mu_{N,\vec{r}}^{\mathrm{ss}} - \mu_G^{\boldsymbol{m}} \right\|_{\mathrm{TV}} \leq \frac{\varepsilon}{2}.$$

Also, by Theorem 3.5, after running P_N^{ss} for

$$O_{\gamma,\Delta,c_{\mathrm{col}}}\left(Nn \log \frac{Nn}{\varepsilon}\right)$$

steps, the law $\widehat{\mu}$ of the resulting particle configuration satisfies

$$\left\| \widehat{\mu} - \mu_{N,\vec{r}}^{\mathrm{ss}} \right\|_{\mathrm{TV}} \leq \frac{\varepsilon}{2}.$$

Let ν denote the law of the output $\sigma^{(1)}$. Since total variation distance cannot increase

Algorithm 6 Particle System Sampler

Require: Graph $G = (V, E)$ with maximum degree Δ, target marginals $\boldsymbol{m} = (m_v)_{v \in V}$ satisfying (3.17), and accuracy $\varepsilon \in (0, 1)$.

Ensure: A random independent set whose law is within total variation distance ε from $\mu_G^{\boldsymbol{m}}$.

1: Choose
$$N \geq C \frac{n^2 \log(4n)}{\varepsilon^4},$$
where $C = C(\gamma, \Delta, c_{\mathrm{col}})$ is a sufficiently large constant.

2: Set
$$r_v := N m_v \qquad \text{for every } v \in V.$$

3: Use Algorithm 5 to construct an initial configuration
$$\boldsymbol{\sigma}_0 \in \Omega^{\mathrm{ss}}_{N,\vec{r}}.$$

4: Run the particle system chain P^{ss}_N from $\boldsymbol{\sigma}_0$ for
$$O_{\gamma,\Delta,c_{\mathrm{col}}}\left(Nn \log \frac{Nn}{\varepsilon}\right)$$
steps.

5: **return** $\sigma^{(1)}$, where
$$\boldsymbol{\sigma} = (\sigma^{(1)}, \ldots, \sigma^{(N)})$$
is the resulting particle configuration.

under projection to the first particle,

$$\left\| \nu - P_1 \mu^{\mathrm{ss}}_{N,\vec{r}} \right\|_{\mathrm{TV}} = \left\| P_1 \widehat{\mu} - P_1 \mu^{\mathrm{ss}}_{N,\vec{r}} \right\|_{\mathrm{TV}} \leq \frac{\varepsilon}{2}.$$

Therefore,

$$\|\nu - \mu_G^{\boldsymbol{m}}\|_{\mathrm{TV}} \leq \left\| \nu - P_1 \mu^{\mathrm{ss}}_{N,\vec{r}} \right\|_{\mathrm{TV}} + \left\| P_1 \mu^{\mathrm{ss}}_{N,\vec{r}} - \mu_G^{\boldsymbol{m}} \right\|_{\mathrm{TV}} \leq \varepsilon.$$

Finally, by the choice of N,

$$N = O_{\gamma,\Delta,c_{\mathrm{col}}}\left(\frac{n^2 \log n}{\varepsilon^4}\right).$$

The initial construction takes $O_\Delta(Nn)$ time, and the running time of the chain is

$$O_{\gamma,\Delta,c_{\mathrm{col}}}\left(Nn \log \frac{Nn}{\varepsilon}\right).$$

Thus the total running time is

$$\widetilde{O}_{\gamma,\Delta,c_{\mathrm{col}}}\left(\frac{n^3}{\varepsilon^4}\right).$$

□

Remark 3.19. There is also a classical approach to the same algorithmic problem: first,

estimate the fugacity vector $\boldsymbol{\lambda}_m$, and then sample from $\mu_{G,\boldsymbol{\lambda}_m}$ using existing samplers for the hard-core model. We present this approach in Appendix B.1.1.

Bibliography

[ALG21] Nima Anari, Kuikui Liu, and Shayan Oveis Gharan. "Spectral independence in high-dimensional expanders and applications to the hardcore model." In: *SIAM Journal on Computing* 53.6 (2021), FOCS20–1 (cit. on pp. 3, 4, 11).

[Ana+21] Nima Anari, Vishesh Jain, Frederic Koehler, Huy Tuan Pham, and Thuy-Duong Vuong. "Entropic independence II: optimal sampling and concentration via restricted modified log-Sobolev inequalities." In: *arXiv preprint arXiv:2111.03247* (2021) (cit. on pp. 8, 103).

[Ana+22] Nima Anari, Vishesh Jain, Frederic Koehler, Huy T. Pham, and Thuy-Duong Vuong. "Entropic Independence: Optimal Mixing of Down-Up Random Walks." In: *Proceedings of the 54th ACM Symposium on Theory of Computing (STOC)*. 2022 (cit. on pp. 9, 20, 56, 57, 180, 189).

[And+11] C. Andrieu, A. Jasra, A. Doucet, and P. Del Moral. "On nonlinear Markov chain Monte Carlo." In: *Bernoulli* 17.3 (2011), pp. 987–1014 (cit. on p. 14).

[Asa+17] Arash Asadpour, Michel X Goemans, Aleksander Madry, Shayan Oveis Gharan, and Amin Saberi. "An $O(\log n/\log\log n)$-approximation algorithm for the asymmetric traveling salesman problem." In: *Operations Research* 65.4 (2017), pp. 1043–1061 (cit. on p. 12).

[Bas+22] Giada Basile, Dario Benedetto, Lorenzo Bertini, and Emanuele Caglioti. "Large deviations for a binary collision model: energy evaporation." In: *Mathematical Engineering* 5.1 (2022) (cit. on p. 1).

[BB19] Roland Bauerschmidt and Thierry Bodineau. "A very simple proof of the LSI for high temperature spin systems." In: *Journal of Functional Analysis* 276.8 (2019), pp. 2582–2588 (cit. on pp. 8, 58).

[BBD24a] Roland Bauerschmidt, Thierry Bodineau, and Benoıt Dagallier. "Kawasaki dynamics beyond the uniqueness threshold." In: *Probability Theory and Related Fields* (2024), pp. 1–36 (cit. on pp. 8, 9, 20, 54, 56).

[BBD24b] Roland Bauerschmidt, Thierry Bodineau, and Benoıt Dagallier. "Stochastic dynamics and the Polchinski equation: an introduction." In: *Probability Surveys* 21 (2024), pp. 200–290 (cit. on pp. 9, 20).

[BBL09] Julius Borcea, Petter Brändén, and Thomas M. Liggett. "Negative dependence and the geometry of polynomials." In: *Journal of the American Mathematical Society* 22.2 (2009), pp. 521–567 (cit. on pp. 56, 57, 59).

[BBS15] Ellen Baake, Michael Baake, and M. Salamat. "The General Recombination Equation in Continuous Time and Its Solution." In: *Discrete and Continuous Dynamical Systems* 36 (2015), pp. 63–95 (cit. on p. 23).

[BD97] Russ Bubley and Martin Dyer. "Path coupling: A technique for proving rapid mixing in Markov chains." In: *Proceedings 38th Annual Symposium on Foundations of Computer Science*. IEEE. 1997, pp. 223–231 (cit. on pp. 20, 105).

[BE19] Sébastien Bubeck and Ronen Eldan. "The entropic barrier: Exponential families, log-concave geometry, and self-concordance." In: *Mathematics of Operations Research* 44.1 (2019), pp. 264–276 (cit. on p. 30).

[BGS14] Guy Bresler, David Gamarnik, and Devavrat Shah. "Hardness of parameter estimation in graphical models." In: *Advances in Neural Information Processing Systems* 27 (2014) (cit. on pp. 12, 172, 173).

[Bla+22] Antonio Blanca, Pietro Caputo, Zongchen Chen, Daniel Parisi, Daniel Štefankovič, and Eric Vigoda. "On mixing of markov chains: Coupling, spectral independence, and entropy factorization." In: *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 2022, pp. 3670–3692 (cit. on pp. 3, 8).

[BN19] Guy Bresler and Dheeraj Nagaraj. "Stein's method for stationary distributions of Markov chains and application to Ising models." In: (2019) (cit. on p. 105).

[Bol72] Ludwig Boltzmann. "Weitere Studien über das Wärmegleichgewicht unter Gasmolekülen." In: *Sitzungsberichte der Kaiserlichen Akademie der Wissenschaften. Mathematisch-Naturwissenschaftliche Classe* 66 (1872), pp. 275–370 (cit. on p. 1).

[Bon+18] Federico Bonetto, Anna Geisinger, Michael Loss, and Tobias Ried. "Entropy Decay for the Kac Evolution." In: *Communications in Mathematical Physics* 363.3 (2018), pp. 847–875 (cit. on p. 19).

[But14] O. Butkovsky. "On ergodic properties of nonlinear Markov chains and stochastic McKean-Vlasov equations." In: *Theory of Probability & Its Applications* 58.4 (2014), pp. 661–674 (cit. on p. 14).

[BW99] Graham R Brightwell and Peter Winkler. "Graph homomorphisms and phase transitions." In: *Journal of combinatorial theory, series B* 77.2 (1999), pp. 221–262 (cit. on p. 2).

[Cap23] Pietro Caputo. *Lecture Notes on Entropy and Markov Chains*. Available at: http://www.mat.uniroma3.it/users/caputo/entropy.pdf. 2023 (cit. on pp. 8, 57, 62, 118).

[Car+10] Eric A. Carlen, Maria C. Carvalho, Julien Le Roux, Michael Loss, and Cédric Villani. "Entropy and chaos in the Kac model." In: *Kinetic and Related Models* 3.1 (2010), pp. 85–122 (cit. on pp. 1, 8, 9, 19).

[CCG00] Eric A. Carlen, Maria C. Carvalho, and Elisabetta Gabetta. "Central Limit Theorem for Maxwellian Molecules and Truncation of the Wild Expansion." In: *Communications on Pure and Applied Mathematics* 53.3 (2000), pp. 370–397 (cit. on p. 23).

[CCP25] Pietro Caputo, Zongchen Chen, and Daniele Parisi. "Factorizations of relative entropy using stochastic localization." In: *arXiv preprint arXiv:2503.19419* (2025) (cit. on pp. 9, 20, 54–56, 58, 62).

[CD22] Louis-Pierre Chaintron and Anaïs Diez. "Propagation of chaos: a review of models, methods and applications. I. Models and methods." In: *Kinetic and Related Models* 15.6 (2022) (cit. on p. 19).

[CDP09] Pietro Caputo, Paolo Dai Pra, and Gustavo Posta. "Convex entropy decay via the Bochner-Bakry-Emery approach." In: *Annales de l'IHP Probabilités et statistiques*. Vol. 45. 3. 2009, pp. 734–753 (cit. on p. 8).

[CE22] Yat Tin Chen and Ronen Eldan. "Localization schemes: A framework for proving mixing bounds for Markov chains." In: *Proceedings of the IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2022, pp. 110–122 (cit. on pp. 8, 9, 11, 20, 54, 55, 58, 75, 94, 95, 180, 189).

[CFL09] Claudio Castellano, Santo Fortunato, and Vittorio Loreto. "Statistical physics of social dynamics." In: *Reviews of modern physics* 81.2 (2009), pp. 591–646 (cit. on p. 2).

[CGM21] Mary Cryan, Heng Guo, and Georgios Mousa. "Modified Log-Sobolev Inequalities for Strongly Log-Concave Distributions." In: *Annals of Probability* 49 (2021), pp. 506–525 (cit. on pp. 9, 20, 56, 57).

[Che+21] Xiaoyu Chen, Weiming Feng, Yitong Yin, and Xinyuan Zhang. "Rapid Mixing of Glauber Dynamics via Spectral Independence for All Degrees." In: *arXiv preprint arXiv:2105.15005* (2021) (cit. on p. 176).

[Che+22] Xiaoyu Chen, Weiming Feng, Yitong Yin, and Xinyuan Zhang. "Optimal mixing for two-state anti-ferromagnetic spin systems." In: *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2022, pp. 588–599 (cit. on pp. 11, 180, 189).

[Che+25] Xiaoyu Chen, Zongchen Chen, Yitong Yin, and Xinyuan Zhang. "Rapid mixing at the uniqueness threshold." In: *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*. 2025, pp. 879–890 (cit. on p. 11).

[CKV20] L Elisa Celis, Vijay Keswani, and Nisheeth Vishnoi. "Data preprocessing to mitigate bias: A maximum entropy based approach." In: *International Conference on Machine Learning*. PMLR. 2020, pp. 1349–1359 (cit. on p. 12).

[CLL25] Pietro Caputo, Cyril Labbé, and Hubert Lacoin. "Cutoff phenomenon in nonlinear recombinations." In: *Annals of Applied Probability* 35.2 (2025), pp. 1164–1197 (cit. on pp. 9, 19, 23, 36).

[CLV21] Zongchen Chen, Kuikui Liu, and Eric Vigoda. "Optimal mixing of Glauber dynamics: Entropy factorization via high-dimensional expansion." In: *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. 2021, pp. 1537–1550 (cit. on pp. 93, 101, 103, 180, 189).

[CLV23] Zongchen Chen, Kuikui Liu, and Eric Vigoda. "Rapid mixing of Glauber dynamics up to uniqueness via contraction." In: *SIAM Journal on Computing* 52.1 (2023), pp. 196–237 (cit. on pp. 3, 4).

[CM25] Pietro Caputo and Mario Morellini. "Kac's Program and Relative Entropy Decay for Nonlinear Spin-Exchange Dynamics." In: *arXiv preprint arXiv:2511.05223* (2025) (cit. on pp. 3, 18, 21).

[CMR02] Nicoletta Cancrini, Fabio Martinelli, and Cyril Roberto. "The Logarithmic Sobolev Constant of Kawasaki Dynamics under a Mixing Condition Revisited." In: *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques* 38 (2002), pp. 385–436 (cit. on p. 19).

[CP07] Pietro Caputo and Gustavo Posta. "Entropy dissipation estimates in a zero-range dynamics." In: *Probability theory and related fields* 139.1 (2007), pp. 65–87 (cit. on p. 8).

[CP21] Pietro Caputo and Daniel Parisi. "Block factorization of the relative entropy via spatial mixing." In: *Communications in Mathematical Physics* 388.2 (2021), pp. 793–818 (cit. on p. 8).

[CP24] Pietro Caputo and Daniele Parisi. "Nonlinear recombinations and generalized random transpositions." In: *Annales Henri Lebesgue* 7 (2024), pp. 1245–1299 (cit. on pp. 2, 19, 23, 36, 65, 67, 75).

[CR14] Jonathan Cutler and A Jamie Radcliffe. "The maximum number of complete subgraphs in a graph with given maximum degree." In: *Journal of Combinatorial Theory, Series B* 104 (2014), pp. 60–71 (cit. on p. 81).

[CS18] Pietro Caputo and Alistair Sinclair. "Entropy production in nonlinear recombination models." In: *Bernoulli* (2018), pp. 3246–3282 (cit. on pp. 19, 79).

[CS24a] Pietro Caputo and Justin Salez. “Entropy factorization via curvature.” In: *arXiv preprint arXiv:2407.13457* (2024) (cit. on p. 8).

[CS24b] Pietro Caputo and Alistair Sinclair. “Nonlinear Dynamics for the Ising Model.” In: *Communications in Mathematical Physics* 405.11 (2024) (cit. on pp. 7, 9, 10, 19, 20, 23, 25, 28, 34, 124).

[CSV25] Zongchen Chen, Daniel Stefankovic, and Eric Vigoda. *Spectral Independence and Local-to-Global Techniques for Optimal Mixing of Markov Chains*. 2025. arXiv: 2307.13826 [cs.DM]. URL: https://arxiv.org/abs/2307.13826 (cit. on p. 103).

[DGR09] Enrico Dolera, Elisabetta Gabetta, and Eugenio Regazzini. “Reaching the Best Possible Rate of Convergence to Equilibrium for Solutions of Kac’s Equation via Central Limit Theorem.” In: *Annals of Applied Probability* 19 (2009), pp. 186–209 (cit. on p. 1).

[Dia+21] Ilias Diakonikolas, Daniel M. Kane, Alistair Stewart, and Yuxin Sun. “Outlier-Robust Learning of Ising Models Under Dobrushin’s Condition.” In: *arXiv preprint arXiv:2102.02171* (2021) (cit. on pp. 172, 173, 177).

[Dob68] Roland L. Dobrushin. “The description of a random field by means of conditional probabilities and conditions of its regularity.” In: *Theory of Probability and its Applications* 13 (1968), pp. 197–224 (cit. on p. 3).

[DP23] Ewan Davies and Will Perkins. “Approximately counting independent sets of a given size in bounded-degree graphs.” In: *SIAM Journal on Computing* 52.2 (2023), pp. 618–640 (cit. on pp. 8, 16, 17, 20, 133, 174, 183).

[DS96] P DIACONIS and L SALOFF-COSTE. “Logarithmic Sobolev inequalities for finite Markov chains.” In: *The Annals of Applied Probability* 6.3 (1996), pp. 695–750 (cit. on pp. 8, 92, 100, 103).

[Dum23] Hugo Duminil-Copin. “100 years of the (critical) Ising model on the hypercubic lattice.” In: *International Congress of Mathematicians*. European Mathematical Society-EMS-Publishing House GmbH. 2023, pp. 164–210 (cit. on p. 2).

[EFS20] Matthias Erbar, Max Fathi, and André Schlichting. “Entropic Curvature and Convergence to Equilibrium for Mean-Field Dynamics on Discrete Spaces.” In: *ALEA* 17 (2020), pp. 445–471 (cit. on p. 19).

[Fei19] Martin Feinberg. *Foundations of Chemical Reaction Network Theory*. Switzerland: Springer, 2019 (cit. on pp. 7, 14, 19).

[Fre75] David A Freedman. “On tail probabilities for martingales.” In: *the Annals of Probability* (1975), pp. 100–118 (cit. on p. 99).

[FV17] Sacha Friedli and Yvan Velenik. *Statistical mechanics of lattice systems: a concrete mathematical introduction*. Cambridge University Press, 2017 (cit. on p. 2).

[Geo11] Hans-Otto Georgii. *Gibbs measures and phase transitions*. Vol. 9. Walter de Gruyter, 2011 (cit. on p. 2).

[Gol89] David E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Reading, MA: Addison-Wesley, 1989 (cit. on p. 19).

[GŠV16] Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. "Inapproximability of the partition function for the antiferromagnetic Ising and hard-core models." In: *Combinatorics, Probability and Computing* 25.4 (2016), pp. 500–559 (cit. on pp. 2, 4).

[Hay06] Thomas P. Hayes. "A simple condition implying rapid mixing of single-site dynamics on spin systems." In: *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 2006, pp. 39–46 (cit. on p. 3).

[HS23] Jonathan Hermon and Justin Salez. "Modified log-Sobolev inequalities for strong-Rayleigh measures." In: *Annals of Applied Probability* 33.2 (2023), pp. 1501–1514 (cit. on pp. 8, 56).

[Jai+22] Vishesh Jain, Will Perkins, Ashwin Sah, and Mehtaab Sawhney. "Approximate Counting and Sampling via Local Central Limit Theorems." In: *STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing*. ACM, 2022, pp. 1473–1486. DOI: 10.1145/3519935.3519957 (cit. on pp. 126, 127).

[Jai+23] Vishesh Jain, Marcus Michelen, Huy Tuan Pham, and Thuy-Duong Vuong. "Optimal mixing of the down-up walk on independent sets of a given size." In: *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2023, pp. 1665–1681 (cit. on pp. 9, 16, 20, 75, 77, 78, 92, 97, 102, 105, 106, 117, 118, 126, 128).

[Jay57a] E.T. Jaynes. "Information Theory and Statistical Mechanics." In: *Physical Review* 106 (1957), pp. 620–630 (cit. on p. 12).

[Jay57b] E.T. Jaynes. "Information Theory and Statistical Mechanics II." In: *Physical Review* 108 (1957), pp. 171–190 (cit. on p. 12).

[Jer95] Mark Jerrum. "A very simple algorithm for estimating the number of k-colourings of a low-degree graph." In: *Random Structures & Algorithms* 7.2 (1995), pp. 157–165 (cit. on p. 119).

[JPV26] Vishesh Jain, Huy Tuan Pham, and Thuy-Duong Vuong. "Entropic independence via sparse localization." In: *arXiv preprint arXiv:2604.10902* (2026) (cit. on pp. 9, 16, 20, 77, 92, 94, 117, 118).

[Kac56] Mark Kac. "Foundations of kinetic theory." In: *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability*. Vol. 3. University of California Press, 1956, pp. 171–197 (cit. on pp. 1, 8, 9, 15, 19, 25, 114).

[Kac59] Mark Kac. *Probability and Related Topics in Physical Sciences*. Vol. 1. American Mathematical Society, 1959 (cit. on p. 1).

[Kap89] J.N. Kapur. *Maximum-Entropy Models in Science and Engineering*. Wiley, New York, 1989 (cit. on p. 12).

[Kel85] Frank P Kelly. "Stochastic models of computer communication systems." In: *Journal of the Royal Statistical Society Series B: Statistical Methodology* 47.3 (1985), pp. 379–395 (cit. on pp. 2, 4).

[LP17] David A Levin and Yuval Peres. *Markov chains and mixing times*. Vol. 107. American Mathematical Soc., 2017 (cit. on p. 28).

[Mar97] Fabio Martinelli. "Lectures on Glauber Dynamics for Discrete Spin Models." In: *Lectures on Probability Theory and Statistics: Saint-Flour 1997*. Springer, 1997, pp. 93–191 (cit. on pp. 2, 3, 8, 71).

[McK66] Henry P. McKean. "A Class of Markov Processes Associated with Nonlinear Parabolic Equations." In: *Proceedings of the National Academy of Sciences* 56.6 (1966), pp. 1907–1911 (cit. on p. 19).

[McK67] Henry P. McKean. "An exponential formula for solving Boltzmann's equation for a Maxwellian gas." In: *Journal of Combinatorial Theory* 2.3 (1967), pp. 358–382 (cit. on p. 1).

[Mic97] Laurent Miclo. "Remarques sur l'hypercontractivité et l'évolution de l'entropie pour des chaines de Markov finies." In: *Séminaire de Probabilités XXXI*. Vol. 1655. Lecture Notes in Mathematics. Springer, 1997, pp. 136–167 (cit. on p. 118).

[Min+26] Mehrad Abbaszadeh Minab, Pietro Caputo, Zongchen Chen, Mario Morellini, and Alistair Sinclair. "Nonlinear Exchange Dynamics for Independent Sets." In: *arXiv preprint arXiv:2607.29016* (2026) (cit. on pp. 4, 7, 14, 15, 19, 72, 73).

[Mit98] Melanie Mitchell. *An Introduction to Genetic Algorithms*. Cambridge, MA: MIT Press, 1998 (cit. on p. 19).

[MM13] Stéphane Mischler and Clément Mouhot. "Kac's program in kinetic theory." In: *Inventiones Mathematicae* 193.1 (2013), pp. 1–147 (cit. on pp. 8, 19).

[Mon15] Andrea Montanari. "Computational implications of reducing data to sufficient statistics." In: *Electronic Journal of Statistics* 9 (2015), pp. 2370–2390 (cit. on p. 12).

[MSP22] Santiago Mazuelas, Yuan Shen, and Aritz Pérez. "Generalized maximum entropy for supervised classification." In: *IEEE Transactions on Information Theory* 68.4 (2022), pp. 2530–2550 (cit. on p. 12).

[MT06] Ravi Montenegro and Prasad Tetali. "Mathematical aspects of mixing times in Markov chains." In: *Foundations and Trends® in Theoretical Computer Science* 1.3 (2006), pp. 237–354 (cit. on p. 92).

[MWW09] Elchanan Mossel, Dror Weitz, and Nicholas Wormald. "On the hardness of sampling independent sets beyond the tree threshold." In: *Probability Theory and Related Fields* 143.3 (2009), pp. 401–439 (cit. on p. 11).

[NZB17] H. C. Nguyen, R. Zecchina, and J Berg. "Inverse statistical problems: from the inverse Ising problem to data science." In: *Advances in Physics* 66.3 (2017), pp. 197–261. DOI: 10.1080/00018732.2017.1341604 (cit. on p. 12).

[OSS11] S. Oveis-Gharan, A. Saberi, and M. Singh. "A randomized rounding approach to the traveling salesman problem." In: *Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science*. FOCS '11. Palm Springs, CA: IEEE, 2011, pp. 550–559 (cit. on p. 12).

[RR19] Gesine Reinert and Nathan Ross. "Approximating stationary distributions of fast mixing Glauber dynamics, with applications to exponential random graphs." In: (2019) (cit. on p. 105).

[RRS98] Yuval Rabani, Yury Rabinovich, and Alistair Sinclair. "A computational view of population genetics." In: *Random Structures & Algorithms* 12.4 (1998), pp. 313–334 (cit. on pp. 2, 10, 19).

[RSW92] Yury Rabinovich, Alistair Sinclair, and Avi Wigderson. "Quadratic dynamical systems." In: *Proceedings of the 33rd Annual Symposium on Foundations of Computer Science*. 1992, pp. 304–313. DOI: 10.1109/SFCS.1992.267761 (cit. on pp. 2, 7, 19).

[Sly10] Allan Sly. "Computational transition at the uniqueness threshold." In: *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*. IEEE. 2010, pp. 287–296 (cit. on pp. 2, 4).

[SS12] Allan Sly and Nike Sun. "The computational hardness of counting in two-spin models on d-regular graphs." In: *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*. IEEE. 2012, pp. 361–369 (cit. on pp. 4, 133, 137).

[SV14a] Mohit Singh and Nisheeth Vishnoi. "Entropy, Optimization and Counting." In: *Proceedings of the 46th ACM Symposium on Theory of Computing (STOC)*. 2014, pp. 50–59 (cit. on p. 19).

[SV14b] Mohit Singh and Nisheeth K. Vishnoi. "Entropy, optimization and counting." In: *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*. STOC '14. New York, NY: Association for Computing Machinery, 2014, pp. 50–59. DOI: 10.1145/2591796.2591803 (cit. on p. 12).

[SV19] Damian Straszak and Nisheeth K Vishnoi. "Maximum entropy distributions: Bit complexity and stability." In: *Conference on Learning Theory*. PMLR. 2019, pp. 2861–2891 (cit. on p. 12).

[Vil02] Cédric Villani. “A Review of Mathematical Topics in Collisional Kinetic Theory.” In: *Handbook of Mathematical Fluid Dynamics*. Vol. 1. 2002, pp. 71–305 (cit. on pp. 1, 9, 19).

[Vil03] Cédric Villani. “Cercignani’s Conjecture is Sometimes True and Always Almost True.” In: *Communications in Mathematical Physics* 234.3 (2003), pp. 455–490 (cit. on pp. 1, 8, 19).

[Wei06] Dror Weitz. “Counting independent sets up to the tree threshold.” In: *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*. 2006, pp. 140–149 (cit. on pp. 2, 4, 11, 103).

[Wil51] Edmond Wild. “On Boltzmann’s equation in the kinetic theory of gases.” In: *Mathematical Proceedings of the Cambridge Philosophical Society* 47.3 (1951), pp. 602–615 (cit. on p. 1).

[WJ08] Martin J. Wainwright and Michael I. Jordan. “Graphical Models, Exponential Families, and Variational Inference.” In: *Foundations and Trends in Machine Learning* 1.1–2 (2008), pp. 1–305 (cit. on pp. 12, 172, 173).

[Yau96] Horng-Tzer Yau. “Logarithmic Sobolev Inequality for Lattice Gases with Mixing Conditions.” In: *Communications in Mathematical Physics* 181.2 (1996), pp. 367–408 (cit. on pp. 19, 92).

Appendix A

Proof of Lemma 2.34 and Lemma 2.35

For all $\sigma \in \Omega_{N,\alpha}$, define

$$\mathcal{M}_{\geq 2}[\sigma] := \{(j,k) \in \{2,\dots,N\} \times [n] : N^+((j,k)) \cap \sigma = \emptyset\}, \tag{A.1}$$

and denote by $M_{\geq 2}[\sigma]$ the cardinality of $\mathcal{M}_{\geq 2}[\sigma]$.

Proof of Lemma 2.34

First, assume that (η_0, η_0') are given by the following rule:

0) Sample $\bar{\eta} \sim \nu_\ell$.

1) Select uniformly at random two pairs (i,ℓ') and (j,k) in $[N] \times [n]$.

2) If either (i,ℓ') or (j,k) is equal to $u := (1,\ell)$, define

$$\eta_0 := \bar{\eta}, \qquad \eta_0' := \bar{\eta}^{i,\ell';j,k}.$$

If both (i,ℓ') and (j,k) are different from u, there are two cases:

i) One of i or j is equal to 1; then define $\eta_0 = \eta_0' = \bar{\eta}$.

ii) Both i and j are different from 1; then define $\eta_0 = \eta_0' = \bar{\eta}^{i,\ell';j,k}$.

We call Γ_0 the above coupling of (η_0, η_0'). Note that $\eta_0 \sim \nu_\ell$, and $\eta_0' \sim \nu_\ell P_{\nu,N}$. Therefore,

$$\left|\mathbb{E}_{\nu_\ell}[P_{\nu_\ell,N}h - P_{\nu,N}h]\right| \leq \frac{2}{Nn}\left|\mathbb{E}\,\mathbb{E}_{(\eta_0,\eta_0')\sim\Gamma_0}\Big[\sum_{t=0}^{\infty}\left(F(\eta_t) - F(\eta_t')\right) \mid \eta_0 \neq \eta_0'\Big]\right|, \tag{A.2}$$

where we used the bound $\mathbb{P}[\eta_0 \neq \eta_0'] \leq \frac{2}{nN}$, and for a fixed realization of η_0, η_0' we use an arbitrary coupling of η_t, η_t'.

On the event $\{\eta_0 \neq \eta_0'\}$, there exists (j,k) with $j \geq 2$ such that $(\eta_0)_k(j) = 0$ and $(\ell,k)_{\eta_0(1),\eta_0(j)}$ is satisfied; in words, (j,k) is an empty and available vertex. According to the definition given in (A.1), the law $L \in \mathcal{P}(\Omega_{N,\alpha}^2)$ of $(\eta_0,\eta_0') \sim \Gamma_0$ conditioned on $\{\eta_0 \neq \eta_0'\}$ is given by

$$L(\eta_0,\eta_0') = \nu_\ell(\eta_0) \sum_{\substack{j\geq 2\\ k\in[n]}} \frac{\mathbf{1}\{(j,k) \in \mathcal{M}_{\geq 2}[\eta_0]\}}{M_{\geq 2}[\eta_0]} \mathbf{1}\{\eta_0' = \eta_0^{1,\ell;j,k}\}.$$

We may compare the measures L and Γ. For all $\eta_0 \in \Omega_{N,\alpha}$ such that $\eta_0(1) = \tau^\ell$ and all $(j,k) \in \mathcal{M}_{\geq 2}[\eta_0]$, by definition one has

$$L(\eta_0,\eta_0^{1,\ell;j,k}) \lesssim_{\gamma,\delta,\Delta} \Gamma(\eta_0,\eta_0^{1,\ell;j,k}) \quad \Longleftarrow \quad \frac{\nu_\ell(\eta_0)}{M_{\geq 2}[\eta_0]} \lesssim_{\gamma,\delta,\Delta} \frac{\bar\mu_{N,\alpha}(\eta_0)}{(N-1)n}. \tag{A.3}$$

where we used the short-hand $\bar\mu_{N,\alpha} = \mu_{N,\alpha}(\cdot \mid \sigma(1) = \tau^\ell,\, \sigma(j)_{k'} = 0,\, \forall k' \in N^+(k))$.

We now focus on proving

$$\frac{\nu_\ell(\eta_0)}{M_{\geq 2}[\eta_0]} \lesssim_{\gamma,\delta,\Delta} \frac{\bar\mu_{N,\alpha}(\eta_0)}{(N-1)n}. \tag{A.4}$$

Recalling the definition

$$\alpha^\tau = \frac{N}{N-1}\alpha - \frac{|\tau|}{(N-1)n},$$

we have

$$\begin{aligned} M_{\geq 2}[\eta_0] &\geq (N-1)n - (N-1)n\,\alpha^{\tau^\ell}(\Delta+1) \\ &= (N-1)n\Big(1 - \frac{N}{N-1}\alpha(\Delta+1) + \frac{|\tau|(\Delta+1)}{(N-1)n}\Big) \\ &\geq (N-1)n(1 - \frac{N\alpha(1+\Delta)}{N-1} + \frac{1}{N-1}) \\ &= (N-1)n(\frac{N(1-\alpha(1+\Delta))}{N-1}) \end{aligned} \tag{A.5}$$

and using $\alpha \leq (1-\delta)\alpha_c(\Delta)$ together with $\alpha_c(\Delta) \leq 0,74/\Delta$ and $\Delta \geq 3$, we have

$$\begin{aligned} 1 - \alpha(\Delta+1) &\geq 1 - 0,74(1-\delta)(\Delta+1)/\Delta \\ &= 1 - 0,74(1-\delta) - 0,74(1-\delta)/\Delta \\ &= 0,26 + 0,74\delta - 0,74(1-\delta)/\Delta \\ &\geq 0,26 + 0,74\delta - 0,74/3 \geq 1/100 \end{aligned}$$

which returns $M_{\geq 2}[\eta_0] \geq n(N-1)/100$.

Our goal is to show that

$$R(\ell,\eta_0,\alpha) = \frac{\nu_\ell(\eta_0)}{\bar\mu_{N,\alpha}(\eta_0)} = O_{\gamma,\delta,\Delta}(1). \tag{A.6}$$

Now fix a subcritical fugacity λ such that $\alpha(\mu_{\lambda,G}) = \alpha$. Then

$$\begin{aligned}\bar{\mu}_{N,\alpha}(\eta_0) &= \frac{\mu_{N,\alpha}(\sigma = \eta_0)}{\mu_{N,\alpha}(\sigma(1) = \tau^\ell,\, \sigma(j)_{k'} = 0,\, \forall k' \in N^+(k))} \\ &= \frac{\mu_{\lambda,G}^{\otimes N}(\sigma = \eta_0)}{\mu_{\lambda,G}^{\otimes N}(\sigma(1) = \tau^\ell,\, \sigma(j)_{k'} = 0,\, \forall k' \in N^+(k),\, \sigma \in \Omega_{N,\alpha})},\end{aligned}$$

and similarly,

$$\nu_\ell(\eta_0) = \frac{\mu_{\lambda,G}^{\otimes N}(\sigma = \eta_0)}{\mu_{\lambda,G}^{\otimes N}(\sigma(1) = \tau^\ell,\, \sigma \in \Omega_{N,\alpha},)}.$$

We have,

$$\begin{aligned}R(\ell, \eta_0, \alpha) &= \frac{\mu_{\lambda,G}^{\otimes N}(\sigma(1) = \tau^\ell,\, \sigma(j)_{k'} = 0,\, \forall k' \in N^+(k),\, \sigma \in \Omega_{N,\alpha})}{\mu_{\lambda,G}^{\otimes N}(\sigma(1) = \tau^\ell,\, \sigma \in \Omega_{N,\alpha})} \\ &= \sum_{\eta : \eta_{N^+(k)} = 0} \mu_{\lambda,G}(\eta) \frac{\mu_{\lambda,G}^{\otimes N-2}\left(\Omega_{N-2,\alpha^{\tau^\ell,\eta}}\right)}{\mu_{\lambda,G}^{\otimes N-1}\left(\Omega_{N-1,\alpha^{\tau^\ell}}\right)}\end{aligned}$$

where $\alpha^{\tau^\ell,\eta} = \frac{N}{N-2}\alpha - \frac{|\tau^\ell|+|\eta|}{(N-2)n}$ and $\alpha^{\tau^\ell} = \frac{N}{N-1}\alpha - \frac{|\tau^\ell|}{(N-1)n}$. By Proposition 2.28 and since $N \geq n$ since uniformly in $\tau, \eta \in \Omega_G$ holds that

$$|\alpha^{\tau,\eta} - \alpha|^2 (N-2)n = \Big|\frac{|\tau| + |\eta|}{2n} - \alpha\Big|^2 \frac{4n}{(N-2)} = O(1),$$

then, $\frac{\mu_{\lambda,G}^{\otimes N-2}\left(\Omega_{N-2,\alpha^{\tau,\eta}}\right)}{\mu_{\lambda,G}^{\otimes N-1}\left(\Omega_{N,\alpha^{\tau}}\right)} = \Theta_{\gamma,\delta,\Delta}(1)$. We finally obtain

$$R(\ell, \eta_0, \alpha) \sim_{\gamma,\delta,\Delta} \mu_{\lambda,G}(\sigma_{N_+(k)} = 0) = O_{\gamma,\delta,\Delta}(1).$$

This ends the proof of (A.4)

It is clear that (A.2) together with (A.3) yields

$$\big|\mathbb{E}_{\nu_\ell}[P_{\nu_\ell,N} h - P_{\nu,N} h]\big| \lesssim_{\gamma,\delta,\Delta} \frac{1}{Nn}\Big|\sum_{t=0}^{\infty} \mathbb{E}\, \mathbb{E}_{(\eta_0,\eta_0')\sim\Gamma}\Big[\big(F(\eta_t) - F(\eta_t')\big)\Big]\Big|, \tag{A.7}$$

To complete the proof, we show how (A.7) implies (2.56). Define

$$\xi_t := \mathbb{E}\mathbb{E}_{(\eta_0,\eta_0')\sim\Gamma}\big[F(\eta_t) + F(\eta_t')\big].$$

Note that, uniformly in t,

$$\xi_t = \mathbb{E}\mathbb{E}_{(\eta_0,\eta_0')\sim\Gamma}[F(\eta_t)] + \mathbb{E}\mathbb{E}_{(\eta_0,\eta_0')\sim\Gamma}[F(\eta_t')] \lesssim_{\gamma,\delta,\Delta} \mathbb{E}_\nu[F], \tag{A.8}$$

since $P_{\nu,N}$ is reversible with respect to ν and the marginals of Γ can be estimated from

above with a constant multiple of ν (with constants depending only on γ, δ, Δ), which follows from the local central limit theorem similar to (A.6).

Let $c > 0$ be a positive constant. Using the Cauchy–Schwarz inequality, (A.7), we have

$$\begin{aligned}\frac{\big|\mathbb{E}_{\nu_\ell}[P_{\nu_\ell,N}h - P_{\nu,N}h]\big|^2}{\mathbb{E}_\nu[F]} &\lesssim_{\gamma,\delta,\Delta} \frac{1}{(Nn)^2\mathbb{E}_\nu[F]}\Big|\sum_{t=0}^{\infty}\mathbb{E}\mathbb{E}_{(\eta_0,\eta_0')\sim\Gamma}[F(\eta_t) - F(\eta_t')]\Big|^2 \\ &\le \frac{1}{(Nn)^2\mathbb{E}_\nu[F]}\Big[\sum_{t=0}^{\infty} e^{-ct/nN}\xi_t\Big] \\ &\quad\times\Big[\sum_{t=0}^{\infty} e^{ct/nN}\frac{\big|\mathbb{E}\mathbb{E}_{(\eta_0,\eta_0')\sim\Gamma}[F(\eta_t) - F(\eta_t')]\big|^2}{\xi_t}\Big].\end{aligned}$$

Using (A.8), we obtain

$$\sum_{t=0}^{\infty} e^{-ct/nN}\xi_t \lesssim_{\gamma,\delta,\Delta} \frac{nN}{c}\,\mathbb{E}_\nu[F].$$

Next, by another application of the Cauchy–Schwarz inequality,

$$\begin{aligned}\big|\mathbb{E}\mathbb{E}_{(\eta_0,\eta_0')\sim\Gamma}[F(\eta_t) - F(\eta_t')]\big|^2 &\le \mathbb{E}\mathbb{E}_{(\eta_0,\eta_0')\sim\Gamma}\Big[\big(\sqrt{F(\eta_t)} - \sqrt{F(\eta_t')}\big)^2\Big] \\ &\quad\times \mathbb{E}\mathbb{E}_{(\eta_0,\eta_0')\sim\Gamma}\Big[\big(\sqrt{F(\eta_t)} + \sqrt{F(\eta_t')}\big)^2\Big] \\ &\le 2\,\xi_t\,\mathbb{E}\mathbb{E}_{(\eta_0,\eta_0')\sim\Gamma}\Big[\big(\sqrt{F(\eta_t)} - \sqrt{F(\eta_t')}\big)^2\Big].\end{aligned}$$

Combining the above estimates, we conclude that

$$\frac{\big|\mathbb{E}_{\nu_\ell}[P_{\nu_\ell,N}h - P_{\nu,N}h]\big|^2}{\mathbb{E}_\nu[F]} \lesssim_{\gamma,\delta,\Delta} \frac{1}{cNn}\sum_{t=0}^{\infty} e^{ct/nN}\mathbb{E}\mathbb{E}_{(\eta_0,\eta_0')\sim\Gamma}\Big[\big(\sqrt{F(\eta_t)} - \sqrt{F(\eta_t')}\big)^2\Big],$$

which proves (2.56). This ends the proof of Lemma 2.34. □

Proof of Lemma 2.35

We recall the definition in (2.55)

$$H(c,\Gamma) := \sum_{t=0}^{\infty} e^{ct/nN}\mathbb{E}\Big[\Big(\sqrt{F(\eta_t)} - \sqrt{F(\eta_t')}\Big)^2\Big].$$

The quantity $H(c,\Gamma)$ depends on the specific choice of the coupling between η_t and η_t'. The inequality stated in Lemma 2.57 follows from a suitable construction of such a coupling. It is sufficient to specify the coupling at time $t = 1$ and only for pairs of independent sets $\eta_0, \eta_0' \in \mathrm{supp}(\Gamma)$. The extension to all remaining cases can then be obtained by standard arguments, for instance by applying the Gluing Lemma.

Let $\eta_0, \eta_0' \in \mathrm{supp}(\Gamma)$. By definition, the pair satisfies $\eta_0 \sim \eta_0'$. Denoting $u = (1,\ell)$, we have $\eta_0(1)_\ell = 1$ and there exists a site $v = (i_v, \ell_v)$ with $\ell_v > 1$ such that $\eta_0' = \eta_0^{(u;v)}$. We

now describe the coupling.

1) Select uniformly at random two pairs (i, ℓ') and (j, k).

2) Perform the following update, with the convention that the exchange is carried out only if it produces two admissible independent sets; otherwise, the update is rejected.

   i) If $(i, \ell') = u$, set
   $$\eta_1 = \eta_0^{(u;j,k)}, \qquad \eta_1' = (\eta_0')^{(v;j,k)}.$$
   If $(i, \ell') = v$, set
   $$\eta_1 = \eta_0^{(v;j,k)}, \qquad \eta_1' = (\eta_0')^{(u;j,k)}.$$
   By symmetry, the same definitions apply when (j, k) coincides with u or v.

   ii) If both (i, ℓ') and (j, k) are different from u and v, we distinguish the following cases:

      A) If either $i = 1$ or $j = 1$, reject the update and set
      $$\eta_1 = \eta_0, \qquad \eta_1' = \eta_0'.$$

      B) If both i and j are different from 1, set
      $$\eta_1 = \eta_0^{(i,\ell';j,k)}, \qquad \eta_1' = (\eta_0')^{(i,\ell';j,k)}.$$

We now prove that if $\alpha \le \frac{1}{C\Delta^2}$, for $C \ge 241$ and $c \le 1/1000$

$$H(c, \Gamma) \le Nn\, \mathbb{E}_{(\eta_0, \eta_0') \sim \Gamma}[(\sqrt{F(\eta_0)} - \sqrt{F(\eta_0')})^2]. \tag{A.9}$$

The inequality (A.9) is obtained by analyzing the evolution at time $t = 1$, which leads to a recursive inequality for $H(c, \Gamma)$, see (A.11) below. Consider therefore the coupled pair (η_1, η_1') at time $t = 1$.

Exactly one of the following three situations may occur:

(1) $\eta_1 = \eta_1'$;

(2) $\eta_1 = (\eta_1')^{(u;v)}$;

(3) $\eta_1 = \eta_0^{(w;w')}$ and $\eta_1' = \eta_0'$, where $w = (i_w, \ell_w)$ and $w' = (i_{w'}, \ell_{w'})$ belong to $\{2, \dots, N\} \times [n]$, with $i_w = i_v$, $\ell_w \in N(\ell_v)$, and $(\eta_0)_w = (\eta_0)_{w'} = 1$.

The third case corresponds to subcase (B) in the coupling construction, where the update is admissible for η_0 but rejected for η_0'. This occurs because the presence of a particle at v prevents the insertion of a particle at a neighboring site in η_0'.

Let p_1, p_2, and p_3 denote the probabilities of cases (1), (2), and (3), respectively. Observe that, on the event corresponding to the second case, the law of (η_1, η_1') coincides with Γ.

Let Γ' denote the law of (η_1, η_1') conditioned on the third case. We then obtain

$$H(c,\Gamma) = \mathbb{E}_{(\eta_0,\eta_0')\sim\Gamma}\Big[\big(\sqrt{F(\eta_0)} - \sqrt{F(\eta_0')}\big)^2\Big] + p_2 e^{c/nN} H(c,\Gamma) + p_3 e^{c/nN} H(c,\Gamma').$$

The key estimate needed to close the argument is the following comparison result.

Proposition A.1. *Suppose that $\alpha \le \frac{1}{C\Delta^2}$ with $C \ge 5$. Then*

$$H(c,\Gamma') \le 120\Delta\, H(c,\Gamma). \tag{A.10}$$

Assuming Proposition A.1, we can complete the proof of (A.9). Indeed, combining the previous identity with (A.10) yields

$$\begin{aligned} H(c,\Gamma) \le\ & \mathbb{E}_{(\eta_0,\eta_0')\sim\Gamma}\Big[\big(\sqrt{F(\eta_0)} - \sqrt{F(\eta_0')}\big)^2\Big] \\ & + e^{c/nN}\big(1 - p_1 + (120\Delta - 1)p_3\big) H(c,\Gamma). \end{aligned} \tag{A.11}$$

Proof of (A.9) assuming Proposition A.1

Assume that the density satisfies $\alpha \le \frac{1}{C\Delta^2}$ with $C \ge 241$. We estimate the probabilities p_1, p_2, and p_3 and substitute the resulting bounds into (A.11).

The probability p_3 of the third case satisfies

$$\frac{\alpha^{\tau^\ell}}{nN} \le p_3 \le \frac{\alpha^{\tau^\ell}\Delta}{Nn}.$$

The probability p_1 of the first case satisfies

$$\frac{2}{Nn} \ge p_1 \ge \frac{1}{Nn}(1 - \alpha^{\tau^\ell}\Delta).$$

Indeed, this event corresponds to two possible choices of $(i,\ell) \in \{u,v\}$ and to a choice of (j,k) among admissible vertices:

$$\begin{aligned} \frac{2}{Nn} \ge p_1 &= \frac{2}{Nn}\Big(\mathbb{P}[(j,k) \in \mathcal{M}_{\ge 2}[\eta_0]] + \mathbb{P}[(j,k) : (\eta_0)_k(j) = 1]\Big) \\ &\ge \frac{2}{nN}\Big(\frac{N-1}{N}(1 - \alpha^{\tau^\ell}(\Delta+1)) + \frac{N-1}{N}\alpha^{\tau^\ell}\Big) = \frac{1}{nN}(1 - \alpha^{\tau^\ell}\Delta), \end{aligned}$$

where $\mathcal{M}_{\ge 2}[\eta_0]$ is defined in (A.1).

As a consequence, since $p_2 = 1 - p_1 - p_3$, we obtain

$$1 - \frac{1}{nN}(2 + \alpha^{\tau^\ell}) \le p_2 \le 1 - \frac{1}{nN}(1 - \alpha^{\tau^\ell}(\Delta - 1)).$$

If $\alpha \le 1/(C\Delta^2)$ with $C \ge 241$, then

$$p_1 - (120\Delta - 1)p_3 \ge \frac{1}{2}\Big(1 - \frac{1}{1000}\Big)\frac{1}{Nn}.$$

In the end, if $c < 1/1000$, then

$$e^{c/Nn}\Big(1 - \frac{1}{2}\Big(1 - \frac{1}{1000}\Big)\frac{1}{Nn}\Big) \leq 1 - \frac{1}{2Nn},$$

which yields the claim of Lemma (A.9).

It therefore remains to prove Proposition A.1.

Proof of Proposition A.1

Let $\eta_1, \eta_1' \in \Omega_{N,\alpha}$ be sampled according to Γ'. The construction starts from $\eta_0, \eta_0' \sim \Gamma$, with $\eta_0' = \eta_0^{(u;v)}$ for some vertex v such that $j_v \geq 2$. Choose uniformly at random a neighbor $w = (j_w, \ell_w)$ of v; by construction, $j_w = j_v$. Then choose a site $w' = (j_{w'}, \ell_{w'})$ such that $(\eta_0)_{\ell_{w'}}(j_{w'}) = 1$, and set $\eta_1 = \eta_0^{(w;w')}$ and $\eta_1' = \eta_0'$.

The discrepancy of the pair (η_1, η_1') is therefore increased by two with respect to that of (η_0, η_0'), i.e. $|\eta_1 \setminus \eta_1'| = 2$ while $|\eta_0 \setminus \eta_0'| = 1$, see figure A. However, the distribution Γ is supported only on pairs with the same discrepancy. Hence, in order to compare the distributions Γ and Γ', it is necessary to introduce auxiliary random variables that restore the correct discrepancy.

We now describe such a construction.

Let $z = (\ell_z, j_z)$, with $j_z \geq 2$, be chosen uniformly at random among the available empty sites of η_1; equivalently, $z \in \mathcal{M}_{\geq 2}(\eta_1')$ (see definition (A.1)).

Consider the following sequence of configurations:

$$\begin{aligned} \eta_1 \to X_0 = (\eta_1)^{(u;z)} \to X_1 = (X_0)^{(w;u)} \to X_2 = (X_1)^{(u;w')} \to \\ \to \eta_0 = (X_2)^{(u;z)} \to \eta_0' = (\eta_0)^{(u;v)} \to \eta_1' = (\eta_0)^{(u,v)}. \end{aligned} \tag{A.12}$$

Our claim is that the distributions $Law(\eta_1, X_0)$, $Law(X_1, X_0)$, $Law(X_1, X_2)$, and $Law(\eta_0, X_2)$ are all bounded above by $24\Delta\,\Gamma$. Once this claim is established, the conclusion follows immediately. Indeed, by the Cauchy–Schwarz inequality,

$$\begin{aligned} H(c, \Gamma') &\leq H\big(c, Law(\eta_1, X_0)\big) + H\big(c, Law(X_1, X_0)\big) + H\big(c, Law(X_1, X_2)\big) \\ &\quad + H\big(c, Law(\eta_0, \eta_1)\big) + H\big(c, Law(\eta_0, \eta_1')\big) \\ &\leq 120\Delta\, H(c, \Gamma). \end{aligned}$$

In what follows, we explicitly write the laws of the auxiliary random variables and compare each of them with Γ. The comparison relies on change-of-measure arguments based on the Local Central Limit Theorem (Proposition 2.28) and on the bounds provided by Proposition A.2. For clarity, it is useful to keep in mind the supports of the auxiliary random variables defined through the sequence (A.12).

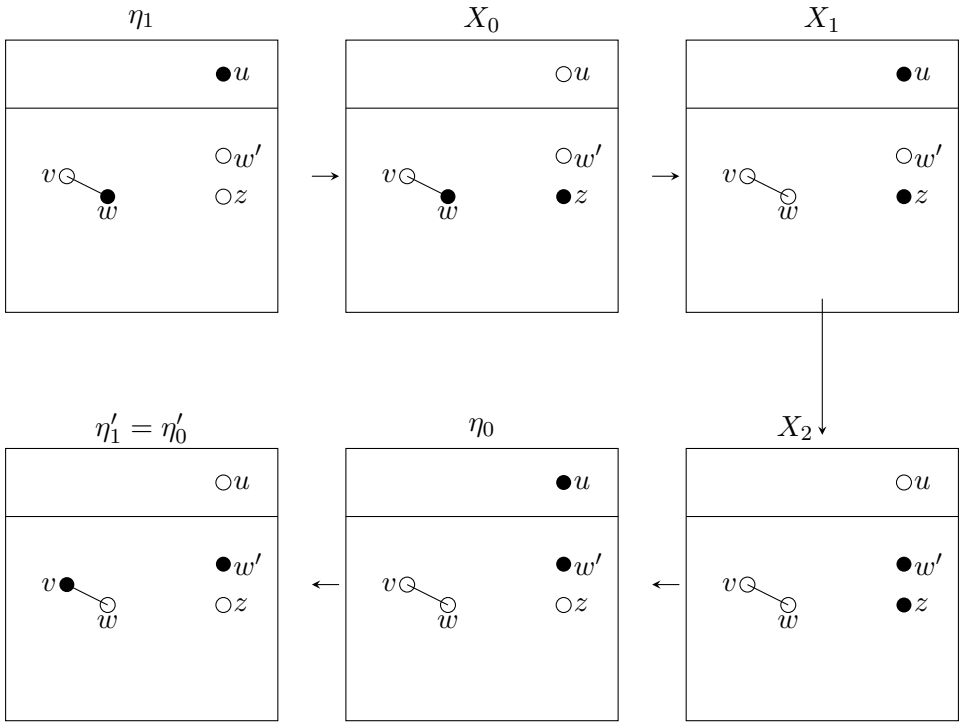


Figure A.1: Schematic representation of the sequence of independent sets appearing in (A.12). All configurations belong to the space $\Omega_{N,\alpha}$ and differ only by the occupation status of the randomly selected vertices v, w, w', z and fixed vertex $u = (1, \ell)$. Filled vertices represent occupied sites, while empty vertices represent unoccupied sites; edges between vertices indicate adjacency relations in the underlying graph. The top row, where the distinguished site u appears, represents the particle configuration $\sigma(1)$. Keeping this picture in mind is useful for the study of the comparison arguments; in particular, observe that the support of the measures involved is generated by varying the vertices v, w, w', z.

Law of η_0.

$$\mathbb{P}(\eta_0 = \sigma) = \frac{1}{(N-1)n} \sum_{v:\, j_v \geq 2} \mu^v_{N,\alpha}(\sigma),$$

where

$$\mu^v_{N,\alpha}(\sigma) = \mu_{N,\alpha}\big(\sigma \mid \sigma(1) = \tau^\ell,\, \sigma_{N^+(\ell_v)}(j_v) = 0\big).$$

The measure $\mu^v_{N,\alpha}$ leaves the site $v = (i_v, \ell_v)$ available. A key point in the comparison arguments below is that the choice of this available vertex can be modified without significantly perturbing the reference measure $\mu_{N,\alpha}$. We now make this statement precise.

Consider two vertices z' and z'' such that $j_{z'}, j_{z''} \geq 2$. We compare the corresponding conditional measures $\mu^{z'}_{N,\alpha}$ and $\mu^{z''}_{N,\alpha}$ and claim that

$$\frac{\mu^{z'}_{N,\alpha}}{\mu^{z''}_{N,\alpha}} \leq 12. \tag{A.13}$$

Let λ^{τ^ℓ} denote the fugacity of the hardcore model μ such that $\alpha(\mu) = \alpha^{\tau^\ell}$. Suppose

that $\sigma \in \mathrm{supp}(\mu^{z'}_{N,\alpha}) \cap \mathrm{supp}(\mu^{z''}_{N,\alpha})$. Then

$$\mu^{z'}_{N,\alpha}(\sigma) \in \frac{\mu^{\otimes N}(\sigma)}{\mu(\tau^\ell)\,\mu\Big(\sigma_{N^+(\ell_{z'})}=0\Big)} \left[\min_{s\in[n]} \mu^{\otimes(N-2)}\Big(\Omega_{N-2,\alpha^{\eta^\ell,s}}\Big)^{-1}, \min_{s\in[n]} \mu^{\otimes(N-2)}\Big(\Omega_{N-2,\alpha^{\eta^\ell,s}}\Big)^{-1}\right].$$

As a consequence,

$$\frac{\mu^{z'}_{N,\alpha}}{\mu^{z''}_{N,\alpha}} \le \max_{0\le s,s'\le n} \frac{\mu^{\otimes(N-2)}\Big(\Omega_{N-2,\alpha^{\eta^\ell,s}}\Big)}{\mu^{\otimes(N-2)}\Big(\Omega_{N-2,\alpha^{\eta^\ell,s'}}\Big)} \frac{\mu(\sigma_{N^+(\ell_{z'})}=0)}{\mu(\sigma_{N^+(\ell_{z''})}=0)}. \tag{A.14}$$

Here, for $s \in [n]$, we define

$$\alpha^{\tau^\ell,s} = \frac{N-1}{N-2}\alpha^{\tau^\ell} - \frac{s}{(N-2)n}.$$

By Proposition 2.28,

$$\frac{\mu^{\otimes(N-2)}\Big(\Omega_{N-2,\alpha^{\eta^\ell,s}}\Big)}{\mu^{\otimes(N-2)}\Big(\Omega_{N-2,\alpha^{\eta^\ell,s'}}\Big)} \le \exp\left(-\frac{(N-2)n\big(|\alpha^{\eta^\ell}-\alpha^{\eta^\ell,s}|^2 - |\alpha^{\eta^\ell}-\alpha^{\eta^\ell,s'}|^2\big)}{2v^2_{N-2}}\right) + \frac{C_{\gamma,\delta,\Delta}}{\sqrt{(N-2)n}}$$

$$\le 2,$$

for $N \ge n$ sufficiently large.

Finally, by Proposition A.2 (see below), we obtain

$$\frac{\mu^{z'}_{N,\alpha}}{\mu^{z''}_{N,\alpha}} \le 2 \max_{k,k'\in[n]} \frac{\mu(\sigma_{N^+(k)}=0)}{\mu(\sigma_{N^+(k')}=0)} \le 2\Big[(1+\lambda^{\tau^\ell})^{\Delta-1}+1\Big]. \tag{A.15}$$

By Lemma 2.12, the parameters λ^{τ^ℓ} and α^{τ^ℓ} satisfy $\alpha^{\tau^\ell} \ge \frac{\lambda^{\tau^\ell}}{1+\lambda^{\tau^\ell}(1+\Delta)}$, since $\alpha(\mu_{G,\lambda^{\tau^\ell}}) = \alpha^{\tau^\ell}$. Moreover, if $\alpha \le 1/(C\Delta^2)$, then $\alpha^{\tau^\ell} \le \frac{1}{N-1} + \frac{1}{C\Delta^2} \le \frac{1}{(C-1)\Delta^2}$ for $N \ge C(C-1)\Delta$. In particular, $\lambda^{\tau^\ell} \le \frac{3}{(C-1)\Delta^2}$. Substituting this bound into (A.15) yields (A.13).

We now proceed to the analysis of the law of η_1.

Law of η_1. By definition, η_1 is obtained from η_0 by moving a particle from a site w' to a site w, where w is chosen uniformly at random from the neighborhood of v. Throughout the remainder of the proof, we implicitly assume that $j_v, j_w, j_{w'} \ge 2$, $j_w = j_v$, and $\ell_w \in N(\ell_v)$, that is, $w \in N(v)$; for simplicity, we do not repeat these assumptions every time.

$$\begin{aligned}\mathbb{P}(\eta_1=\sigma) &= \frac{1}{(N-1)n}\sum_{\sigma',v}\mu^v_{N,\alpha}(\sigma')\sum_{w'}\frac{\mathbf{1}(\sigma'_{\ell_{w'}}(j_{w'})=1)}{(N-1)n\alpha^{\tau^\ell}}\sum_{w\in N(v)}\frac{\mathbf{1}(\sigma=(\sigma')^{(w;w')})}{|N(v)|}\\ &= \frac{1}{(N-1)n}\sum_{v,w,w'}\frac{\mu^v_{N,\alpha}(\sigma^{(w;w')})\mathbf{1}(\sigma_{N^+(\ell_{w'})}=0,\sigma_{\ell_w}(j_w)=1)}{|N(v)|(N-1)n\alpha^{\tau^\ell}}.\end{aligned}$$

Law of (η_1, X_0). Given η_1, v, w, w', the configuration X_0 is obtained by selecting uniformly a vertex z among the available sites in $\mathcal{M}_{\geq 2}(\eta_1)\setminus(N^+(w')\cup N^+(v))$, and setting $X_0 = \eta_1^{(u;z)}$. Thus,

$$\mathbb{P}(\eta_1 = \sigma,\, X_0 = \sigma') = \sum_{v,w,w',z} \frac{\mu_{N,\alpha}^v(\sigma^{(w;w')})\mathbf{1}(\sigma_{N^+(\ell_{w'})} = 0, \sigma_{\ell_w}(j_w) = 1)}{|N(v)|(N-1)^2 n^2 \alpha^{\tau^\ell}} \times \frac{\mathbf{1}(z \in \mathcal{M}_{\geq 2}(\sigma) \setminus (N^+(w') \cup N^+(v)))}{|\mathcal{M}_{\geq 2}(\sigma) \setminus (N^+(w') \cup N^+(v))|} \mathbf{1}(\sigma' = \sigma^{(u;z)}).$$

Law of X_0. Taking the second marginal of the previous law,

$$\mathbb{P}(X_0 = \sigma) = \sum_{v,w,w',z} \frac{\mu_{N,\alpha}^v((\sigma^{(u;z)})^{(w;w')})\mathbf{1}(\sigma_{N^+(\ell_{w'})} = 0, \sigma_{\ell_w}(j_w) = 1)}{|N(v)|(N-1)^2 n^2 \alpha^{\tau^\ell}} \times \frac{\mathbf{1}(z \in \mathcal{M}_{\geq 2}((\sigma^{(u;z)})) \setminus (N^+(w') \cup N^+(v)))}{|\mathcal{M}_{\geq 2}((\sigma^{(u;z)})) \setminus (N^+(w') \cup N^+(v))|}.$$

Law of (X_0, X_1). Given v, w, w', z and X_0, define $X_1 = X_0^{(w;u)}$. Then

$$\mathbb{P}(X_0 = \sigma, X_1 = \sigma') = \sum_{v,w,w',z} \frac{\mu_{N,\alpha}^v((\sigma^{(u;z)})^{(w;w')})\mathbf{1}(\sigma_{N^+(\ell_{w'})} = 0, \sigma_{\ell_w}(j_w) = 1)}{|N(v)|(N-1)^2 n^2 \alpha^{\tau^\ell}} \times \frac{\mathbf{1}(z \in \mathcal{M}_{\geq 2}((\sigma^{(u;z)})) \setminus (N^+(w') \cup N^+(v)))}{|\mathcal{M}_{\geq 2}((\sigma^{(u;z)})) \setminus (N^+(w') \cup N^+(v))|} \mathbf{1}(\sigma' = \sigma^{(w;u)})$$

Law of X_1.

$$\mathbb{P}(X_1 = \sigma') = \sum_{v,w,w',z} \frac{\mu_{N,\alpha}^v(\sigma^{z,w'})\mathbf{1}(\sigma_{N^+(w')} = 0, \sigma_{N(w)} = 0, \sigma_\ell(1) = 1)}{|N(v)|(N-1)^2 n^2 \alpha^{\tau^\ell}} \times \frac{\mathbf{1}(z \in \mathcal{M}_{\geq 2}(\sigma^{(w;z)}) \setminus (N^+(w') \cup N^+(v)))}{|\mathcal{M}_{\geq 2}(\sigma^{(w;z)}) \setminus (N^+(w') \cup N^+(v))|}.$$

Law of (X_1, X_2). Given v, w, w', z and X_1, by definition $X_2 = (X_1)^{(u;w')}$

$$\mathbb{P}(X_1 = \sigma, X_2 = \sigma') = \sum_{v,w,w',z} \frac{\mu_{N,\alpha}^v(\sigma^{z,w'})\mathbf{1}(\sigma_{N^+(w')} = 0, \sigma_{N(w)} = 0, \sigma_x(1) = 1)}{|N(v)|(N-1)^2 n^2 \alpha^{\tau^\ell}} \times \frac{\mathbf{1}(z \in \mathcal{M}_{\geq 2}(\sigma^{(w;z)}) \setminus (N^+(w') \cup N^+(v)))}{|\mathcal{M}_{\geq 2}(\sigma^{(w;z)}) \setminus (N^+(w') \cup N^+(v))|} \mathbf{1}(\sigma' = \sigma^{(u;w')}).$$

Law of X_2.

$$\mathbb{P}(X_2=\sigma)=\sum_{v,w,w',z}\frac{\mu^v_{N,\alpha}(\sigma^{u,z})\mathbf{1}(\sigma_{\ell_{w'}}(j_{w'})=1,\sigma_{N(w)}=0)}{|N(v)|(N-1)^2n^2\alpha^{\tau^\ell}}\times \\ \frac{\mathbf{1}(z\in\mathcal{M}_{\geq 2}((\sigma^{(w;z)})^{(u;w')})\setminus(N^+(w')\cup N^+(v)))}{|\mathcal{M}_{\geq 2}((\sigma^{(w;z)})^{(u;w')})\setminus(N^+(w')\cup N^+(v))|}.$$

Law of (X_2,η_0). Given v,w,w',z and X_1, by definition $\eta_0=(X_2)^{(u;v)}$

$$\mathbb{P}(X_2=\sigma,\eta_0=\sigma')=\sum_{v,w,w',z}\frac{\mu^v_{N,\alpha}(\sigma^{u,z})\mathbf{1}(\sigma_{\ell_{w'}}(j_{w'})=1,\sigma_{N(w)}=0)}{|N(v)|(N-1)^2n^2\alpha^{\tau^\ell}}\times \\ \frac{\mathbf{1}(z\in\mathcal{M}_{\geq 2}((\sigma^{(w;z)})^{(u;w')})\setminus(N^+(w')\cup N^+(v)))}{|\mathcal{M}_{\geq 2}((\sigma^{(w;z)})^{(u;w')})\setminus(N^+(w')\cup N^+(v))|}\mathbf{1}(\sigma'=\sigma^{(z;u)}).$$

Finally, the law of (η_0,η_1') is Γ, since in the third case we have $\eta_1'=\eta_0'$.

Comparison between the law of (η_1,X_0) and Γ. Explicitly, the law of (η_1,X_0) is

$$\mathbb{P}(\eta_1=\sigma,X_0=\sigma')=\sum_{v,w,w',z}\frac{\mu^v_{N,\alpha}(\sigma^{(w;w')})\mathbf{1}(\sigma_{N^+(\ell_{w'})}=0,\sigma_{\ell_w}(j_w)=1)}{|N(v)|(N-1)^2n^2\alpha^{\tau^\ell}}\times \\ \frac{\mathbf{1}(z\in\mathcal{M}_{\geq 2}(\sigma)\setminus(N^+(w')\cup N^+(v)))}{|\mathcal{M}_{\geq 2}(\sigma)\setminus(N^+(w')\cup N^+(v))|}\mathbf{1}(\sigma'=\sigma^{(u;z)}),$$

where the quantity in the sum is non-zero only if $\sigma_{N^+(z)}=0$. Therefore, σ lies in the support of $\mu^z_{N,\alpha}$, and by homogeneity we can apply estimate (A.13). This gives

$$\mathbb{P}(\eta_1=\sigma,X_0=\sigma')(\sigma,\sigma')\leq\frac{12}{(N-1)n}\sum_z\mu^z_{N,\alpha}(\sigma)\mathbf{1}(\sigma'=\sigma^{(u;z)})\times \\ \sum_{v,w,w'}\frac{\mathbf{1}(w'\in N(v))}{|\mathcal{M}_{\geq 2}(\sigma)\setminus(N^+(w')\cup N^+(v))|}\frac{\mathbf{1}(\sigma_{N^+(\ell_{w'})}=0,\sigma_{\ell_w}(j_w)=1)}{|N(v)|(N-1)n\alpha^{\tau^\ell}}.$$

Proceeding similarly to (A.5) and using the hypothesis $\alpha\leq\frac{1}{C\Delta^2}$, which implies $\alpha^{\tau^\ell}\leq\frac{1}{(C-1)\Delta}$, we get for $C\geq 5$

$$|\mathcal{M}_{\geq 2}(\sigma)\setminus(N^+(w')\cup N^+(v))|\geq(N-1)n\big[1-(\Delta+1)(\alpha^{\tau^\ell}+2/(N-1)n)\big]\geq\frac{1}{2}(N-1)n.$$

Moreover,

$$\sum_{v,w}\mathbf{1}(w\in N(v),\sigma_{\ell_w}(j_w)=1)=\sum_{v,w}\mathbf{1}(v\in N(w),\sigma_{\ell_w}(j_w)=1)\leq\Delta N(n-1)\alpha^{\tau^\ell}.$$

Finally, we obtain

$$\begin{aligned}
&\sum_{v,w,w'} \frac{\mathbf{1}(w' \in N(v))}{|\mathcal{M}_{\geq 2}(\sigma) \setminus (N^+(w') \cup N^+(v))|} \frac{\mathbf{1}(\sigma_{N^+(\ell_{w'})} = 0, \sigma_{\ell_w}(j_w) = 1)}{|N(v)|(N-1)n\alpha^{\tau^\ell}} \\
&\leq \frac{2}{(N-1)^2 n^2 \alpha^{\tau^\ell}} \sum_{w'} \mathbf{1}(\sigma_{N^+(\ell_{w'})} = 0) \sum_{v,w} \frac{\mathbf{1}(w \in N(v), \sigma_{\ell_w}(j_w) = 1)}{|N(v)|} \\
&\leq \frac{2}{(N-1)^2 n^2 \alpha^{\tau^\ell}} (N-1)^2 n^2 \alpha^{\tau^\ell} \Delta = 2\Delta,
\end{aligned}$$

which leads to

$$\mathbb{P}(\eta_1 = \sigma, X_0 = \sigma') \leq 24\Delta\, \Gamma(\sigma, \sigma').$$

Comparison between the law of (X_1, X_0) and Γ. Note that $\sigma' = \sigma^{(w;u)}$ is in the support of $\mu^v_{N,\alpha}$ and $\mu^w_{N,\alpha}$, allowing a change of measure as in Proposition (A.2). Moreover, z must satisfy $\sigma'_{\ell_z}(j_z) = 1$. Using the previous observation and the lower bound $|\mathcal{M}_{\geq 2}(\sigma^{(u;z)}) \setminus (N^+(w') \cup N^+(v))| \geq \frac{1}{2}(N-1)n$, we obtain

$$\begin{aligned}
\mathbb{P}(X_1 = \sigma, X_0 = \sigma') &\leq \frac{24}{(N-1)^2 n^2} \sum_{w : j_w \geq 2} \mu^w_{N,\alpha}(\sigma') \mathbf{1}(\sigma = \sigma'^{(w;u)}) \times \\
&\quad \sum_{v,w',z} \frac{\mathbf{1}(\sigma_{N^+(w')} = 0, \sigma_{\ell_w}(j_w) = 1, v \in N(w), \sigma'_{N^+(z)} = 0)}{|N(v)|(N-1)n\alpha^{\tau^\ell}} \\
&\leq 24 \frac{1}{(N-1)n} \sum_{w : j_w \geq 2} \mu^w_{N,\alpha}(\sigma') \mathbf{1}(\sigma = \sigma'^{(w;u)}),
\end{aligned}$$

where the second sum is estimated by $\Delta(N-1)^2 n^2 \alpha^{\tau^\ell}$, giving

$$\mathbb{P}(X_1 = \sigma, X_0 = \sigma') \leq 24\Delta\, \Gamma(\sigma, \sigma').$$

Comparison between the law of (X_1, X_2) and Γ. The sum over z is non-zero only for z such that $\sigma_{\ell_z}(j_z) = 1$. Proceeding as before, we have

$$\begin{aligned}
\mathbb{P}(X_1 = \sigma, X_2 = \sigma') &\leq \frac{24}{(N-1)^2 n^2} \sum_{w' : j_{w'} \geq 2} \mu^{w'}_{N,\alpha}(\sigma) \mathbf{1}(\sigma' = \sigma^{(u;w)}) \times \\
&\quad \sum_{v,w,z} \frac{\mathbf{1}(\sigma_{\ell_z}(j_z) = 1, w \in N(v))}{(N-1)n\alpha^{\tau^\ell}} \\
&\leq \frac{24\Delta}{(N-1)n} \sum_{w' : j_{w'} \geq 2} \mu^{w'}_{N,\alpha}(\sigma) \mathbf{1}(\sigma' = \sigma^{(u;w)}),
\end{aligned}$$

leading to

$$\mathbb{P}(X_1 = \sigma, X_2 = \sigma') \leq 24\Delta\, \Gamma(\sigma, \sigma').$$

Comparison between the law of (η_0, X_2) and Γ. Since σ' is in the support of $\mu^z_{N,\alpha}$,

we can change the measure and reorganize the sum as before:

$$\mathbb{P}(\eta_0 = \sigma, X_2 = \sigma') \leq \frac{24}{(N-1)^2 n^2} \sum_{z : j_z \geq 2} \mu_{N,\alpha}^z(\sigma') \mathbf{1}(\sigma' = \sigma^{(u;z)}) \times \sum_{v,w,w'} \frac{\mathbf{1}(\sigma_{\ell_{w'}}(j_{w'}) = 1, \sigma_{N(w)} = 0)}{|N(v)|(N-1)n}.$$

The last sum is bounded by $\Delta(N-1)n$, giving

$$\mathbb{P}(\eta_0 = \sigma, X_2 = \sigma') \leq 24\Delta\, \Gamma(\sigma, \sigma').$$

This concludes the comparison that we need. □

We conclude this subsection with an estimate we used in the previous proof.

Proposition A.2. *Consider a graph G with n vertices and maximum degree $\Delta \geq 3$. Then, for all $\tau \in \Omega_G$ and for all $\ell \in [n]$ such that $\tau_k = 0$ for all $k \in N^+(\ell)$, we have*

$$\frac{1}{(1+\lambda)^\Delta + \lambda} \leq \mathbb{P}_{\sigma \sim \mu_{\lambda,G}}(\sigma_{N^+(\ell)} = 0 \mid \sigma_\Lambda = \tau_\Lambda) \leq \frac{1}{1+\lambda}, \tag{A.16}$$

where $\Lambda = (N^+(\ell))^c$ denotes the complement of $N^+(\ell)$.

Proof. Let σ be an independent set of the graph G sampled according to $\mu_{G,\lambda}$. Conditioned on the configuration outside $N^+(\ell)$ being equal to some τ_Λ, the restriction $\widetilde{\sigma} = \{\sigma_k : k \in N^+(\ell)\}$ is again an independent set, distributed according to a hardcore model with fugacity λ, with some vertices forced to be 0, on a graph with maximum degree Δ.

Let $\widetilde{Z}_\lambda$ denote the partition function of this conditioned hardcore model. Then,

$$(1+\lambda)^\Delta + \lambda \geq \widetilde{Z}_\lambda \geq 1 + \lambda,$$

which yields the estimate (A.16). Here, we used that $\widetilde{Z}_\lambda$ is maximized on the star graph with ℓ as the center. This concludes the proof of proposition. □

Appendix B

B.1 The marginal polytope

Let

$$\mathcal{P}_{\mathrm{IS}} := \mathrm{conv}\{\sigma : \sigma \in \Omega_G\} \subseteq \mathbb{R}^V$$

be the independent set polytope. Equivalently, $\mathcal{P}_{\mathrm{IS}}$ is the set of all feasible one vertex marginal vectors of probability measures supported on Ω_G. Also, we write $\mathcal{P}_{\mathrm{IS}}^{\circ}$ for its interior.

We first note that the marginal assumption used in **??** implies that the prescribed vector belongs to the interior of this polytope.

Lemma B.1. *Let $\boldsymbol{m} = (m_v)_{v \in V}$ satisfy*

$$0 \leq m_v \leq \frac{1}{\Delta + 1} \qquad \forall v \in V.$$

Then $\boldsymbol{m} \in \mathcal{P}_{\mathrm{IS}}$. Moreover, if for some $\gamma, \delta > 0$,

$$\gamma \leq m_v \leq \frac{1 - \delta}{\Delta + 1} \qquad \forall v \in V,$$

then $\boldsymbol{m} \in \mathcal{P}_{\mathrm{IS}}^{\circ}$.

Proof. The first statement follows from the initialization construction in Algorithm 3, which gives a probability distribution on Ω_G with one-vertex marginals $\boldsymbol{m}$ whenever $\max_v m_v \leq 1/(\Delta + 1)$.

For the second statement, choose

$$\eta := \frac{1}{2} \min\left\{\gamma, \frac{\delta}{\Delta + 1}\right\}.$$

If $\|\boldsymbol{m} - \boldsymbol{m}'\|_\infty \leq \eta$, then

$$0 \leq m'_v \leq \frac{1}{\Delta + 1} \qquad \forall v \in V.$$

Hence every such $\boldsymbol{m}'$ belongs to $\mathcal{P}_{\mathrm{IS}}$ by the first part. Thus an ℓ^∞-neighborhood of $\boldsymbol{m}$ is contained in $\mathcal{P}_{\mathrm{IS}}$, and so $\boldsymbol{m} \in \mathcal{P}_{\mathrm{IS}}^{\circ}$. □

By the standard canonical parameter correspondence for finite exponential families,

for $\boldsymbol{m} \in \mathcal{P}_{\mathrm{IS}}^{\circ}$, there is a unique fugacity vector $\boldsymbol{\lambda} \in (0,\infty)^V$ whose hard-core measure has marginal vector $\boldsymbol{m}$; see, e.g., [WJ08]. Thus, Theorem B.1 verifies that the marginal assumptions used in the main text indeed correspond to a well-defined hard-core Gibbs measure $\mu_G^{\boldsymbol{m}}$.

B.1.1 Gradient descent for prescribed marginals

In this section, we answer the same algorithmic question addressed in Section 3.2.3, this time using an optimization approach. The method we present is an adaptation of the approach of [Dia+21].

Note that the hard-core model can be viewed as a canonical exponential family. In general, an exponential family has the form

$$P_{\boldsymbol{\theta}}(x) = \exp(\langle \boldsymbol{\theta}, T(x)\rangle - A(\theta)),$$

where $T(x)$ is the sufficient statistic and $A(\theta)$ is the log-partition function. For the hard-core model, the sufficient statistic is simply the independent-set vector σ. In the notation of the preliminaries, for $\boldsymbol{\theta} \in \mathbb{R}^V$,

$$\mu_{G,\boldsymbol{\theta}}(\sigma) = \exp\left(\langle \boldsymbol{\theta}, \sigma\rangle - A_G(\boldsymbol{\theta})\right), \qquad \sigma \in \Omega_G.$$

Equivalently, $\theta_v = \log \lambda_v$, where λ_v is the fugacity of v.

Let $\mathcal{P}_{\mathrm{IS}}$ be the independent set polytope introduced in Appendix B.1. In Appendix B.1 we noted that, for every $\boldsymbol{m} \in \mathcal{P}_{\mathrm{IS}}^{\circ}$, there is a unique finite canonical parameter $\boldsymbol{\theta}_{\boldsymbol{m}} \in \mathbb{R}^V$ such that

$$\mathbb{E}_{\mu_{G,\boldsymbol{\theta}_{\boldsymbol{m}}}}[\sigma] = \boldsymbol{m}.$$

This is the usual backward map from mean parameters to canonical parameters; see, for example, [WJ08; BGS14].

The problem of finding the fugacities corresponding to a prescribed marginal vector $\boldsymbol{m}$ is therefore the problem of computing this backward map. The standard dual objective for this problem is

$$F_{\boldsymbol{m}}(\boldsymbol{\theta}) := A_G(\boldsymbol{\theta}) - \langle \boldsymbol{\theta}, \boldsymbol{m}\rangle.$$

Indeed, the Fenchel conjugate of A_G is

$$A_G^*(\boldsymbol{m}) := \sup_{\boldsymbol{\theta}\in\mathbb{R}^V} \left\{\langle \boldsymbol{m}, \boldsymbol{\theta}\rangle - A_G(\boldsymbol{\theta})\right\}.$$

For $\boldsymbol{m} \in \mathcal{P}_{\mathrm{IS}}^{\circ}$, the supremum is attained at the unique parameter $\boldsymbol{\theta}_{\boldsymbol{m}}$ satisfying

$$\nabla A_G(\boldsymbol{\theta}_{\boldsymbol{m}}) = \boldsymbol{m}.$$

Thus, for the target marginal vector $\boldsymbol{m}$, the desired log-fugacity vector is characterized by

$$\boldsymbol{\theta}_{\boldsymbol{m}} \in_{\boldsymbol{\theta}\in\mathbb{R}^V} \{\langle \boldsymbol{m}, \boldsymbol{\theta}\rangle - A_G(\boldsymbol{\theta})\}.$$

Equivalently,

$$\boldsymbol{\theta}_{\boldsymbol{m}} \in_{\boldsymbol{\theta}\in\mathbb{R}^V} \{A_G(\boldsymbol{\theta}) - \langle \boldsymbol{\theta}, \boldsymbol{m}\rangle\} =_{\boldsymbol{\theta}\in\mathbb{R}^V} F_{\boldsymbol{m}}(\boldsymbol{\theta}).$$

Thus minimizing $F_{\boldsymbol{m}}$ is exactly the convex-dual formulation of recovering the canonical parameter whose Gibbs measure has marginals $\boldsymbol{m}$. Projected gradient descent for this type of exponential-family objective, including with approximate gradients, has been used in related works; see, for example, [WJ08; BGS14; Dia+21].

We now compute the derivative of $F_{\boldsymbol{m}}$. For every $v \in V$,

$$\frac{\partial}{\partial \theta_v} A_G(\boldsymbol{\theta}) = \mathbb{E}_{\mu_{G,\boldsymbol{\theta}}}[\sigma_v].$$

Hence

$$\nabla F_{\boldsymbol{m}}(\boldsymbol{\theta}) = \mathbb{E}_{\mu_{G,\boldsymbol{\theta}}}[\sigma] - \boldsymbol{m}.$$

Therefore, the minimizer $\boldsymbol{\theta}_* = \boldsymbol{\theta}_{\boldsymbol{m}}$ of $F_{\boldsymbol{m}}$ satisfies

$$\mathbb{E}_{\mu_{G,\boldsymbol{\theta}_*}}[\sigma] = \boldsymbol{m}.$$

Thus gradient descent on $F_{\boldsymbol{m}}$ has a simple interpretation: if the current marginal of a vertex v is larger than the target value m_v, then the update decreases the corresponding log-fugacity θ_v; if it is smaller than m_v, then the update increases θ_v. In this way, the algorithm iteratively adjusts the fugacities to reduce the discrepancy between the current marginals and the target marginals.

Moreover,

$$\nabla^2 F_{\boldsymbol{m}}(\boldsymbol{\theta}) = \mathrm{Cov}_{\mu_{G,\boldsymbol{\theta}}}(\sigma).$$

Consequently, a uniform lower bound on the minimum eigenvalue of $\mathrm{Cov}_{\mu_{G,\boldsymbol{\theta}}}(\sigma)$ gives strong convexity, while a uniform upper bound on the maximum eigenvalue gives smoothness. Once these two bounds are available, we can apply projected gradient descent to $F_{\boldsymbol{m}}$. In our setting, the gradient need not be computed exactly; instead, we estimate $\mathbb{E}_{\mu_{G,\boldsymbol{\theta}}}[\sigma]$ by approximately sampling from $\mu_{G,\boldsymbol{\theta}}$.

Throughout this section, we assume that

$$\gamma \le m_v \le (1-\delta)\alpha_c(\Delta) \qquad \forall v \in V, \tag{B.1}$$

where $\gamma, \delta > 0$ are fixed constants. The lower bound in (B.1) will be used to obtain a uniform lower bound on the minimum eigenvalue of the covariance matrix, and hence strong convexity. Note that since $\alpha_c(\Delta) < 1/(\Delta+1)$, (B.1) implies that there exists $\delta' > 0$

such that

$$\max_{v \in V} m_v \le \frac{1-\delta'}{\Delta+1}.$$

Hence, by Theorem B.1, the assumption implies $\boldsymbol{m} \in \mathcal{P}_{\mathrm{IS}}^{\circ}$.

We now state the consequence of (B.1) for the fugacities. Let $\mu = \mu_{G,\boldsymbol{\lambda_m}}$. For every $v \in V$, we have

$$m_v = (\lambda_{\boldsymbol{m}})_v \, \mu(\sigma_u = 0 \text{ for all } u \in v)\,.$$

Therefore, since the probability above is at most one,

$$(\lambda_{\boldsymbol{m}})_v \ge m_v \ge \gamma.$$

On the other hand, by the union bound,

$$\mu(\sigma_u = 0 \text{ for all } u \in v) \ge 1 - \sum_{u \in v} m_u \ge 1 - (\Delta+1)\alpha_c(\Delta).$$

Hence

$$(\lambda_{\boldsymbol{m}})_v = \frac{m_v}{\mu(\sigma_u = 0 \text{ for all } u \in v)} \le \frac{(1-\delta)\alpha_c(\Delta)}{1-(\Delta+1)\alpha_c(\Delta)}.$$

Since

$$\lambda_c(\Delta) = \frac{\alpha_c(\Delta)}{1-(\Delta+1)\alpha_c(\Delta)},$$

we conclude that

$$(\lambda_{\boldsymbol{m}})_v \le (1-\delta)\lambda_c(\Delta) \qquad \text{for all } v \in V.$$

Thus the desired fugacity vector lies uniformly below the uniqueness threshold. In particular, if we define

$$\lambda_{\max} := (1-\delta)\lambda_c(\Delta),$$

then

$$\gamma \le (\lambda_{\boldsymbol{m}})_v \le \lambda_{\max} \qquad \text{for all } v \in V.$$

Equivalently, the desired log-fugacity vector $\boldsymbol{\theta}_* = \boldsymbol{\theta}_{\boldsymbol{m}}$ belongs to the box

$$\mathcal{D} := \left\{ \boldsymbol{\theta} \in \mathbb{R}^V : \log\gamma \le \theta_v \le \log\lambda_{\max} \text{ for all } v \in V \right\}.$$

We will run projected gradient descent on this set $\mathcal{D}$.

The next lemma gives the strong convexity of $F_{\boldsymbol{m}}$ on $\mathcal{D}$. The proof uses the same conditioning idea as in [DP23]: after choosing a large independent set, we condition on its complement, so that the remaining available vertices behave as independent Bernoulli random variables.

Lemma B.2. *For every $\boldsymbol{\theta} \in \mathcal{D}$, one has*

$$\mathrm{Cov}_{\mu_{G,\boldsymbol{\theta}}}(\sigma) \succeq \ell I,$$

where

$$\ell := \frac{\gamma}{(\Delta+1)(1+\lambda_{\max})^{\Delta+2}}.$$

Consequently, $F_{\boldsymbol{m}}$ is ℓ-strongly convex on $\mathcal{D}$.

Proof. Fix $\boldsymbol{\theta} \in \mathcal{D}$, and write $\lambda_v = e^{\theta_v}$. Then

$$\gamma \le \lambda_v \le \lambda_{\max} \qquad \forall v \in V.$$

It is enough to show that for every $a \in \mathbb{R}^V$,

$$\mathrm{Var}_{\mu_{G,\boldsymbol{\theta}}}\left(\sum_{v\in V} a_v \sigma_v\right) \ge \ell \|a\|_2^2.$$

Since G has maximum degree at most Δ, we can color it using $\Delta+1$ colors. Let $V = V_1 \cup \ldots \cup V_{\Delta+1}$ be the corresponding partition of V into independent sets. Choose $i \in [\Delta+1]$ such that

$$\sum_{v\in V_i} a_v^2 \ge \frac{1}{\Delta+1}\|a\|_2^2.$$

By the law of total variance,

$$\mathrm{Var}_{\mu_{G,\boldsymbol{\theta}}}\left(\sum_{v\in V} a_v\sigma_v\right) \ge \mathbb{E}_{\mu_{G,\boldsymbol{\theta}}}\left[\mathrm{Var}_{\mu_{G,\boldsymbol{\theta}}}\left(\sum_{v\in V_i} a_v\sigma_v \,\middle|\, \sigma_{V\setminus V_i}\right)\right].$$

Since V_i is an independent set, once $\sigma_{V\setminus V_i}$ is fixed, the variables $(\sigma_v)_{v\in V_i}$ are conditionally independent. Moreover, for each $v \in V_i$, if some neighbor of v is occupied, then $\sigma_v = 0$ deterministically; otherwise σ_v is Bernoulli with parameter $\lambda_v/(1+\lambda_v)$. Therefore,

$$\mathrm{Var}_{\mu_{G,\boldsymbol{\theta}}}\left(\sigma_v \,\middle|\, \sigma_{V\setminus V_i}\right) = \mathbf{1}\{\sigma_u = 0 \text{ for all } u \in v\}\frac{\lambda_v}{(1+\lambda_v)^2}.$$

Hence

$$\begin{aligned}&\mathbb{E}_{\mu_{G,\boldsymbol{\theta}}}\left[\mathrm{Var}_{\mu_{G,\boldsymbol{\theta}}}\left(\sum_{v\in V_i} a_v\sigma_v \,\middle|\, \sigma_{V\setminus V_i}\right)\right]\\ &\quad= \sum_{v\in V_i} a_v^2\, \mu_{G,\boldsymbol{\theta}}(\sigma_u = 0 \text{ for all } u \in v)\,\frac{\lambda_v}{(1+\lambda_v)^2}.\end{aligned}$$

We now lower bound the probability that all neighbors of v are unoccupied. Reveal the vertices of v one at a time. At each step, regardless of the previously revealed values, the conditional probability that the next vertex is unoccupied is at least $1/(1+\lambda_{\max})$. Indeed, if all its neighbors are unoccupied, its conditional probability of being occupied is at most $\lambda_{\max}/(1+\lambda_{\max})$, and otherwise it is forced to be unoccupied. Therefore,

$$\mu_{G,\boldsymbol{\theta}}(\sigma_u = 0 \text{ for all } u \in v) \ge \frac{1}{(1+\lambda_{\max})^{\Delta}}.$$

Also,

$$\frac{\lambda_v}{(1+\lambda_v)^2} \geq \frac{\gamma}{(1+\lambda_{\max})^2}.$$

Thus,

$$\mu_{G,\boldsymbol{\theta}}(\sigma_u = 0 \text{ for all } u \in v)\, \frac{\lambda_v}{(1+\lambda_v)^2} \geq \frac{\gamma}{(1+\lambda_{\max})^{\Delta+2}}.$$

Combining the previous estimates gives

$$\begin{aligned}\operatorname{Var}_{\mu_{G,\boldsymbol{\theta}}}\left(\sum_{v\in V} a_v\sigma_v\right) &\geq \frac{\gamma}{(1+\lambda_{\max})^{\Delta+2}} \sum_{v\in V_i} a_v^2 \\ &\geq \frac{\gamma}{(\Delta+1)(1+\lambda_{\max})^{\Delta+2}}\|a\|_2^2.\end{aligned}$$

Since this holds for every $a \in \mathbb{R}^V$, we conclude that

$$\operatorname{Cov}_{\mu_{G,\boldsymbol{\theta}}}(\sigma) \succeq \frac{\gamma}{(\Delta+1)(1+\lambda_{\max})^{\Delta+2}} I.$$

Finally, since

$$\nabla^2 F_{\boldsymbol{m}}(\boldsymbol{\theta}) = \operatorname{Cov}_{\mu_{G,\boldsymbol{\theta}}}(\sigma),$$

the same bound implies that $F_{\boldsymbol{m}}$ is ℓ-strongly convex on $\mathcal{D}$. □

We now prove the corresponding smoothness bound. We use the complete spectral independence for the hard-core model below the uniqueness threshold established in [Che+21]. Since every $\boldsymbol{\theta} \in \mathcal{D}$ satisfies

$$e^{\theta_v} \leq \lambda_{\max} = (1-\delta)\lambda_c(\Delta) \qquad \text{for all } v \in V,$$

the measure $\mu_{G,\boldsymbol{\theta}}$ can be viewed as the hard-core model with uniform fugacity $\lambda_{\max}$, followed by local fields

$$\phi_v := \frac{e^{\theta_v}}{\lambda_{\max}} \leq 1.$$

Thus complete spectral independence applies uniformly throughout $\mathcal{D}$.

Lemma B.3. *Assume $\Delta \geq 3$. Then there exists a constant $L = L(\Delta, \delta)$ such that for every $\boldsymbol{\theta} \in \mathcal{D}$, one has*

$$\operatorname{Cov}_{\mu_{G,\boldsymbol{\theta}}}(\sigma) \preceq LI.$$

Consequently, $F_{\boldsymbol{m}}$ is L-smooth on $\mathcal{D}$.

Proof. Fix $\boldsymbol{\theta} \in \mathcal{D}$, and for $u \neq v$, define the signed influence

$$J_{u,v} := \mu_{G,\boldsymbol{\theta}}(\sigma_v = 1 \mid \sigma_u = 1) - \mu_{G,\boldsymbol{\theta}}(\sigma_v = 1 \mid \sigma_u = 0),$$

and set $J_{v,v} = 0$. By complete spectral independence of the hard-core model below the uniqueness threshold [Che+21, Lemma 8.4], there exists $C_{\mathrm{SI}} = C_{\mathrm{SI}}(\Delta, \delta)$ such that

$\alpha_{\mathrm{sp}}(|J|) \le C_{\mathrm{SI}}$, where $|J|$ denotes the entrywise absolute value of J, and α_{sp} denotes spectral radius.

Let

$$\mathcal{V}_{\boldsymbol{\theta}} := \mathrm{diag}(\mathrm{Var}_{\mu_{G,\boldsymbol{\theta}}}(\sigma_v) : v \in V).$$

For $u \neq v$, we have

$$\mathrm{Cov}_{\mu_{G,\boldsymbol{\theta}}}(\sigma_u, \sigma_v) = \mathrm{Var}_{\mu_{G,\boldsymbol{\theta}}}(\sigma_u) \left[\mu_{G,\boldsymbol{\theta}}(\sigma_v = 1 \mid \sigma_u = 1) - \mu_{G,\boldsymbol{\theta}}(\sigma_v = 1 \mid \sigma_u = 0)\right].$$

This means that

$$\mathrm{Cov}_{\mu_{G,\boldsymbol{\theta}}}(\sigma) = \mathcal{V}_{\boldsymbol{\theta}}(I + J),$$

where I is the identity matrix. Equivalently,

$$\mathcal{V}_{\boldsymbol{\theta}}^{-1/2} \mathrm{Cov}_{\mu_{G,\boldsymbol{\theta}}}(\sigma) \mathcal{V}_{\boldsymbol{\theta}}^{-1/2} = \mathcal{V}_{\boldsymbol{\theta}}^{1/2}(I + J)\mathcal{V}_{\boldsymbol{\theta}}^{-1/2}.$$

The matrix on the left is symmetric, and it is similar to $I + J$. Hence its eigenvalues are the same as the eigenvalues of $I + J$. Also, note that $\alpha_{\mathrm{sp}}(J) \le \alpha_{\mathrm{sp}}(|J|) \le C_{\mathrm{SI}}$. Therefore, the max eigen value of $\mathcal{V}_{\boldsymbol{\theta}}^{-1/2} \mathrm{Cov}_{\mu_{G,\boldsymbol{\theta}}}(\sigma) \mathcal{V}_{\boldsymbol{\theta}}^{-1/2}$ is at most $1 + C_{\mathrm{SI}}$. Thus

$$\mathcal{V}_{\boldsymbol{\theta}}^{-1/2} \mathrm{Cov}_{\mu_{G,\boldsymbol{\theta}}}(\sigma) \mathcal{V}_{\boldsymbol{\theta}}^{-1/2} \preceq (1 + C_{\mathrm{SI}})I.$$

Multiplying on the left and right by $\mathcal{V}_{\boldsymbol{\theta}}^{1/2}$ gives

$$\mathrm{Cov}_{\mu_{G,\boldsymbol{\theta}}}(\sigma) \preceq (1 + C_{\mathrm{SI}})\mathcal{V}_{\boldsymbol{\theta}}.$$

Finally, since each σ_v is $\{0, 1\}$-valued, $\mathrm{Var}_{\mu_{G,\boldsymbol{\theta}}}(\sigma_v) \le \frac{1}{4}$, and hence

$$\mathrm{Cov}_{\mu_{G,\boldsymbol{\theta}}}(\sigma) \preceq \frac{1 + C_{\mathrm{SI}}}{4} I.$$

Taking $L := \frac{1+C_{\mathrm{SI}}}{4}$ gives $\mathrm{Cov}_{\mu_{G,\boldsymbol{\theta}}}(\sigma) \preceq LI$. Since $\nabla^2 F_{\boldsymbol{m}}(\boldsymbol{\theta}) = \mathrm{Cov}_{\mu_{G,\boldsymbol{\theta}}}(\sigma)$, we conclude that $F_{\boldsymbol{m}}$ is L-smooth on $\mathcal{D}$. □

We now apply projected gradient descent on $\mathcal{D}$. The projection onto $\mathcal{D}$ is explicit. For $z \in \mathbb{R}^V$,

$$\Pi_{\mathcal{D}}(z)_v = \min\left\{\log \lambda_{\max}, \max\{\log \gamma, z_v\}\right\}.$$

Thus there is no projection error in our application.

We consider the projected gradient descent iterates

$$\boldsymbol{\theta}^{(t+1)} = \Pi_{\mathcal{D}}\left(\boldsymbol{\theta}^{(t)} - \frac{1}{L}\widehat{g}^{(t)}\right),$$

where $\widehat{g}^{(t)}$ is an approximation to $\nabla F_{\boldsymbol{m}}(\boldsymbol{\theta}^{(t)})$. We use the following guarantee, which is the specialization of [Dia+21, Algorithm 3 and Claim D.2] to our setting.

Lemma B.4. *Let $\boldsymbol{\theta}_{\boldsymbol{m}}$ be the unique minimizer of $F_{\boldsymbol{m}}$ over $\mathcal{D}$. Suppose that*

$$\left\|\widehat{g}^{(t)} - \nabla F_{\boldsymbol{m}}(\boldsymbol{\theta}^{(t)})\right\|_2 \le \eta_1 \qquad \text{for all } t.$$

Then after $T = O\left(\frac{L}{\ell}\log\frac{\operatorname{diam}(\mathcal{D})}{\eta}\right)$ iterations, one has

$$\left\|\boldsymbol{\theta}^{(T)} - \boldsymbol{\theta}_{\boldsymbol{m}}\right\|_2 \le \eta + \frac{\eta_1/L}{1-\sqrt{1-\ell/L}}.$$

In particular, since $1-\sqrt{1-\ell/L} \ge \frac{\ell}{2L}$, we have

$$\left\|\boldsymbol{\theta}^{(T)} - \boldsymbol{\theta}_{\boldsymbol{m}}\right\|_2 \le \eta + \frac{2\eta_1}{\ell}.$$

We now explain how accurate the output parameter needs to be. For any $\widehat{\boldsymbol{\theta}} \in \mathcal{D}$, we have

$$D_{\mathrm{KL}}\left(\mu_{G,\boldsymbol{\theta}_{\boldsymbol{m}}} \,\middle\|\, \mu_{G,\widehat{\boldsymbol{\theta}}}\right) = A_G(\widehat{\boldsymbol{\theta}}) - A_G(\boldsymbol{\theta}_{\boldsymbol{m}}) - \left\langle \nabla A_G(\boldsymbol{\theta}_{\boldsymbol{m}}), \widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_{\boldsymbol{m}}\right\rangle.$$

Since A_G is L-smooth on $\mathcal{D}$, it follows that

$$D_{\mathrm{KL}}\left(\mu_{G,\boldsymbol{\theta}_{\boldsymbol{m}}} \,\middle\|\, \mu_{G,\widehat{\boldsymbol{\theta}}}\right) \le \frac{L}{2}\left\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_{\boldsymbol{m}}\right\|_2^2.$$

Therefore, by Pinsker's inequality,

$$\left\|\mu_{G,\widehat{\boldsymbol{\theta}}} - \mu_{G,\boldsymbol{\theta}_{\boldsymbol{m}}}\right\|_{\mathrm{TV}} \le \frac{\sqrt{L}}{2}\left\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_{\boldsymbol{m}}\right\|_2.$$

Note that our algorithm first runs projected gradient descent with approximate gradients to find an estimate $\widehat{\boldsymbol{\theta}}$. At iteration t, it estimates the gradient by sampling from $\mu_{G,\boldsymbol{\theta}^{(t)}}$. The algorithm then uses Glauber dynamics to generate an approximate sample from $\mu_{G,\widehat{\boldsymbol{\theta}}}$. Hence there are three sources of error in the output: the event that some gradient estimate is inaccurate, the parameter estimation error $\left\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_{\boldsymbol{m}}\right\|_2$, and the final sampling error from approximately sampling $\mu_{G,\widehat{\boldsymbol{\theta}}}$. The following lemma, assuming that all gradient estimates are accurate, chooses the gradient accuracy and number of iterations so that the combined contribution of the parameter error and the final sampling error is at most $\varepsilon/2$.

Lemma B.5. *Let $\varepsilon > 0$, and set*

$$\eta := \frac{\varepsilon}{4\sqrt{L}}, \qquad \eta_1 := \frac{\ell\varepsilon}{8\sqrt{L}}.$$

Suppose that for every $t = 0, \dots, T-1$, $\left\|\widehat{g}^{(t)} - \nabla F_{\boldsymbol{m}}(\boldsymbol{\theta}^{(t)})\right\|_2 \le \eta_1$, where $T = O\left(\frac{L}{\ell}\log\frac{\operatorname{diam}(\mathcal{D})}{\eta}\right)$. Then the output $\widehat{\boldsymbol{\theta}} := \boldsymbol{\theta}^{(T)}$ satisfies

$$\left\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_{\boldsymbol{m}}\right\|_2 \le \frac{\varepsilon}{2\sqrt{L}}.$$

Consequently,

$$\left\|\mu_{G,\widehat{\boldsymbol{\theta}}} - \mu_G^{\boldsymbol{m}}\right\|_{\mathrm{TV}} \le \frac{\varepsilon}{4}.$$

If, in addition, the final sample is generated from a distribution within total variation distance $\varepsilon/4$ from $\mu_{G,\widehat{\boldsymbol{\theta}}}$, then the final output distribution is within total variation distance $\varepsilon/2$ from $\mu_G^{\boldsymbol{m}}$.

Proof. By Theorem B.4, after T iterations we have

$$\left\|\boldsymbol{\theta}^{(T)} - \boldsymbol{\theta}_{\boldsymbol{m}}\right\|_2 \le \eta + \frac{2\eta_1}{\ell}.$$

By our choice of η and η_1,

$$\eta + \frac{2\eta_1}{\ell} = \frac{\varepsilon}{4\sqrt{L}} + \frac{\varepsilon}{4\sqrt{L}} = \frac{\varepsilon}{2\sqrt{L}}.$$

Thus

$$\left\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_{\boldsymbol{m}}\right\|_2 \le \frac{\varepsilon}{2\sqrt{L}}.$$

Using the total variation bound proved above, we get

$$\left\|\mu_{G,\widehat{\boldsymbol{\theta}}} - \mu_{G,\boldsymbol{\theta}_{\boldsymbol{m}}}\right\|_{\mathrm{TV}} \le \frac{\sqrt{L}}{2}\left\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_{\boldsymbol{m}}\right\|_2 \le \frac{\varepsilon}{4}.$$

Since $\mu_{G,\boldsymbol{\theta}_{\boldsymbol{m}}} = \mu_G^{\boldsymbol{m}}$, this gives

$$\left\|\mu_{G,\widehat{\boldsymbol{\theta}}} - \mu_G^{\boldsymbol{m}}\right\|_{\mathrm{TV}} \le \frac{\varepsilon}{4}.$$

Finally, if the last Glauber-dynamics sample is drawn from a distribution within total variation distance $\varepsilon/4$ from $\mu_{G,\widehat{\boldsymbol{\theta}}}$, then by the triangle inequality the final output distribution is within total variation distance $\varepsilon/2$ from $\mu_G^{\boldsymbol{m}}$. □

Throughout the gradient descent we need to estimate the gradient. Recall that

$$\nabla F_{\boldsymbol{m}}(\boldsymbol{\theta}) = \mathbb{E}_{\mu_{G,\boldsymbol{\theta}}}[\sigma] - \boldsymbol{m}.$$

The estimation of $\mathbb{E}_{\mu_{G,\boldsymbol{\theta}}}[\sigma]$ will be done by sampling from $\mu_{G,\boldsymbol{\theta}}$ multiple times and then taking the average. Therefore, we need to make sure that we can sample from $\mu_{G,\boldsymbol{\theta}}$ for every $\boldsymbol{\theta} \in \mathcal{D}$.

Since every $\boldsymbol{\theta} \in \mathcal{D}$ satisfies

$$e^{\theta_v} \le \lambda_{\max} = (1-\delta)\lambda_c(\Delta) \qquad \text{for all } v \in V,$$

the fugacity vector from which we need to sample is uniformly below the uniqueness threshold, with slack δ. Therefore, the Glauber dynamics for $\mu_{G,\boldsymbol{\theta}}$ mixes to total variation distance ζ in

$$O_{\Delta,\delta}\left(n \log \frac{n}{\zeta}\right)$$

steps [CLV21; Ana+22; Che+22; CE22]. We note that, although the results in these works are presented only for uniform fugacity below the critical threshold with slack δ, they can be easily extended to non-uniform fugacity vectors as well. This extension assumes that each entry is at most $(1-\delta)\lambda_c(\Delta)$, consistent with our case. Therefore, for every sampling accuracy $\zeta > 0$, we can sample from a distribution within total variation distance ζ from $\mu_{G,\boldsymbol{\theta}}$ in

$$O_{\Delta,\delta}\left(n \log \frac{n}{\zeta}\right)$$

steps of Glauber dynamics.

The following lemma shows that if the number of samples that are taken to estimate the gradient at each iteration is sufficiently large, then the event that all the gradient estimates are accurate occurs with probability at least $1 - \frac{\varepsilon}{2}$. Thus, together with Theorem B.5, it shows that our algorithm outputs an independent set whose law is within ε variation distance of the target measure.

Lemma B.6. *Let $\varepsilon \in (0,1)$ and $\eta_1 > 0$. At each step $t = 0, \dots, T-1$, suppose that we generate independent samples $\sigma^{(t,1)}, \dots, \sigma^{(t,M)}$ such that each sample has law within total variation distance $\frac{\eta_1}{2\sqrt{n}}$ from $\mu_{G,\boldsymbol{\theta}^{(t)}}$. Define*

$$\widehat{g}^{(t)} := \frac{1}{M}\sum_{i=1}^{M} \sigma^{(t,i)} - \boldsymbol{m}.$$

If $M \geq \frac{2n}{\eta_1^2} \log \frac{4nT}{\varepsilon}$, then

$$\Pr\left(\left\|\widehat{g}^{(t)} - \nabla F_{\boldsymbol{m}}(\boldsymbol{\theta}^{(t)})\right\|_2 \leq \eta_1 \text{ for all } t = 0, \dots, T-1\right) \geq 1 - \frac{\varepsilon}{2}.$$

Proof. For each t, let

$$\mathcal{E}_t := \left\{\left\|\widehat{g}^{(t)} - \nabla F_{\boldsymbol{m}}(\boldsymbol{\theta}^{(t)})\right\|_2 \leq \eta_1\right\}.$$

We show that $\Pr(\mathcal{E}_t^c) \leq \frac{\varepsilon}{2T}$ for every t, and then take a union bound.

Fix t, and condition on all randomness before step t. Then $\boldsymbol{\theta}^{(t)}$ is fixed. Write

$$q_v^{(t)} := \mathbb{E}_{\mu_{G,\boldsymbol{\theta}^{(t)}}}[\sigma_v], \qquad v \in V.$$

For each $v \in V$, set

$$\overline{q}_v^{(t)} := \frac{1}{M}\sum_{i=1}^{M} \mathbb{E}[\sigma_v^{(t,i)}],$$

where the expectation is with respect to the actual law of the samples used at step t. Since each sample is within total variation distance $\eta_1/(2\sqrt{n})$ from $\mu_{G,\boldsymbol{\theta}^{(t)}}$, we have

$$\left|\overline{q}_v^{(t)} - q_v^{(t)}\right| \leq \frac{\eta_1}{2\sqrt{n}}.$$

Moreover, the random variables $\sigma_v^{(t,1)}, \dots, \sigma_v^{(t,M)}$ are independent and take values in $\{0,1\}$.

Hence, by Hoeffding's inequality,

$$\Pr\left(\left|\frac{1}{M}\sum_{i=1}^{M}\sigma_v^{(t,i)} - \overline{q}_v^{(t)}\right| > \frac{\eta_1}{2\sqrt{n}}\right) \le 2\exp\left(-\frac{M\eta_1^2}{2n}\right).$$

By the union bound over $v \in V$, with probability at least

$$1 - 2n\exp\left(-\frac{M\eta_1^2}{2n}\right),$$

we have, for every $v \in V$,

$$\left|\frac{1}{M}\sum_{i=1}^{M}\sigma_v^{(t,i)} - q_v^{(t)}\right| \le \frac{\eta_1}{\sqrt{n}}.$$

On this event,

$$\left\|\frac{1}{M}\sum_{i=1}^{M}\sigma^{(t,i)} - \mathbb{E}_{\mu_{G,\boldsymbol{\theta}^{(t)}}}[\sigma]\right\|_2 \le \eta_1.$$

Since

$$\nabla F_{\boldsymbol{m}}(\boldsymbol{\theta}^{(t)}) = \mathbb{E}_{\mu_{G,\boldsymbol{\theta}^{(t)}}}[\sigma] - \boldsymbol{m},$$

this implies the event $\mathcal{E}_t$. Therefore,

$$\Pr(\mathcal{E}_t^c) \le 2n\exp\left(-\frac{M\eta_1^2}{2n}\right).$$

By the choice of M,

$$2n\exp\left(-\frac{M\eta_1^2}{2n}\right) \le \frac{\varepsilon}{2T}.$$

Finally, taking a union bound over $t = 0, \dots, T-1$ gives

$$\Pr\left(\bigcap_{t=0}^{T-1}\mathcal{E}_t\right) \ge 1 - \frac{\varepsilon}{2}.$$

This proves the lemma. □

We now state the resulting algorithm.

Algorithm 7 Sampling from the hard-core model with prescribed marginals

Require: Graph $G = (V, E)$ with maximum degree $\Delta \geq 3$, target marginals $\boldsymbol{m} = (m_v)_{v \in V}$ satisfying (B.1) with parameters $\gamma, \delta > 0$, and accuracy $\varepsilon \in (0, 1)$.

Ensure: A random independent set whose law is within total variation distance ε from $\mu_G^{\boldsymbol{m}}$.

Set

$$\lambda_{\max} := (1 - \delta)\lambda_c(\Delta).$$

Set

$$\mathcal{D} := \left\{ \boldsymbol{\theta} \in \mathbb{R}^V : \log \gamma \leq \theta_v \leq \log \lambda_{\max} \text{ for all } v \in V \right\}.$$

Set ℓ, L to be the constants introduced in Theorem B.2 and Theorem B.3, respectively.

Define

$$\eta := \frac{\varepsilon}{4\sqrt{L}}, \qquad \eta_1 := \frac{\ell \varepsilon}{8\sqrt{L}}.$$

Set

$$T = O\left(\frac{L}{\ell} \log \frac{\operatorname{diam}(\mathcal{D})}{\eta} \right).$$

Choose an integer M satisfying

$$M \geq \frac{2n}{\eta_1^2} \log \frac{4nT}{\varepsilon}.$$

Choose an arbitrary initial point $\boldsymbol{\theta}^{(0)} \in \mathcal{D}$.

for $t = 0, \ldots, T - 1$ **do**

  Independently generate samples

$$\sigma^{(t,1)}, \ldots, \sigma^{(t,M)}$$

  such that each sample has law within total variation distance at most $\eta_1/(2\sqrt{n})$ from $\mu_{G,\boldsymbol{\theta}^{(t)}}$.

  Define

$$\widehat{g}^{(t)} := \frac{1}{M} \sum_{i=1}^{M} \sigma^{(t,i)} - \boldsymbol{m}.$$

  Update

$$\boldsymbol{\theta}^{(t+1)} = \Pi_{\mathcal{D}} \left(\boldsymbol{\theta}^{(t)} - \frac{1}{L} \widehat{g}^{(t)} \right).$$

end for

Run Glauber dynamics for $\mu_{G,\boldsymbol{\theta}^{(T)}}$ to total variation distance $\varepsilon/4$.

return the resulting independent set.

Theorem B.7. *Let $\delta, \gamma \in (0, 1)$. Let $G = (V, E)$ be a graph on n vertices with maximum degree $\Delta \geq 3$, and let $\boldsymbol{m}$ be a marginal vector satisfying*

$$\gamma \leq m_v \leq (1 - \delta)\alpha_c(\Delta) \qquad \forall v \in V.$$

Then, for every $\varepsilon \in (0,1)$, Algorithm 7 outputs a random independent set whose law is within total variation distance ε from $\mu_G^{\boldsymbol{m}}$. Moreover, its running time is

$$\widetilde{O}_{\Delta,\delta,\gamma}\left(\frac{n^2}{\varepsilon^2}\right).$$

Proof. Let

$$\mathcal{E} := \left\{ \left\| \widehat{g}^{(t)} - \nabla F_{\boldsymbol{m}}(\boldsymbol{\theta}^{(t)}) \right\|_2 \le \eta_1 \text{ for all } t = 0, \dots, T-1 \right\}.$$

By Theorem B.6, the choice of M in Algorithm 7 implies that

$$\Pr(\mathcal{E}) \ge 1 - \frac{\varepsilon}{2}.$$

On the event $\mathcal{E}$, Theorem B.5 implies that the law of the final output is within total variation distance $\varepsilon/2$ from $\mu_G^{\boldsymbol{m}}$. Therefore, if ν denotes the law of the output, then

$$\|\nu - \mu_G^{\boldsymbol{m}}\|_{\mathrm{TV}} \le \Pr(\mathcal{E}^c) + \|\nu(\cdot \mid \mathcal{E}) - \mu_G^{\boldsymbol{m}}\|_{\mathrm{TV}} \le \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon.$$

The running time bound follows from the choices of T and M in Algorithm 7, together with the $O_{\Delta,\delta}(n \log(n/\zeta))$ mixing bound for Glauber dynamics below the uniqueness threshold to within ζ total variation distance. Indeed, T is logarithmic in n and $1/\varepsilon$ up to constants depending on Δ, δ, γ, and

$$M = \widetilde{O}_{\Delta,\delta,\gamma}\left(\frac{n}{\varepsilon^2}\right).$$

Each approximate sample costs $\widetilde{O}_{\Delta,\delta,\gamma}(n)$ Glauber updates. Hence the total running time is

$$\widetilde{O}_{\Delta,\delta,\gamma}\left(\frac{n^2}{\varepsilon^2}\right).$$

□

B.1.2 Bisection for prescribed density

In this section we use the bisection method to answer the same algorithmic question addressed in Section 3.2.1: given a target density α^*, approximately sample from the uniform fugacity hard-core Gibbs measure with density α^*. Our bisection method is an adaptation of the binary search approach used in [DP23] for sampling from the uniform distribution on fixed-size independent sets.

Following the canonical parametrization of exponential families, we use $\theta := \log \lambda$ and write

$$Z_G(\theta) := \sum_{\sigma \in \Omega_G} e^{\theta|\sigma|}, \qquad \mu_{G,\theta}(\sigma) = \frac{e^{\theta|\sigma|}}{Z_G(\theta)} \qquad \forall \sigma \in \Omega_G.$$

Note that

$$(\log Z_G(\theta))' = \mathbb{E}_{\sigma\sim\mu_{G,\theta}}[|\sigma|], \qquad (\log Z_G(\theta))'' = \mathrm{Var}_{\sigma\sim\mu_{G,\theta}}(|\sigma|).$$

Since $\mathrm{Var}_{\sigma\sim\mu_{G,\theta}}(|\sigma|) > 0$, the map $\theta \mapsto \alpha(\mu_{G,\theta})$ is strictly increasing. Moreover, it is a bijection from $(-\infty, \infty)$ to $(0, i_{\max}/n)$, where $i_{\max}$ is the size of a maximum independent set in G. Therefore, for every $\alpha \in (0, i_{\max}/n)$, there is a unique $\theta(\alpha)$ such that the density of $\mu_{G,\theta(\alpha)}$ is exactly α.

Consider the function

$$q(\theta) := \mathbb{E}_{\sigma\sim\mu_{G,\theta}}[|\sigma|].$$

Since

$$q'(\theta) = \mathrm{Var}_{\sigma\sim\mu_{G,\theta}}(|\sigma|) > 0,$$

the map $\theta \mapsto q(\theta)$ is strictly increasing. Hence the unique parameter θ_* satisfying

$$q(\theta_*) = n\alpha^*$$

can be approximated by the bisection method. Thus, we first approximate θ_*, and then approximately sample from the Gibbs measure $\mu_{G,\widehat{\theta}}$, where $\widehat{\theta}$ is the output of the bisection procedure.

Throughout this section, we assume that

$$\gamma \le \alpha^* \le (1-\delta)\alpha_c(\Delta), \tag{B.2}$$

where $\gamma, \delta > 0$ are fixed constants. We need to find the range of fugacities corresponding to densities satisfying (B.2). For any such density α, let $\lambda(\alpha)$ be the corresponding fugacity. Then

$$n\alpha = \sum_{v\in V} \mu_{G,\lambda(\alpha)}(\sigma_v = 1) = \sum_{v\in V} \lambda(\alpha)\, \mu_{G,\lambda(\alpha)}\left(\sigma_u = 0 \text{ for all } u \in v\right) \le n\lambda(\alpha).$$

Therefore,

$$\lambda(\alpha) \ge \alpha \ge \gamma.$$

On the other hand, by the union bound,

$$\begin{aligned}
n\alpha &= \lambda(\alpha) \sum_{v\in V} \mu_{G,\lambda(\alpha)}\left(\sigma_u = 0 \text{ for all } u \in v\right) \\
&\ge \lambda(\alpha) \sum_{v\in V} \left(1 - \sum_{u\in v} \mu_{G,\lambda(\alpha)}(\sigma_u = 1)\right) \\
&\ge \lambda(\alpha) \left(n - (\Delta+1) \sum_{u\in V} \mu_{G,\lambda(\alpha)}(\sigma_u = 1)\right) \\
&\ge \lambda(\alpha) n (1 - (\Delta+1)\alpha_c(\Delta)).
\end{aligned}$$

Hence, using $\alpha \le (1-\delta)\alpha_c(\Delta)$, we get

$$\lambda(\alpha) \le \frac{(1-\delta)\alpha_c(\Delta)}{1-(\Delta+1)\alpha_c(\Delta)} = (1-\delta)\lambda_c(\Delta).$$

Therefore, defining

$$\lambda_{\max} := (1-\delta)\lambda_c(\Delta),$$

the corresponding range for the parameter θ is the interval

$$\mathcal{I} := [\log \gamma, \log \lambda_{\max}].$$

We will apply the bisection method on $\mathcal{I}$. In Theorems B.8 and B.9, we show lower and upper bounds on $q'(\theta)$. The lower bound ensures that approximate signs of $q(\theta) - n\alpha^*$ are reliable away from θ_*, while the upper bound is used to convert parameter error into total variation error.

Lemma B.8. *For every $\theta \in \mathcal{I}$, one has*

$$\mathrm{Var}_{\mu_{G,\theta}}(|\sigma|) \ge \ell n,$$

where

$$\ell := \frac{\gamma}{(\Delta+1)(1+\lambda_{\max})^{\Delta+2}}.$$

Consequently, $q'(\theta)$ is lower bounded by ℓn on $\mathcal{I}$.

Proof. Let $\boldsymbol{\theta} := (\theta)_{v\in V}$. By the definition of $\mathcal{I}$, we have $\boldsymbol{\theta} \in \mathcal{D}$, where $\mathcal{D}$ is the box used in Appendix B.1.1. Also, $\mu_{G,\theta} = \mu_{G,\boldsymbol{\theta}}$. By Theorem B.2, we have

$$\mathrm{Cov}_{\mu_{G,\boldsymbol{\theta}}}(\sigma) \succeq \ell I.$$

Taking the all-ones vector $\mathbf{1} \in \mathbb{R}^V$, this gives

$$\mathrm{Var}_{\mu_{G,\theta}}(|\sigma|) = \mathbf{1}\,\mathrm{Cov}_{\mu_{G,\boldsymbol{\theta}}}(\sigma)\mathbf{1} \ge \ell\|\mathbf{1}\|_2^2 = \ell n.$$

Since $q'(\theta) = \mathrm{Var}_{\mu_{G,\theta}}(|\sigma|)$, this proves the claim. □

Lemma B.9. *There exists a constant*

$$L = L(\Delta, \delta)$$

such that for every $\theta \in \mathcal{I}$, one has

$$\mathrm{Var}_{\mu_{G,\theta}}(|\sigma|) \le Ln.$$

Consequently, $q'(\theta)$ is upper bounded by Ln on $\mathcal{I}$.

Proof. Let $\boldsymbol{\theta} := (\theta)_{v \in V}$. By the definition of $\mathcal{I}$, we have $\boldsymbol{\theta} \in \mathcal{D}$, where $\mathcal{D}$ is the box used in Appendix B.1.1. Also, $\mu_{G,\theta} = \mu_{G,\boldsymbol{\theta}}$. By Theorem B.3, there exists a constant $L = L(\Delta, \delta)$ such that

$$\mathrm{Cov}_{\mu_{G,\boldsymbol{\theta}}}(\sigma) \preceq LI.$$

Taking the all-ones vector $\mathbf{1} \in \mathbb{R}^V$, this gives

$$\mathrm{Var}_{\mu_{G,\theta}}(|\sigma|) = \mathbf{1}\,\mathrm{Cov}_{\mu_{G,\boldsymbol{\theta}}}(\sigma)\mathbf{1} \leq L\|\mathbf{1}\|_2^2 = Ln.$$

Since $q'(\theta) = \mathrm{Var}_{\mu_{G,\theta}}(|\sigma|)$, this proves the claim. □

The following lemma gives the number of bisection steps needed to approximate θ_*, assuming access to approximate evaluations of q on $\mathcal{I}$. Recall that the bisection method proceeds as follows. Start with the interval $\mathcal{I}$. At each step, we query the midpoint θ_t of the current interval and estimate $q(\theta_t)$ by $\widehat{q}_t$. If

$$|\widehat{q}_t - n\alpha^*|$$

is small, then we stop and output θ_t. Otherwise, since q is increasing, we keep the left half if $\widehat{q}_t > n\alpha^*$, and keep the right half if $\widehat{q}_t < n\alpha^*$. The lower bound on q' ensures that stopping in the first case already gives a good approximation to θ_*, while in the second case the sign is reliable.

Lemma B.10. *Assume that at each bisection step t, with midpoint θ_t, we have an estimate $\widehat{q}_t$ satisfying*

$$|\widehat{q}_t - q(\theta_t)| \leq \eta_1.$$

Consider the following bisection procedure. If

$$|\widehat{q}_t - n\alpha^*| \leq \eta_1,$$

then stop and output $\widehat{\theta} = \theta_t$. Otherwise, keep the left or right half according to the sign of $\widehat{q}_t - n\alpha^*$. If the procedure does not stop early, then after

$$T = \left\lceil \log_2 \frac{\mathrm{diam}(\mathcal{I})}{\eta} \right\rceil$$

bisection steps, output any point $\widehat{\theta}$ from the final interval. Then the output satisfies

$$|\widehat{\theta} - \theta_*| \leq \max\left\{\eta, \frac{2\eta_1}{\ell n}\right\}.$$

Proof. Suppose first that the algorithm stops at some step t. Then

$$|\widehat{q}_t - n\alpha^*| \leq \eta_1.$$

Since $|\widehat{q}_t - q(\theta_t)| \le \eta_1$, we get

$$|q(\theta_t) - n\alpha^*| \le |q(\theta_t) - \widehat{q}_t| + |\widehat{q}_t - n\alpha^*| \le 2\eta_1.$$

By Theorem B.8, we have $q'(\theta) \ge \ell n$ on $\mathcal{I}$. Since $q(\theta_*) = n\alpha^*$, this implies

$$\ell n|\theta_t - \theta_*| \le |q(\theta_t) - q(\theta_*)| = |q(\theta_t) - n\alpha^*| \le 2\eta_1.$$

Therefore,

$$|\theta_t - \theta_*| \le \frac{2\eta_1}{\ell n}.$$

Since $\widehat{\theta} = \theta_t$, this gives the desired bound in the case where the algorithm stops early.

It remains to consider the case where the algorithm does not stop during the first T steps. Then, at every step $t < T$,

$$|\widehat{q}_t - n\alpha^*| > \eta_1.$$

Since $|\widehat{q}_t - q(\theta_t)| \le \eta_1$, the quantities $\widehat{q}_t - n\alpha^*$ and $q(\theta_t) - n\alpha^*$ have the same sign. Hence every bisection decision is correct, and θ_* remains inside the maintained interval throughout the procedure.

After

$$T = \left\lceil \log_2 \frac{\operatorname{diam}(\mathcal{I})}{\eta} \right\rceil$$

steps, the final interval has length at most η. Since θ_* belongs to this interval, any output $\widehat{\theta}$ from the final interval satisfies

$$|\widehat{\theta} - \theta_*| \le \eta.$$

Combining the two cases, we obtain

$$|\widehat{\theta} - \theta_*| \le \max\left\{\eta, \frac{2\eta_1}{\ell n}\right\}.$$

This completes the proof. □

For $\theta \in \mathcal{I}$, let

$$A(\theta) := \log Z_G(\theta).$$

Then

$$D_{\mathrm{KL}}\left(\mu_{G,\theta_*} \,\|\, \mu_{G,\theta}\right) = A(\theta) - A(\theta_*) - A'(\theta_*)(\theta - \theta_*).$$

By Theorem B.9, the function A is Ln-smooth on $\mathcal{I}$, and thus

$$D_{\mathrm{KL}}\left(\mu_{G,\theta_*} \,\|\, \mu_{G,\theta}\right) \le \frac{Ln}{2}(\theta - \theta_*)^2.$$

Therefore, by Pinsker's inequality,

$$\|\mu_{G,\theta_*} - \mu_{G,\theta}\|_{\mathrm{TV}} \le \frac{\sqrt{Ln}}{2}|\theta - \theta_*|.$$

Consequently, to make the error coming from parameter estimation at most $\varepsilon/4$ in total variation, it is enough to ensure that

$$|\theta - \theta_*| \le \frac{\varepsilon}{2\sqrt{Ln}}.$$

If we then sample from a distribution within total variation distance $\varepsilon/4$ from $\mu_{G,\theta}$, the final output distribution is within $\varepsilon/2$ of μ_{G,θ_*}.

The following, assuming that the estimates of the expected size $q(\theta_t)$ used in the bisection procedure are accurate, chooses the estimation accuracy and number of bisection steps so that the combined contribution of the parameter error and the final sampling error is at most $\varepsilon/2$. The remaining $\varepsilon/2$ will be used to account for the probability that one of the expected-size estimates fails.

Lemma B.11. *Let $\varepsilon > 0$, and set*

$$\eta := \frac{\varepsilon}{2\sqrt{Ln}}, \qquad \eta_1 := \frac{\ell n \varepsilon}{4\sqrt{Ln}}.$$

Suppose that at every bisection step t, with midpoint θ_t, we have an estimate $\widehat{q}_t$ satisfying

$$|\widehat{q}_t - q(\theta_t)| \le \eta_1.$$

Run the bisection procedure of Theorem B.10 for at most

$$T = \left\lceil \log_2 \frac{\mathrm{diam}(\mathcal{I})}{\eta} \right\rceil$$

steps, and let $\widehat{\theta}$ be its output. Then

$$\left|\widehat{\theta} - \theta_*\right| \le \frac{\varepsilon}{2\sqrt{Ln}}.$$

Consequently,

$$\left\|\mu_{G,\widehat{\theta}} - \mu_{G,\theta_*}\right\|_{\mathrm{TV}} \le \frac{\varepsilon}{4}.$$

If, in addition, the final sample is generated from a distribution within total variation distance $\varepsilon/4$ from $\mu_{G,\widehat{\theta}}$, then the final output distribution is within total variation distance $\varepsilon/2$ from μ_{G,θ_*}.

Proof. By Theorem B.10, the output $\widehat{\theta}$ satisfies

$$|\widehat{\theta} - \theta_*| \le \max\left\{\eta, \frac{2\eta_1}{\ell n}\right\}.$$

With our choices of η and η_1, we have

$$\eta = \frac{\varepsilon}{2\sqrt{Ln}}, \qquad \frac{2\eta_1}{\ell n} = \frac{\varepsilon}{2\sqrt{Ln}}.$$

Therefore,

$$|\widehat{\theta} - \theta_*| \le \frac{\varepsilon}{2\sqrt{Ln}}.$$

By the smoothness discussion above,

$$\left\|\mu_{G,\widehat{\theta}} - \mu_{G,\theta_*}\right\|_{\mathrm{TV}} \le \frac{\sqrt{Ln}}{2} \cdot \frac{\varepsilon}{2\sqrt{Ln}} = \frac{\varepsilon}{4}.$$

Finally, if the distribution used to generate the final sample is within total variation distance $\varepsilon/4$ from $\mu_{G,\widehat{\theta}}$, then by the triangle inequality its distance from μ_{G,θ_*} is at most $\frac{\varepsilon}{4} + \frac{\varepsilon}{4} = \frac{\varepsilon}{2}$. This completes the proof. □

We first compute an approximation $\widehat{\theta} \in \mathcal{I}$ to θ_* using the bisection method, and then sample from the hard-core model with fugacity $e^{\widehat{\theta}}$. Note that, since

$$q(\theta) = \mathbb{E}_{\mu_{G,\theta}}[|\sigma|],$$

throughout the bisection process we also need to sample from $\mu_{G,\theta}$ in order to estimate $\mathbb{E}_{\mu_{G,\theta}}[|\sigma|]$.

Since $\theta \in \mathcal{I}$, we have

$$e^{\theta} \le \lambda_{\max} = (1-\delta)\lambda_c(\Delta).$$

Thus the fugacity is uniformly below the uniqueness threshold, with slack δ. Hence, by the rapid-mixing theorem for Glauber dynamics for the hard-core model below the uniqueness threshold [CLV21; Ana+22; Che+22; CE22], for every $\zeta > 0$ we can sample from a distribution within total variation distance ζ from $\mu_{G,\theta}$ in

$$O_{\Delta,\delta}\left(n \log \frac{n}{\zeta}\right)$$

steps of Glauber dynamics.

The following lemma shows that by taking sufficiently many approximate samples from μ_{G,θ_t} at each bisection step, the estimates of the expected size $q(\theta_t)$ are accurate throughout the bisection procedure with probability at least $1 - \frac{\varepsilon}{2}$. Together with Theorem B.11, this proves that the algorithm outputs an independent set whose law is within total variation distance ε of the target measure.

Lemma B.12. *Let $\varepsilon \in (0,1)$ and $\eta_1 > 0$. Set*

$$\zeta := \min\left\{\frac{\eta_1}{2n}, \frac{L}{n}\right\}.$$

At each bisection step $t = 0, \dots, T-1$, suppose that we generate independent samples

$$\sigma^{(t,1)}, \dots, \sigma^{(t,M)}$$

such that each sample has law within total variation distance ζ from μ_{G,θ_t}. Write $M = BK$, and partition the samples into K blocks of size B. For each $j = 1, \dots, K$, define

$$Y_j^{(t)} := \frac{1}{B} \sum_{i=(j-1)B+1}^{jB} |\sigma^{(t,i)}|.$$

Let

$$\widehat{q}_t := \operatorname{median}\{Y_1^{(t)}, \dots, Y_K^{(t)}\}.$$

If

$$B \geq \frac{32Ln}{\eta_1^2}, \qquad K \geq 8 \log \frac{2T}{\varepsilon},$$

then

$$\Pr\left(|\widehat{q}_t - q(\theta_t)| \leq \eta_1 \text{ for all } t = 0, \dots, T-1\right) \geq 1 - \frac{\varepsilon}{2}.$$

Proof. For each t, let

$$\mathcal{E}_t := \left\{|\widehat{q}_t - q(\theta_t)| \leq \eta_1\right\}.$$

We show that

$$\Pr(\mathcal{E}_t^c) \leq \frac{\varepsilon}{2T}$$

for every t, and then take a union bound.

Fix t, and condition on all randomness before step t. Then θ_t is fixed. Write

$$q_t := q(\theta_t) = \mathbb{E}_{\mu_{G,\theta_t}}[|\sigma|].$$

By Theorem B.9,

$$\operatorname{Var}_{\mu_{G,\theta_t}}(|\sigma|) \leq Ln.$$

Fix a block j, and set

$$\overline{q}_j^{(t)} := \mathbb{E}[Y_j^{(t)}],$$

where the expectation is with respect to the actual laws of the samples used at step t. Since each sample is within total variation distance ζ from μ_{G,θ_t}, and since $0 \leq |\sigma| \leq n$, we have

$$\left|\overline{q}_j^{(t)} - q_t\right| \leq n\zeta \leq \frac{\eta_1}{2}.$$

Moreover, if X denotes one of the random variables $|\sigma^{(t,i)}|$ in the block, then

$$\operatorname{Var}(X) \leq \mathbb{E}\left[(X - q_t)^2\right] \leq \operatorname{Var}_{\mu_{G,\theta_t}}(|\sigma|) + n^2\zeta \leq 2Ln,$$

where we used $\zeta \le L/n$. Since the samples in the block are independent,

$$\mathrm{Var}(Y_j^{(t)}) \le \frac{2Ln}{B}.$$

Therefore, by Chebyshev's inequality,

$$\Pr\left(\left|Y_j^{(t)} - \overline{q}_j^{(t)}\right| > \frac{\eta_1}{2}\right) \le \frac{8Ln}{B\eta_1^2} \le \frac{1}{4}.$$

It follows that

$$\Pr\left(\left|Y_j^{(t)} - q_t\right| > \eta_1\right) \le \frac{1}{4}.$$

Thus each block is good, meaning

$$\left|Y_j^{(t)} - q_t\right| \le \eta_1,$$

with probability at least $3/4$. Since the blocks are independent, Hoeffding's inequality gives

$$\Pr\left(|\widehat{q}_t - q_t| > \eta_1\right) \le \exp\left(-\frac{K}{8}\right).$$

By the choice of K,

$$\exp\left(-\frac{K}{8}\right) \le \frac{\varepsilon}{2T}.$$

Therefore,

$$\Pr(\mathcal{E}_t^c) \le \frac{\varepsilon}{2T}.$$

Finally, taking a union bound over $t = 0, \ldots, T-1$ gives

$$\Pr\left(\bigcap_{t=0}^{T-1} \mathcal{E}_t\right) \ge 1 - \frac{\varepsilon}{2}.$$

This proves the lemma. □

We now state the resulting algorithm.

Algorithm 8 Bisection sampler for the hard-core model with prescribed density

Require: Graph $G = (V, E)$ with maximum degree $\Delta \geq 3$, target density α^* satisfying (B.2) with parameters $\gamma, \delta > 0$, and accuracy $\varepsilon \in (0, 1)$.

Ensure: A random independent set whose law is within total variation distance ε from μ_{G,θ_*}, where $q(\theta_*) = n\alpha^*$.

Set $\lambda_{\max} := (1-\delta)\lambda_c(\Delta)$ and $\mathcal{I} := [\log\gamma, \log\lambda_{\max}]$.

Let ℓ, L be the constants from Theorems B.8 and B.9.

Set

$$\eta := \frac{\varepsilon}{2\sqrt{Ln}}, \qquad \eta_1 := \frac{\ell n \varepsilon}{4\sqrt{Ln}}, \qquad \zeta := \min\left\{\frac{\eta_1}{2n}, \frac{L}{n}\right\}.$$

Set

$$T := \left\lceil \log_2 \frac{\operatorname{diam}(\mathcal{I})}{\eta} \right\rceil.$$

Choose integers B, K with

$$B \geq \frac{32Ln}{\eta_1^2}, \qquad K \geq 8 \log \frac{2T}{\varepsilon},$$

and set $M := BK$.

Initialize $a_0 := \log\gamma$ and $b_0 := \log\lambda_{\max}$.

for $t = 0, \ldots, T-1$ **do**

  Set $\theta_t := (a_t + b_t)/2$.

  Generate independent samples $\sigma^{(t,1)}, \ldots, \sigma^{(t,M)}$, each within total variation distance ζ from μ_{G,θ_t}.

  Partition the samples into K blocks of size B, and define

$$Y_j^{(t)} := \frac{1}{B} \sum_{i=(j-1)B+1}^{jB} |\sigma^{(t,i)}|, \qquad j = 1, \ldots, K.$$

  Set

$$\widehat{q}_t := \operatorname{median}\{Y_1^{(t)}, \ldots, Y_K^{(t)}\}.$$

  if $|\widehat{q}_t - n\alpha^*| \leq \eta_1$ **then**

    Set $\widehat{\theta} := \theta_t$ and terminate the loop.

  else if $\widehat{q}_t > n\alpha^*$ **then**

    Set $a_{t+1} := a_t$ and $b_{t+1} := \theta_t$.

  else

    Set $a_{t+1} := \theta_t$ and $b_{t+1} := b_t$.

  end if

end for

if the loop was not terminated early **then**

  Set $\widehat{\theta} := (a_T + b_T)/2$.

end if

Run Glauber dynamics for $\mu_{G,\widehat{\theta}}$ to total variation distance $\varepsilon/4$.

return the resulting independent set.

Theorem B.13. *Let $\gamma, \delta \in (0, 1)$. Let $G = (V, E)$ be a graph on n vertices with maximum*

degree $\Delta \geq 3$, and let α^* be a density satisfying

$$\gamma \leq \alpha^* \leq (1-\delta)\alpha_c(\Delta).$$

Then, for every $\varepsilon \in (0,1)$, Algorithm 8 outputs a random independent set whose law is within total variation distance ε from μ_{G,θ_*}, where $q(\theta_*) = n\alpha^*$. Moreover, its running time is

$$\widetilde{O}_{\Delta,\delta,\gamma}\left(\frac{n}{\varepsilon^2}\right).$$

Proof. Let $\mathcal{E}$ be the event that every estimate actually used by the bisection procedure satisfies

$$|\widehat{q}_t - q(\theta_t)| \leq \eta_1.$$

By Theorem B.12, the choice of M in Algorithm 8 implies that

$$\Pr(\mathcal{E}) \geq 1 - \frac{\varepsilon}{2}.$$

On the event $\mathcal{E}$, the deterministic bisection guarantee in Theorem B.11 implies that the law of the final output is within total variation distance $\varepsilon/2$ from μ_{G,θ_*}. Therefore, if ν denotes the law of the output, then

$$\|\nu - \mu_{G,\theta_*}\|_{\mathrm{TV}} \leq \Pr(\mathcal{E}^c) + \|\nu(\cdot \mid \mathcal{E}) - \mu_{G,\theta_*}\|_{\mathrm{TV}} \leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon.$$

The running time bound follows from the choices of T, M, B, K in Algorithm 8, together with the $O_{\Delta,\delta}(n\log(n/\zeta))$ mixing bound for Glauber dynamics below the uniqueness threshold to within ζ total variation distance. Indeed, T and K are logarithmic in n and $1/\varepsilon$, up to constants depending on Δ, δ, γ, and

$$B = O_{\Delta,\delta,\gamma}\left(\frac{1}{\varepsilon^2}\right).$$

Thus

$$M = BK = \widetilde{O}_{\Delta,\delta,\gamma}\left(\frac{1}{\varepsilon^2}\right).$$

Each approximate sample costs $\widetilde{O}_{\Delta,\delta,\gamma}(n)$ Glauber updates. Hence the total running time is

$$\widetilde{O}_{\Delta,\delta,\gamma}\left(\frac{n}{\varepsilon^2}\right).$$

□

B.1.3 Missing proofs from Section 3.2.2

Proof of Theorem 3.12. Similar to Appendix B.1.2, we use the log-fugacity coordinate $\theta = \log\lambda$. We write θ_α for the log-fugacity parameter corresponding to α. Recall from

Appendix B.1.2 that

$$A(\theta_\alpha) := \log Z_G(\theta_\alpha), \qquad A'(\theta_\alpha) = n\alpha, \qquad A''(\theta_\alpha) = \mathrm{Var}_{\sigma\sim\mu_G^\alpha}(|\sigma|).$$

Also, the interval $[\alpha, \hat{\alpha}]$ corresponds to an interval $[\theta_\alpha, \theta_{\hat{\alpha}}]$ for the log-fugacity parameters. Note that the upper bound assumption on the variance for densities in $[\alpha, \hat{\alpha}]$ shows that A is Cn-smooth on $[\theta_\alpha, \hat{\theta}_\alpha]$. Thus, as we proved in Appendix B.1.2, we have

$$\|\mu_G^\alpha - \mu_G^{\hat{\alpha}}\|_{\mathrm{TV}} \le \frac{\sqrt{Cn}}{2}|\theta_\alpha - \theta_{\hat{\alpha}}|.$$

Moreover, note that by the lower bound assumption on the variance and the mean value theorem, there exists a density $\beta \in [\alpha, \hat{\alpha}]$ such that

$$|n\hat{\alpha} - n\alpha| = \mathrm{Var}_{\sigma\sim\mu_G^\beta}(|\sigma|)|\theta_\alpha - \theta_{\hat{\alpha}}| \ge cn|\theta_\alpha - \theta_{\hat{\alpha}}|.$$

Thus, we conclude that

$$\|\mu_G^\alpha - \mu_G^{\hat{\alpha}}\|_{\mathrm{TV}} \le \frac{\sqrt{Cn}}{2c}|\alpha - \hat{\alpha}|.$$

□