Electronic Journal of Probability

Electron. J. Probab. **21** (2016), no. 29, 1–22.
ISSN: 1083-6489 DOI: 10.1214/16-EJP4321

Dynamics of lattice triangulations on thin rectangles

Pietro Caputo* Fabio Martinelli† Alistair Sinclair‡
Alexandre Stauffer§

Abstract

We consider random lattice triangulations of $n \times k$ rectangular regions with weight $\lambda^{|\sigma|}$ where $\lambda > 0$ is a parameter and $|\sigma|$ denotes the total edge length of the triangulation. When $\lambda \in (0,1)$ and k is fixed, we prove a tight upper bound of order n^2 for the mixing time of the edge-flip Glauber dynamics. Combined with the previously known lower bound of order $\exp(\Omega(n^2))$ for $\lambda > 1$ [3], this establishes the existence of a dynamical phase transition for thin rectangles with critical point at $\lambda = 1$.

Keywords: lattice triangulation; Glauber dynamics; mixing times.
AMS MSC 2010: 60K35.
Submitted to EJP on May 22, 2015, final version accepted on April 4, 2016.

1 Introduction

Consider an $n \times k$ lattice rectangle $\Lambda^0_{n,k} = \{0,1,\dots,n\} \times \{0,1,\dots,k\}$ in the plane. A triangulation of $\Lambda^0_{n,k}$ is defined as a maximal set of non-crossing edges (straight line segments), each of which connects two points of $\Lambda^0_{n,k}$ and passes through no other point. See Figure 1 for an example.

Call $\Omega(n,k)$ the set of all triangulations of $\Lambda^0_{n,k}$. All $\sigma \in \Omega(n,k)$ have the same number of edges and the set of midpoints of the edges of σ does not depend on σ. Thus, we may view $\sigma \in \Omega(n,k)$ as a collection of edges $\{\sigma_x,\ x \in \Lambda_{n,k}\}$ indexed by $\Lambda_{n,k}$, where

$$\Lambda_{n,k} := \{0, \tfrac{1}{2}, 1, \tfrac{3}{2}, \dots, n - \tfrac{1}{2}, n\} \times \{0, \tfrac{1}{2}, 1, \tfrac{3}{2}, \dots, k - \tfrac{1}{2}, k\} \setminus \Lambda^0_{n,k},$$

is the set of all midpoints. Moreover, any element $\sigma \in \Omega(n,k)$ is *unimodular*, i.e., each triangle in σ has area $\frac{1}{2}$; see, e.g., [8, 6, 3] for these standard structural properties. If an

*Department of Mathematics and Physics, University of Roma Tre, Largo San Murialdo 1, 00146 Roma, Italy. E-mail: caputo@mat.uniroma3.it

†Department of Mathematics and Physics, University of Roma Tre, Largo San Murialdo 1, 00146 Roma, Italy. E-mail: martin@mat.uniroma3.it

‡Computer Science Division, University of California, Berkeley CA 94720-1776, U.S.A. E-mail: sinclair@cs.berkeley.edu. Supported in part by NSF grant 1420934.

§Department of Mathematical Sciences, University of Bath, U.K. E-mail: a.stauffer@bath.ac.uk. Supported in part by a Marie Curie Career Integration Grant PCIG13-GA-2013-618588 DSRELIS.

Figure 1: *Two triangulations of a* 5×3 *rectangle*

edge σ_x of σ is the diagonal of a parallelogram, then it is said to be *flippable*: one can delete this edge and add the opposite diagonal to obtain a new triangulation $\sigma' \in \Omega(n,k)$. In this case σ, σ' differ by a single *diagonal flip* and are said to be *adjacent*. The corresponding graph with vertex set $\Omega(n,k)$, and edges between adjacent triangulations, called the *flip graph*, is known to be connected and to have interesting structural properties; see [8, 3] and references therein.

We consider the following model of random triangulations. Fix $\lambda \in (0,\infty)$ and define a probability measure μ on $\Omega(n,k)$ by

$$\mu(\sigma) = \frac{\lambda^{|\sigma|}}{Z},$$

where $Z = \sum_{\sigma' \in \Omega(n,k)} \lambda^{|\sigma'|}$ and $|\sigma|$ is the total ℓ_1 length of the edges in σ, i.e., the sum of the horizontal and vertical lengths of each edge. The case $\lambda = 1$ is the uniform distribution, while $\lambda < 1$ (respectively, $\lambda > 1$) favors triangulations with shorter (respectively, longer) edges. We refer to [3] and references therein for background and motivation concerning this choice of weights.

A natural way to simulate triangulations distributed according to μ is to use the edge-flip *Glauber dynamics* defined as follows. In state σ, pick a midpoint $x \in \Lambda_{n,k}$ uniformly at random; if the edge σ_x is flippable to edge σ'_x (producing a new triangulation σ'), then flip it with probability

$$\frac{\mu(\sigma')}{\mu(\sigma') + \mu(\sigma)} = \frac{\lambda^{|\sigma'_x|}}{\lambda^{|\sigma'_x|} + \lambda^{|\sigma_x|}}, \tag{1.1}$$

else do nothing. Since the flip graph is connected, this defines an irreducible Markov chain on $\Omega(n,k)$, and the flip probabilities (1.1) ensure that the chain is reversible with respect to μ. Hence the dynamics converges to the stationary distribution μ. We analyze convergence to stationarity via the standard notion of *mixing time*, defined by

$$T_{\mathrm{mix}} = \inf\left\{t \in \mathbb{N} : \max_{\sigma \in \Omega(n,k)} \|p^t(\sigma,\cdot) - \mu\| \le 1/4\right\},$$

where $p^t(\sigma,\cdot)$ denotes the distribution after t steps when the initial state is σ, and $\|\nu - \mu\| = \frac{1}{2}\sum_{\sigma \in \Omega(n,k)} |\nu(\sigma) - \mu(\sigma)|$ is the usual total variation distance between two distributions μ, ν.

As discussed in [3], there is empirical evidence that the value $\lambda = 1$ represents a *critical point* separating the *sub-critical regime* $\lambda \in (0,1)$, characterized by rapid decay of both equilibrium and dynamical correlations, from the *super-critical regime* $\lambda > 1$, characterized by the emergence of long-range correlations and a dramatic slowdown in the convergence to equilibrium. We substantiated this picture by showing that there exist constants $C > 0$ and $\lambda_1 \in (0,1)$ such that

$$T_{\mathrm{mix}} \le Ckn(k+n),$$

for all $k, n \in \mathbb{N}$ and for all $\lambda \le \lambda_1$; see [3, Theorem 5.1]. This estimate is based on a coupling argument that requires λ to be sufficiently small; in particular, $\lambda_1 = 1/8$

suffices. We conjectured in [3] that the mixing time should satisfy $T_{\mathrm{mix}} = O(kn(k+n))$ throughout the sub-critical regime $\lambda \in (0,1)$. However, except for the special case $k = 1$, establishing even an arbitrary polynomial bound on T_{mix} in the whole region $\lambda < 1$ has turned out to be very challenging. Regarding the super-critical regime, by [3, Theorem 6.1 and Theorem 6.2] it is known that, for $\lambda > 1$, one has $T_{\mathrm{mix}} = \exp(\Omega(k+n))$ for all k, n, and that $T_{\mathrm{mix}} = \exp(\Omega(n^2/k))$ if $n > k^2$.

In this paper we establish the conjectured behavior for all $\lambda < 1$ in the case of "thin" rectangles, i.e., the case when k is fixed and n is large.

Theorem 1.1. *For any $\lambda \in (0,1)$, $k \in \mathbb{N}$, there exists a constant $C = C(\lambda, k) > 0$ such that the mixing time of the Glauber dynamics for $n \times k$ triangulations satisfies $T_{\mathrm{mix}} \leq C\, n^2$ for all $n \geq 1$.*

We remark that the above bound is sharp up to the value of the constant C since it is known that $T_{\mathrm{mix}} \geq C_0 kn(k+n)$ for some positive constant C_0 for any $k, n \in \mathbb{N}$ and any $\lambda > 0$; see [3, Proposition 6.3]. However, as a function of k the constant C in Theorem 1.1 can be exponentially large, and thus the interest of this bound is limited to the case of thin rectangles.

In the special case $k = 1$, the above theorem can be obtained by a direct coupling argument; see [3, Theorem 5.3]. Moreover, it is interesting to observe that in the case $k = 1$ the set of triangulations is in 1-1 correspondence with the set of configurations of a lattice path, and that diagonal flips are equivalent to so-called mountain/valley flips in the lattice path representation. Weighted versions of lattice path models have been studied extensively in the past (see, e.g., [4, 7]), and it is tempting to analyze the $n \times k$ triangulation model as a multi-path system with k interacting lattice paths. While this can be done in principle, it turns out that the interaction between the paths is technically very complex. Even the case $k = 2$ apparently does not allow for significant simplification with this representation.

The proof of Theorem 1.1 will rely crucially on some recent developments by one of us [13] based on a Lyapunov function approach to the sub-critical regime $\lambda \in (0,1)$. As detailed in subsequent sections, the main results of [13] will be used first to show that after $T = O(n^2)$ steps of the chain we can reduce the problem to a restricted chain on a "good" set of triangulations, each edge of which never exceeds logarithmic length, and then to show that distant regions in our thin rectangles can be decoupled with an exponentially small error. This will enable us to set up a recursive scheme for functional inequalities related to mixing time such as the *logarithmic Sobolev inequality*. The recursion, based on a bisection approach for the relative entropy functional inspired by the spin system analysis of [10, 5], allows us to reduce the scale from $n \times k$ down to $\mathrm{polylog}(n) \times k$. Once we reach the $\mathrm{polylog}(n) \times k$ scale, we use a refinement from [2] of the classical *canonical paths argument* [12]. This allows one to obtain an upper bound on the relaxation time of a Markov chain in terms of the congestion ratio restricted to a subspace Ω' and the time the chain needs to visit Ω' with large probability. Here we use a further crucial input from [13] permitting us to identify a "canonical" subset of triangulations Ω' such that after $T = O(n^2)$ the chain enters Ω' with large probability and such that the chain restricted to Ω' has small congestion ratio. A detailed high-level overview of the proof will be given in Section 4.1.

The rest of the paper is organized as follows. In Section 2, we first recall some important tools from [3] and then formulate the main ingredients we need from [13]. Then, in Section 3 we develop the applications of improved canonical path techniques to our setting. In Section 4 we discuss the recursive scheme for the log-Sobolev inequality and prove Theorem 1.1.

2 Main tools

2.1 Triangulations with boundary conditions

We will often consider subsets of $\Omega(n,k)$ consisting of triangulations in which some edges are kept fixed, or "frozen"; we call these *constraint edges*. Formally, let $\Lambda' \subset \Lambda_{n,k}$ denote a subset of the midpoints, and fix a collection of non-crossing edges $\{\tau_y,\ y \in \Lambda'\}$, i.e., straight lines with midpoints in Λ' each of which connects two points of $\Lambda^0_{n,k}$ and passes through no other point of $\Lambda^0_{n,k}$. If $\sigma \in \Omega(n,k)$ satisfies $\{\sigma_y = \tau_y\ ,\ y \in \Lambda'\}$, we say that σ is *compatible* with the constraint edges τ. We interpret the constraint edges τ as a *boundary condition*.

We shall actually need a more general notion of boundary condition, in order to deal with the possibility of constraint edges whose midpoints lie outside the rectangle $\Lambda^0_{n,k}$. Let N be an integer and consider the set $Q^0_{N,n,k} = \{-N,\dots,n+N\} \times \{0,\dots,k\}$, i.e., a $(2N+n) \times k$ rectangle containing $\Lambda^0_{n,k}$, and let $Q_{N,n,k}$ denote the set of midpoints of a triangulation of $Q^0_{N,n,k}$. Fix a triangulation $\widehat{\tau}$ of the region $Q^0_{N,n,k}$ and call τ the set of edges obtained from $\widehat{\tau}$ by deleting some or all edges $\widehat{\tau}_x$ with midpoint $x \in \Lambda_{n,k}$. Thus, τ is a set of constraint edges for triangulations of $Q^0_{N,n,k}$ such that all edges with midpoints in $Q_{N,n,k} \setminus \Lambda_{n,k}$ are assigned. Given constraint edges τ as above, we define $\Omega^\tau(n,k)$ as the set of all triangulations σ of $Q^0_{N,n,k}$ that are compatible with τ. Since the parameter N will play no essential role in what follows we often omit it from our notation. Since all elements of $\Omega^\tau(n,k)$ have the same edges at midpoints in $Q_{N,n,k} \setminus \Lambda_{n,k}$, one can also view a triangulation $\sigma \in \Omega^\tau(n,k)$ as an assignment of edges to midpoints in $\Lambda_{n,k}$ with certain constraints. Note that while the midpoint of a non-constraint edge of a triangulation $\sigma \in \Omega^\tau(n,k)$ is always contained in $\Lambda_{n,k}$, its endpoints need not be contained in $\Lambda^0_{n,k}$; we refer to Lemma 3.4 below for a quantitative statement on the smallest rectangle containing all non-constraint edges of any $\sigma \in \Omega^\tau(n,k)$ in terms of the length of the largest edge in τ.

The random triangulation σ with boundary condition τ is the random variable $\sigma \in \Omega^\tau(n,k)$ with distribution

$$\mu^\tau(\sigma) = \frac{\lambda^{|\sigma|}}{Z}, \tag{2.1}$$

where $Z = \sum_{\sigma' \in \Omega^\tau(n,k)} \lambda^{|\sigma'|}$. We sometimes write μ instead of μ^τ and Ω instead of $\Omega^\tau(n,k)$ if there is no need to stress the dependence on the constraint edges. We say that there is *no boundary condition* when $N = 0$ and the set of constraint edges τ is empty. In this case $\Omega^\tau(n,k)$ coincides with $\Omega(n,k)$, the set of all triangulations of $\Lambda^0_{n,k}$.

2.2 Ground states

It is a fact that for any set of constraint edges τ, the set of triangulations $\Omega^\tau(n,k)$ that are compatible with τ is non-empty. Among the compatible triangulations, we are particularly interested in those with minimal ℓ_1-edge length, which we call *ground state triangulations*. These are the triangulations of maximum weight in (2.1) when $\lambda < 1$, and they play a central role in our analysis. In the absence of boundary conditions, the ground state triangulations are trivial: every edge is either horizontal or vertical or a unit diagonal, so in particular the ground state is unique up to flipping of the unit diagonals. The presence of constraint edges can change the ground state considerably. However, the following result from [3, Lemma 3.4] reveals the strikingly simple structure of ground states for any set of contraints.

Lemma 2.1. [Ground State Lemma] *Given any set of constraint edges, the ground state triangulation is unique (up to possible flipping of unit diagonals), and can be constructed by placing each edge in its minimal length configuration consistent with the constraints, independent of the other edges.*

Given a set of constraint edges, we denote by $\bar{\sigma}$ the unique ground state triangulation. (An arbitrary choice of the available unit diagonal orientations is understood in this notation.) If no confusion arises, we omit to specify the dependence on the constraint edges. An important structural property of triangulations with constraint edges, which follows from Lemma 2.1, is that from any triangulation σ compatible with τ one can reach the ground state $\bar{\sigma}$ with a path in the flip graph with the property that no flip increases the length of an edge.

2.3 The Glauber dynamics

The Glauber dynamics in the presence of a boundary condition τ is defined as before (see equation (1.1)), with the modification that the midpoint x to be updated is picked uniformly at random among all midpoints of non-constraint edges. For any $\lambda > 0$, this defines an irreducible Markov chain on $\Omega^\tau(n, k)$ that is reversible w.r.t. the stationary distribution μ^τ (see [3] for details). It was shown in [3, Theorem 5.1] that for some constants $C > 0$ and $\lambda_1 \in (0, 1)$, the mixing time of this chain in an $n \times k$ rectangle satisfies $T_{\mathrm{mix}} \leq Ckn(k + n)$ uniformly in the choice of the constraint edges, whenever $\lambda \leq \lambda_1$. We also conjectured in [3] that the $O(kn(k + n))$ mixing time should hold for all $\lambda \in (0, 1)$.

2.4 Key ingredients from [13]

We gather in Lemmas 2.2–2.5 below some estimates from [13] that will be crucial in our analysis; for the proofs see [13]. Note that these estimates are valid throughout the sub-critical regime $\lambda \in (0, 1)$.

The first lemma applies to the case where there are no constraint edges, so that the ground state is trivial. It follows from [13, Corollary 7.4], and establishes that after running the Markov chain for $O(n^2)$ steps, the ℓ_1-length of a given edge has an exponential tail. For a given initial triangulation $\sigma = \sigma^0$, we denote by σ^t the triangulation after t steps of the chain, and denote by $\mathbb{P}$ the probability measure induced by the evolution of the chain.

Lemma 2.2. *Fix $\lambda \in (0, 1)$. There exist positive constants $c_1 = c_1(\lambda)$ and $c_2 = c_2(\lambda)$ such that for $n \geq k \geq 1$, for any $t \geq c_1 n^2$, any $\ell > 0$, any midpoint $x \in \Lambda_{n,k}$, and any initial triangulation $\sigma \in \Omega(n, k)$:*

$$\mathbb{P}\left(|\sigma_x^t| \geq \ell\right) \leq c_1 \exp(-c_2 \ell).$$

The next lemma deals with the evolution in the presence of constraint edges τ, and follows from [13, Theorem 7.3]. We denote by $\bar{\sigma}_x$ the ground state edge at x (compatible with τ). Given $\sigma \in \Omega^\tau(n, k)$ and $y \in \Lambda_{n,k}$, we write $\sigma_y \cap \bar{\sigma}_x \neq \emptyset$ if the edge σ_y crosses $\bar{\sigma}_x$ (not including the case where σ_y and $\bar{\sigma}_x$ intersect only at their endpoints).

Lemma 2.3. *Fix $\lambda \in (0, 1)$. There exist positive constants $c_1 = c_1(\lambda)$ and $c_2 = c_2(\lambda)$ such that the following holds for any $n \geq k \geq 1$, any set of constraint edges τ, and any midpoint $x \in \Lambda_{n,k}$. Let M be the ℓ_1 length of the largest edge in any triangulation $\sigma \in \Omega^\tau(n, k)$. Then, for any $t \geq c_1 kn(M + \log n)$, and any $\ell \geq 0$, we have*

$$\mathbb{P}\left(\bigcup_{y \in \Lambda_{n,k}} \left\{\sigma_y^t \cap \bar{\sigma}_x \neq \emptyset\right\} \cap \left\{|\sigma_y^t| \geq |\bar{\sigma}_x| + \ell\right\}\right) \leq c_1 \exp(-c_2 \ell). \tag{2.2}$$

Next we give a rough upper bound on the number of small edges intersecting a given ground state edge. We assume that a set of constraint edges τ is given. For any triangulation $\sigma \in \Omega^\tau(n, k)$, any ground state edge g, and any $\ell \in \mathbb{Z}^+$, define

$$I_g(\sigma, \ell) = \{\sigma_x\,,\ x \in \Lambda_{n,k} \colon\ \sigma_x \cap g \neq \emptyset \text{ and } |\sigma_x| \leq |g| + \ell\}\,.$$

We denote by $|I_g(\sigma,\ell)|$ the cardinality of $I_g(\sigma,\ell)$. For a proof of the lemma below, see [13, Proposition 4.4].

Lemma 2.4. *Let g be a ground state edge, and let $\sigma \in \Omega^\tau(n,k)$ be a triangulation.*

i) If $\sigma_x \cap g \neq \emptyset$ then $|\sigma_x| \geq |g|$, with strict inequality when the midpoint of g is not x.

ii) For any $\ell \geq 1$, all midpoints of edges in $I_g(\sigma,\ell)$ are contained in the ball of radius 2ℓ centered at the midpoint of g.

iii) There exists a universal $c > 0$ such that for any $\ell \geq 1$ we have

$$|I_g(\sigma,\ell)| \leq c\,\ell^2\,, \quad \text{and} \quad \Big|\bigcup_\sigma I_g(\sigma,\ell)\Big| \leq c\,\ell^4.$$

Finally, the lemma below establishes the probability of having a top-to-bottom crossing of unit verticals in a random triangulation σ. By a "top-to-bottom crossing of unit verticals in σ" we mean a straight line of length k made up of k vertical edges in σ each of length 1. The lemma below follows from [13, Theorems 8.1 and 8.2].

Lemma 2.5. *Let $k \in \mathbb{N}$ and $\lambda \in (0,1)$ be fixed. There exist positive constants $c = c(\lambda,k)$, $\delta = \delta(\lambda,k)$ and $m_0 = m_0(\lambda,k)$ such that the following holds. Let R be an $m \times k$ rectangle contained in $[0,n] \times [0,k]$ with $m \geq m_0$. Consider an arbitrary set of constraint edges τ such that no edge from τ intersects R. For any triangulation $\sigma \in \Omega^\tau(n,k)$, let $C_R(\sigma)$ be the number of disjoint top-to-bottom crossings of unit verticals from σ that are inside R. Then,*

$$\mu^\tau\left(C_R(\sigma) \leq \delta\, m\right) \leq e^{-c\,m}.$$

Furthermore, let σ,σ' be two triangulations sampled from the stationary distribution μ given two different sets of constraint edges τ,τ' such that no edge of τ,τ' intersects R. Then, there exists a coupling of σ,σ' such that the probability that they have less than $\delta\,m$ common top-to-bottom crossings of unit verticals is at most $e^{-c\,m}$.

3 Estimates via canonical paths

We recall that the relaxation time T_{rel} is defined as the inverse of the spectral gap of the Markov chain. We start by showing that a direct application of the usual canonical path argument [12] yields an exponential bound on the relaxation time of the Markov chain that is valid for all $\lambda \leq 1$. We recall the well known estimate relating T_{rel} and T_{mix} (see, e.g., [9, Theorem 12.3]):

$$T_{\mathrm{mix}} \leq T_{\mathrm{rel}}(2 + \log(1/\mu_*)), \tag{3.1}$$

where $\mu_* = \min_\sigma \mu(\sigma)$.

Theorem 3.1. *There exists a positive constant C such that for any $\lambda \leq 1$, $n,k \in \mathbb{N}$ and any set of constraint edges τ, the Glauber dynamics on $\Omega^\tau(n,k)$ satisfies*

$$T_{\mathrm{rel}} \leq \exp(Ckn).$$

Before proving the above theorem we recall a useful structural fact. Given a set of constraint edges τ and a midpoint x, consider the set Ω_x^τ of possible values of σ_x, as σ ranges in $\Omega^\tau(n,k)$. Two edges $\sigma_x,\sigma'_x \in \Omega_x^\tau$ are said to be *neighbors* if σ_x is flippable to σ'_x within some triangulation $\sigma \in \Omega^\tau(n,k)$. Then it is known (see, e.g., [3]) that the induced graph with vertex set Ω_x^τ is a tree $\mathcal{G}_x^\tau$, and that for each edge $\sigma_x \in \Omega_x^\tau$ not in ground state there is a unique edge $\sigma'_x \in \Omega_x^\tau$ for which $|\sigma_x| > |\sigma'_x|$ and σ_x,σ'_x are neighbors. It is useful to see $\mathcal{G}_x^\tau$ as a tree rooted at the ground state edges of midpoint x; recall that if there are more than one ground state edges of midpoint x, then there are exactly two and they are the opposite unit diagonals, which are neighbors in $\mathcal{G}_x^\tau$. We then obtain that the shortest path between σ_x and a ground state edge of midpoint x in $\mathcal{G}_x^\tau$ is given by a sequence of edges of decreasing lengths.

In order to describe the shortest path between any two edges $\sigma_x, \sigma'_x \in \Omega_x^\tau$ in $\mathcal{G}_x^\tau$, we introduce one piece of notation. If there is only one ground state edge of midpoint x, let $\upsilon(\sigma_x, \sigma'_x) = \upsilon(\sigma'_x, \sigma_x)$ be the lowest common ancestor of σ_x and σ'_x in the tree $\mathcal{G}_x^\tau$. If there are two ground state edges of midpoint x, but σ_x and σ'_x belong to the same subtree rooted at one of the ground state edges, then as before we define $\upsilon(\sigma_x, \sigma'_x) = \upsilon(\sigma'_x, \sigma_x)$ to be the lowest common ancestor of σ_x and σ'_x. However, if σ_x and σ'_x belong to different subtrees, then we let $\upsilon(\sigma_x, \sigma'_x)$ (resp., $\upsilon(\sigma'_x, \sigma_x)$) be the closest ground state edge to σ_x (resp., σ'_x) in $\mathcal{G}_x^\tau$. Then the shortest path from σ_x to σ'_x in $\mathcal{G}_x^\tau$ is composed of a sequence of edges of decreasing lengths from σ_x to $\upsilon(\sigma_x, \sigma'_x)$, and a sequence of edges of increasing lengths from $\upsilon(\sigma'_x, \sigma_x)$ to σ'_x.

We will make use of the following technical lemma; see [3, Proposition 3.8] for the proof. The lemma below establishes that, given any two triangulations $\sigma, \sigma' \in \Omega^\tau(n,k)$ and any midpoint x, if we observe the different edges of midpoint x that are obtained in the shortest path between σ and σ' in the flip graph, then we encounter the same sequence of edges in the shortest path between σ_x and σ'_x in the tree $\mathcal{G}_x^\tau$.

Lemma 3.2. *Fix a set of constraint edges τ. For any midpoint x and any two triangulations $\sigma, \sigma' \in \Omega^\tau(n,k)$, the distance between σ and σ' in the flip graph is equal to $\sum_{x \in \Lambda_{n,k}} \kappa(\sigma_x, \sigma'_x)$, where $\kappa(\sigma_x, \sigma'_x)$ is the distance between σ_x and σ'_x in the tree $\mathcal{G}_x^\tau$.*

We denote by an *edge-decreasing flip* (resp., edge-increasing flip) a flip that causes an edge to decrease (resp., increase) its length. The proof of the above lemma, given in [3, Proposition 3.8], constructs such a path by performing a sequence of edge-decreasing flips to σ and σ' until obtaining triangulations η, η', respectively, such that $\eta_x = \upsilon(\sigma_x, \sigma'_x)$ and $\eta'_x = \upsilon(\sigma'_x, \sigma_x)$ for all x. This establishes a path from σ to σ' which is first composed by a sequence of edge-decreasing flips (transforming σ into η), followed by flips of unit diagonals (transforming η into η'), and then by a sequence of edge-increasing flips (transforming η' into σ').

Proof of Theorem 3.1. For each pair $\sigma, \sigma' \in \Omega^\tau(n,k)$, let $\Gamma_{\sigma,\sigma'}$ be a shortest path between σ and σ' in the flip graph. From Lemma 3.2 and the properties discussed right after Lemma 3.2, we have that for any triangulation η in the path $\Gamma_{\sigma,\sigma'}$ and any midpoint x,

$$|\eta_x| \le |\sigma_x| \vee |\sigma'_x|. \tag{3.2}$$

We can also assume that $\Gamma_{\sigma,\sigma'}$ is a *monotone* path in the sense that it is composed of a sequence of edge-decreasing flips, until for each midpoint x the edge with that midpoint is $\upsilon(\sigma_x, \sigma'_x)$, followed by a sequence of flips of unit diagonals, and then a sequence of edge-increasing flips.

Now, for any function $f : \Omega \to \mathbb{R}$, we have

$$f(\sigma) - f(\sigma') = \sum_{(\eta,\eta') \in \Gamma_{\sigma,\sigma'}} \nabla_{\eta,\eta'} f,$$

where we employ the notation $\nabla_{\eta,\eta'} f = f(\eta) - f(\eta')$. For simplicity, below we write μ instead of μ^τ and Ω instead of $\Omega^\tau(n,k)$. Thus, using Cauchy-Schwarz, the variance of f with respect to μ satisfies

$$\begin{aligned} \mathrm{Var}(f) &= \frac{1}{2} \sum_{\sigma,\sigma'} \mu(\sigma)\mu(\sigma')(f(\sigma) - f(\sigma'))^2 \\ &\le \frac{1}{2}\, \mathcal{C}(\Omega) \sum_{\eta,\eta' :\, \eta \sim \eta'} \mu(\eta) p(\eta,\eta')(\nabla_{\eta,\eta'} f)^2, \end{aligned}$$

where $p(\eta,\eta')$ is the probability that the Glauber chain goes from η to η' in one step, $\eta\sim\eta'$ denotes that η and η' are adjacent triangulations, and we use the notation

$$\mathcal{C}(\Omega)=\max_{\eta,\eta'\,:\,\eta\sim\eta'}\sum_{\sigma,\sigma'\,:\,(\eta,\eta')\in\Gamma_{\sigma,\sigma'}}\frac{\mu(\sigma)\mu(\sigma')}{\mu(\eta)p(\eta,\eta')}|\Gamma_{\sigma,\sigma'}|,\tag{3.3}$$

for the so-called "congestion ratio." Now assume that $p(\eta,\eta')\geq p(\eta',\eta)$, otherwise use reversibility to write $\mu(\eta)p(\eta,\eta')$ as $\mu(\eta')p(\eta',\eta)$. With this assumption we have that $p(\eta,\eta')\geq\frac{1}{2|\Lambda_{n,k}|}$. Also, from Lemma 3.2 we have $|\Gamma_{\sigma,\sigma'}|=O(nk(n+k))$. This holds because there are $O(nk)$ midpoints, the largest length of an edge is $O(n+k)$, and each flip that is not between two ground state edges must change the length of an edge by at least 1. The key property we use is that (3.2) gives

$$\frac{\mu(\sigma)\mu(\sigma')}{\mu(\eta)}=Z^{-1}\prod_x\lambda^{|\sigma_x|+|\sigma'_x|-|\eta_x|}\leq Z^{-1}\prod_x\lambda^{|\sigma_x|\wedge|\sigma'_x|}\leq 1,$$

where we used the bound

$$Z\geq\prod_x\lambda^{|\bar\sigma_x|}\geq\prod_x\lambda^{|\sigma_x|\wedge|\sigma'_x|}.$$

Plugging this into (3.3), we obtain

$$\mathcal{C}(\Omega)\leq Cnk(n+k)\,|\Lambda_{n,k}|\,|\Omega^\tau(n,k)|^2.\tag{3.4}$$

Using Anclin's bound [1] one has $|\Omega^\tau(n,k)|\leq 2^{|\Lambda_{n,k}|}$. The proof is then concluded by recalling that $T_{\rm rel}$ is the smallest constant γ such that the inequality

$$\mathrm{Var}(f)\leq\frac{\gamma}{2}\sum_{\eta,\eta'\,:\,\eta\sim\eta'}\mu(\eta)p(\eta,\eta')(\nabla_{\eta,\eta'}f)^2$$

holds for all functions $f:\Omega^\tau(n,k)\mapsto\mathbb{R}$. □

3.1 An improved canonical paths argument

Here we establish a first polynomial bound on the relaxation time. The result here can be formulated as follows.

Theorem 3.3. *Fix $\lambda\in(0,1)$ and $k\in\mathbb{N}$. There exists a positive constant $c=c(\lambda,k)$ such that for any boundary condition $\tau=\{\tau_x\}$ such that $|\tau_x|\leq n/4$ for all x, the relaxation time of the Glauber chain in $\Omega^\tau(n,k)$ satisfies*

$$T_{\rm rel}\leq n^c.$$

The strategy of the proof is as follows. We shall identify a subset Ω' of triangulations such that the congestion ratio $\mathcal{C}(\Omega')$ defined as in (3.3) but restricted to Ω' satisfies a polynomial bound, in contrast with the exponential bound in (3.4). Using a key input from [13], we show that the Glauber chain enters the set Ω' with large probability after a *burn-in* time of $T=O(n^2)$ steps. Following an idea already used in [2] we establish the desired upper bound on $T_{\rm rel}$ by combining the above facts.

We start with a deterministic estimate.

Lemma 3.4. *Fix $L>0$ and let $\sigma\in\Omega^\tau(n,k)$ be a triangulation of the $n\times k$ rectangle with boundary condition $\tau=\{\tau_x\}$ such that $|\tau_x|\leq L$ for all x. Then, all edges of σ are contained in the rectangle $[-L,n+L]\times[-L,k+L]$.*

Proof. First, note that the ground state triangulation must satisfy the lemma, because all edges have size at most L. We use the notation σ^x to denote the triangulation obtained from σ by flipping σ_x, and call σ_x an *increasing edge* if it can be flipped and σ_x^x is larger than σ_x. Now it is enough to show that there cannot be an increasing edge σ_x with $x \in \Lambda_{n,k}$ such that $\sigma_x^x \not\subset [-L, n+L] \times [-L, k+L]$ but all edges of σ are inside $[-L, n+L] \times [-L, k+L]$. In order to achieve a contradiction, assume that such an increasing edge σ_x exists and assume that σ_x^x is at the left part of the triangulation (i.e., that its leftmost endpoint has horizontal coordinate smaller than $-L$). Let σ_y, σ_z be the triangle containing σ_x such that the vertex $v = \sigma_y \cap \sigma_z$ has horizontal coordinate smaller than $-L$. Since σ is completely inside $[-L, n+L] \times [-L, k+L]$, we obtain that σ_y and σ_z are constraint edges. Also, since $x \in \Lambda$, σ_x must have one endpoint u of horizontal coordinate at least 0. This gives that $\|v - u\|_1 > L$, and consequently, either σ_y or σ_z has length larger than L, which is a contradiction. $\square$

Next, we formulate a general upper bound on T_{rel} in terms of the congestion ratio of a subset Ω' of the state space Ω, a time T, and the probability needed to reach Ω' within time T. A version of this lemma appears in [2, Theorem 2.4]. For the reader's convenience we give a detailed proof.

Lemma 3.5 (Canonical paths with burn-in time)**.** *Consider a Markov chain with state space Ω, irreducible transition matrix $p(\cdot,\cdot)$ and reversible probability measure μ. Let $\Omega' \subset \Omega$ be a subset so that between each $\sigma, \sigma' \in \Omega'$ there is a path $\Gamma_{\sigma,\sigma'}$ in the Markov chain that is entirely contained in Ω'. Define the congestion ratio*

$$\mathcal{C}(\Omega') = \max_{\eta,\eta' \in \Omega' \colon \eta \sim \eta'} \sum_{\sigma,\sigma' \colon (\eta,\eta') \in \Gamma_{\sigma,\sigma'}} \frac{\mu(\sigma)\mu(\sigma')|\Gamma_{\sigma,\sigma'}|}{\mu(\eta)p(\eta,\eta')},$$

where the sum is over all pairs of states $\sigma, \sigma' \in \Omega'$ so that the path $\Gamma_{\sigma,\sigma'}$ uses the transition (η, η'). Fix $T \in \mathbb{N}$ and let ρ be a lower bound on the probability that at time T the chain is inside Ω', uniformly over the starting state in Ω. Then the relaxation time satisfies

$$T_{\mathrm{rel}} \leq \frac{6\,T^2}{\rho} + \frac{3\,\mathcal{C}(\Omega')}{\rho^2}.$$

Proof. We run the Markov chain for T steps. For $\sigma, \tau \in \Omega$, let $\mu_\sigma(\tau)$ be the probability that, starting from σ, the Markov chain is at τ after T steps. Note that $\mu_\sigma(\Omega') \geq \rho$. For $\sigma, \tau \in \Omega$, and for any path γ of length T in the chain starting at σ and ending at τ, let $\nu_{\sigma,\tau}(\gamma)$ be the conditional probability that, given the initial state σ at time 0 and the final state τ after T steps, the Markov chain traverses the path γ. Then, for any function $f \colon \Omega \to \mathbb{R}$, we have

$$\begin{aligned}\mathrm{Var}(f) &= \frac{1}{2}\sum_{\sigma,\sigma' \in \Omega} \mu(\sigma)\mu(\sigma')(f(\sigma) - f(\sigma'))^2 = \frac{1}{2}\sum_{\sigma,\sigma' \in \Omega}\sum_{\eta,\eta' \in \Omega'} \mu(\sigma)\mu(\sigma')\frac{\mu_\sigma(\eta)\mu_{\sigma'}(\eta')}{\mu_\sigma(\Omega')\mu_{\sigma'}(\Omega')} \times \\ &\quad \times \sum_{\gamma_1,\gamma_2} \nu_{\sigma,\eta}(\gamma_1)\nu_{\sigma',\eta'}(\gamma_2)\Big(\textstyle\sum_{e\in\gamma_1}\nabla_e f + \sum_{e\in\gamma_2}\nabla_e f + \sum_{e\in\Gamma_{\eta,\eta'}}\nabla_e f\Big)^2,\end{aligned}$$

where the three sums inside the parenthesis are over the edges of the paths γ_1, γ_2, and $\Gamma_{\eta,\eta'}$, respectively, with γ_1 being a path from σ to η and γ_2 being a path from η' to σ'. Then, applying Cauchy-Schwarz, we obtain

$$\begin{aligned}\mathrm{Var}(f) &\leq \frac{3}{2}\sum_{\sigma,\sigma' \in \Omega}\sum_{\eta,\eta' \in \Omega'} \mu(\sigma)\mu(\sigma')\frac{\mu_\sigma(\eta)\mu_{\sigma'}(\eta')}{\mu_\sigma(\Omega')\mu_{\sigma'}(\Omega')} \times \\ &\quad \times \sum_{\gamma_1,\gamma_2} \nu_{\sigma,\eta}(\gamma_1)\nu_{\sigma',\eta'}(\gamma_2)\Big(T\textstyle\sum_{e\in\gamma_1}(\nabla_e f)^2 + T\sum_{e\in\gamma_2}(\nabla_e f)^2 + |\Gamma_{\eta,\eta'}|\sum_{e\in\Gamma_{\eta,\eta'}}(\nabla_e f)^2\Big).\end{aligned}$$

We write the right-hand side above as $A_1 + A_2 + A_3$, where

$$A_1 = \frac{3}{2} \sum_{\sigma,\sigma' \in \Omega} \sum_{\eta,\eta' \in \Omega'} \mu(\sigma)\mu(\sigma') \frac{\mu_\sigma(\eta)\mu_{\sigma'}(\eta')}{\mu_\sigma(\Omega')\mu_{\sigma'}(\Omega')} \sum_{\gamma_1,\gamma_2} \nu_{\sigma,\eta}(\gamma_1)\nu_{\sigma',\eta'}(\gamma_2)\, T \sum_{e\in\gamma_1} (\nabla_e f)^2$$

$$A_2 = \frac{3}{2} \sum_{\sigma,\sigma' \in \Omega} \sum_{\eta,\eta' \in \Omega'} \mu(\sigma)\mu(\sigma') \frac{\mu_\sigma(\eta)\mu_{\sigma'}(\eta')}{\mu_\sigma(\Omega')\mu_{\sigma'}(\Omega')} \sum_{\gamma_1,\gamma_2} \nu_{\sigma,\eta}(\gamma_1)\nu_{\sigma',\eta'}(\gamma_2)\, T \sum_{e\in\gamma_2} (\nabla_e f)^2$$

$$A_3 = \frac{3}{2} \sum_{\sigma,\sigma' \in \Omega} \sum_{\eta,\eta' \in \Omega'} \mu(\sigma)\mu(\sigma') \frac{\mu_\sigma(\eta)\mu_{\sigma'}(\eta')}{\mu_\sigma(\Omega')\mu_{\sigma'}(\Omega')} \sum_{\gamma_1,\gamma_2} \nu_{\sigma,\eta}(\gamma_1)\nu_{\sigma',\eta'}(\gamma_2)\, |\Gamma_{\eta,\eta'}| \sum_{e\in\Gamma_{\eta,\eta'}} (\nabla_e f)^2.$$

We start with A_1. Summing over γ_2, σ', η', and using $\sum_{\gamma_2} \nu_{\sigma',\eta'}(\gamma_2) = 1$, we have

$$A_1 = \frac{3}{2} T \sum_{\sigma\in\Omega} \sum_{\eta\in\Omega'} \mu(\sigma) \frac{\mu_\sigma(\eta)}{\mu_\sigma(\Omega')} \sum_{\gamma_1} \nu_{\sigma,\eta}(\gamma_1) \sum_{e\in\gamma_1} (\nabla_e f)^2.$$

Changing the order of the summations, and summing first over all pairs of adjacent states $\tau \sim \tau'$, we get

$$\begin{aligned} A_1 &= \frac{3}{2} T \sum_{\tau,\tau'\in\Omega:\ \tau\sim\tau'} \mu(\tau)p(\tau,\tau')(\nabla_{\tau,\tau'} f)^2 \sum_{\sigma\in\Omega,\eta\in\Omega',\gamma:\ (\tau,\tau')\in\gamma} \frac{\mu(\sigma)\mu_\sigma(\eta)\nu_{\sigma,\eta}(\gamma)}{\mu_\sigma(\Omega')\mu(\tau)p(\tau,\tau')} \\ &\le \frac{3T}{2\rho} \sum_{\tau,\tau'\in\Omega:\ \tau\sim\tau'} \mu(\tau)p(\tau,\tau')(\nabla_{\tau,\tau'} f)^2 \sum_{\sigma\in\Omega,\eta\in\Omega',\gamma:\ (\tau,\tau')\in\gamma} \frac{\mu(\sigma)\mu_\sigma(\eta)\nu_{\sigma,\eta}(\gamma)}{\mu(\tau)p(\tau,\tau')} \\ &\le \frac{3T}{2\rho} \sum_{\tau,\tau'\in\Omega:\ \tau\sim\tau'} \mu(\tau)p(\tau,\tau')(\nabla_{\tau,\tau'} f)^2 \frac{\mathbb{P}_\mu\left(\text{Markov chain traverses } (\tau,\tau') \text{ within } T \text{ steps}\right)}{\mu(\tau)p(\tau,\tau')} \\ &\le \frac{3T}{2\rho} \sum_{\tau,\tau'\in\Omega:\ \tau\sim\tau'} \mu(\tau)p(\tau,\tau')(\nabla_{\tau,\tau'} f)^2 \frac{T\mu(\tau)p(\tau,\tau')}{\mu(\tau)p(\tau,\tau')} = \frac{3T^2}{\rho}\mathcal{D}(f,f), \end{aligned}$$

where $\mathbb{P}_\mu(\cdot)$ denotes the measure induced by the Markov chain started from stationarity, and we use the notation

$$\mathcal{D}(f,f) = \frac{1}{2} \sum_{\tau,\tau'\in\Omega:\ \tau\sim\tau'} \mu(\tau)p(\tau,\tau')(\nabla_{\tau,\tau'} f)^2$$

for the so-called Dirichlet form. For the second term, we have by symmetry that $A_2 = A_1$. For A_3, we use $\rho \le \mu_\sigma(\Omega'), \mu_{\sigma'}(\Omega')$, and sum over $\gamma_1, \gamma_2, \sigma, \sigma'$ to obtain

$$A_3 \le \frac{3}{2\rho^2} \sum_{\eta,\eta'\in\Omega'} \mu(\eta)\mu(\eta')|\Gamma_{\eta,\eta'}| \sum_{e\in\Gamma_{\eta,\eta'}} (\nabla_e f)^2.$$

Changing the order of summations, we get

$$A_3 \le \frac{3\,\mathcal{C}(\Omega')}{\rho^2}\mathcal{D}(f,f).$$

The result now follows since T_{rel} is the smallest constant γ such that the inequality

$$\mathrm{Var}(f) \le \gamma\,\mathcal{D}(f,f)$$

holds for all functions $f : \Omega \mapsto \mathbb{R}$. □

Proof of Theorem 3.3. Let $T = c_1 n^2 k$ for some large enough constant $c_1 = c_1(\lambda) > 0$. Thanks to Lemma 3.4 we may apply Lemma 2.3 with $M = 2n + k$. Thus, for any given

$x \in \Lambda_{n,k}$ and ground-state edge $\bar{\sigma}_x$ with midpoint x, taking $\ell = c_2 \log |\Lambda_{n,k}|$ for some large enough constant $c_2 = c_2(\lambda) > 0$, and taking the union bound over all $x \in \Lambda_{n,k}$ in (2.2) we obtain that the triangulation σ^T at time T, for an arbitrary initial condition σ, satisfies

$$\mathbb{P}\left(\bigcup_{x\in\Lambda_{n,k}} \bigcup_{y\in\Lambda_{n,k}} \left\{\sigma_y^T \cap \bar{\sigma}_x \neq \emptyset\right\} \cap \left\{|\sigma_y^T| > |\bar{\sigma}_x| + \ell\right\}\right) \leq n^{-1}. \tag{3.5}$$

Let

$$\Omega' = \Big\{\sigma\colon \text{ for all } x, y \in \Lambda_{n,k}, |\sigma_x| \leq |\bar{\sigma}_x| + \ell \text{ and } \mathbf{1}\left(\sigma_y \cap \bar{\sigma}_x \neq \emptyset\right) \leq \mathbf{1}\left(|\sigma_y| \leq |\bar{\sigma}_x| + \ell\right)\Big\}.$$

Thus (3.5) implies that $\mathbb{P}\left(\sigma^T \in \Omega'\right) \geq 1 - n^{-1}$. Note that Ω' is a decreasing set in the sense that if $\sigma \in \Omega'$ then for all σ' that can be obtained from σ by performing decreasing flips, we have $\sigma' \in \Omega'$. This allows us to construct a path $\Gamma_{\sigma,\sigma'}$ within Ω' between any pair of triangulations $\sigma, \sigma' \in \Omega'$, as described right after Lemma 3.2.

We now describe the path $\Gamma_{\sigma,\sigma'}$. Fix two triangulations $\sigma, \sigma' \in \Omega'$, and any midpoint $x \in \Lambda_{n,k}$. Note that in order to transform σ_x into a ground state edge g, we only need to perform a sequence of edge-decreasing flips on edges whose midpoints are in $I_g(\sigma, \ell)$ (recall the definition of I_g from Lemma 2.4). This holds because we can construct a triangulation η by taking σ, removing all edges that cross g from σ, adding the edge g and then completing the triangulation with ground state edges given all the other edges already present. In this way, η contains all edges of σ that do not cross g, and can be obtained from σ by a sequence of edge-decreasing flips since its other edges are in ground state. Hence, since $\sigma, \sigma' \in \Omega'$, in order to transform σ_x into σ'_x we do not need to flip edges that do not cross a ground state edge g of midpoint x. So all edges that need to be flipped belong to $I_g(\sigma, \ell) \cup I_g(\sigma', \ell)$.

We will construct the path $\Gamma_{\sigma,\sigma'}$ as follows. Partition $[0, n] \times [0, k]$ into slabs of horizontal width 2ℓ, and number the slabs from left to right. We take the first slab and perform the smallest number of flips needed to σ until all edges of midpoint inside the first slab are equal to σ'. We do this by first doing a sequence of edge-decreasing flips so that, for each midpoint x in the first slab, we obtain the edge $v(\sigma_x, \sigma'_x)$. Then we perform a sequence of flips of unit diagonals and then a sequence of edge-increasing flips until for each such x we obtain the edge σ'_x. From the explanation above, all edges that needed to be flipped in this procedure are contained in $\bigcup_x (I_{g_x}(\sigma, \ell) \cup I_{g_x}(\sigma', \ell))$, where the first union is over all midpoints x inside the first slab, and g_x stands for a ground state edge of midpoint x. By Lemma 2.4 we have that all edges in $\bigcup_{\sigma\in\Omega'} I_{g_x}(\sigma, \ell)$ have midpoint inside a ball of radius 2ℓ centered at x. This implies that only edges of the first and second slabs are flipped in the procedure above. Then we iterate this procedure, finding a sequence of flips that transform σ into σ' slab by slab, from left to right, so that when transforming the ith slab, only edges with midpoints in the ith and $(i + 1)$th slabs need to be flipped, and edges with midpoint in the slabs $1, 2, \ldots, i - 1$ are already equal to σ'. Note that, in each slab, we just perform the minimum number of flips needed to transform that slab into σ', and we do that by first performing all decreasing flips and then all increasing flips.

Our goal is to apply Lemma 3.5, for which we need to bound the value of the congestion ratio $\mathcal{C}(\Omega')$. To do this, consider a pair of adjacent triangulations η, η'. Assume that η, η' differ at an edge of the ith slab. Therefore, if σ, σ' are two triangulations for which the path between them includes the transition (η, η') we know that triangulation η has slabs $1, 2, \ldots, i - 2$ equal to σ' and slabs $i + 2, i + 3, \ldots$ equal to σ. Let ξ be a partial triangulation in Ω' of the first $i - 2$ slabs and m be a partial triangulation in Ω' of the middle slabs so that ξ, m and σ are compatible, meaning that ξ, m and the edges of σ inside slabs $i + 2, i + 3, \ldots$ can coexist to form a full triangulation. Similarly, let ξ' be a partial triangulation in Ω' of the last slabs ($i+2, i+3, \ldots$) and m' be a partial triangulation in Ω' of the middle slabs so that ξ', m' and the edges of σ' inside slabs $1, 2, \ldots, i-2$ are compatible. Assume that $p(\eta, \eta') \geq p(\eta', \eta)$, which implies that $p(\eta, \eta') \geq \frac{1}{2|\Lambda_{n,k}|}$ (otherwise, replace $\mu(\eta)p(\eta, \eta')$ with $\mu(\eta')p(\eta', \eta)$ in $\mathcal{C}(\Omega')$). Let η_i be the part of η inside slabs $i-1, i, i+1$. Then, summing over all ξ, ξ', m, m' as above such that (η, η') is a transition in the path from ξ, m, σ to σ', m', ξ', and noting that the path between σ and σ' has length at most $2\ell|\Lambda_{n,k}|$, we obtain the following upper bound for $\mathcal{C}(\Omega')$:

$$\mathcal{C}(\Omega') \leq 4\ell|\Lambda_{n,k}|^2 \sum_{\xi,\xi',m,m'} \frac{\lambda^{|\xi|+|\xi'|+|m|+|m'|-|\eta_i|}}{Z_{\Omega'}},$$

where $Z_{\Omega'} = \sum_{\sigma \in \Omega'} \lambda^{|\sigma|}$. Instead of summing over m, m', we will sum over triangulations m'' of the middle slabs that are compatible with both ξ and ξ' and are to be interpreted as $m \wedge m'$. Given m'', we sum over m, m' that can be obtained from m'' by increasing flips and such that (η, η') is a transition in the path from ξ, m, σ to σ', m', ξ'. Let $A(m'', m, m', \eta)$ be the indicator that all four of them are compatible, as described above. When $A(m'', m, m', \eta) = 1$ we have that $|\eta_x| \leq |m_x| \vee |m'_x|$ for any midpoint x in the middle slabs. Hence, $|m''| + |m \setminus m''| + |m' \setminus m''| \geq |\eta_i|$, which gives

$$\mathcal{C}(\Omega') \leq 4\ell|\Lambda_{n,k}|^2 \sum_{\xi,\xi',m''} \frac{\lambda^{|\xi|+|\xi'|+|m''|}}{Z_{\Omega'}} \sum_{m,m'} \lambda^{|m''|+|m\setminus m''|+|m'\setminus m''|-|\eta_i|} A(m'', m, m', \eta).$$

Since $\lambda < 1$, we can simply use Anclin's bound [1] saying that the number of triangulations of an $\ell \times k$ region with arbitrary constraint edges is at most $2^{3k\ell}$ to obtain that

$$\mathcal{C}(\Omega') \leq 4\ell|\Lambda_{n,k}|^2 2^{6k\ell} \sum_{\xi,\xi',m''} \frac{\lambda^{|\xi|+|\xi'|+|m''|}}{Z_{\Omega'}} \leq 4\ell|\Lambda_{n,k}|^2 2^{6k\ell}.$$

Plugging everything into Lemma 3.5 completes the proof. □

4 Proof of Theorem 1.1

4.1 High-level overview

The proof is composed of three main ingredients: (i) a good ensemble, (ii) a decay of correlation analysis, and (iii) a recursion for the logarithmic Sobolev inequality.

The good ensemble. The first step is to show that uniformly over the initial condition, with high probability, for all times $t \in [T, T+n^2]$, with $T = O(n^2)$, the Markov chain stays within a subset $\widetilde{\Omega}$ of triangulations where all edges have length at most $C \log n$ for some constant $C > 0$. We will call this subset the *good ensemble*. This result will be a consequence of the tail estimate of Lemma 2.2. Therefore, we will couple our evolution in the time interval $t \in [T, T+n^2]$ with the Markov chain restricted to the good ensemble, which evolves as before, by attempting to flip edges chosen uniformly at random, but with the suppression of any edge flip that would render an edge longer than $C \log n$. The structural properties of triangulations imply that this Markov chain is irreducible. Moreover, the reversible probability measure is given by $\widetilde{\mu} = \mu(\cdot \mid \widetilde{\Omega})$, the measure μ conditioned on the event $\sigma \in \widetilde{\Omega}$. Since μ and $\widetilde{\mu}$ can be coupled with high probability, it is sufficient to analyze convergence to equilibrium for the restricted chain, and to show that the latter mixes in time $T' = O(n^2)$. We will actually prove that the restricted chain mixes in time $T' = n\mathrm{polylog}(n)$. For the rest of this discussion we assume that we are working with the Markov chain restricted to the good ensemble $\widetilde{\Omega}$.

Decay of correlations. We split the set of midpoints $\Lambda_{n,k}$ into two intersecting slabs Λ_ℓ and Λ_r, where Λ_ℓ contains all midpoints with horizontal coordinate smaller than $n/2 + 2C \log n$ and $\Lambda_{\rm r}$ contains all midpoints with horizontal coordinate at least $n/2 - 2C \log n$. Note that $\Lambda_\ell \cap \Lambda_{\rm r}$ is a slab of height k and horizontal width $4C \log n$. Let $\mathcal{F}_{\rm r}, \mathcal{F}_\ell$ be the σ-algebras generated by the edges with midpoints in $\Lambda_{\rm r} \setminus \Lambda_\ell, \Lambda_\ell \setminus \Lambda_{\rm r}$ respectively. We want to show that, conditional on any event $F \in \mathcal{F}_{\rm r}$, the distribution of the edges in $\Lambda_\ell \setminus \Lambda_{\rm r}$ is not affected much, and similarly for events $F \in \mathcal{F}_\ell$. The intuition for this is that the intersection $\Lambda_\ell \cap \Lambda_{\rm r}$ of the slabs is large enough to allow correlations from $\Lambda_\ell \setminus \Lambda_{\rm r}$ to decay. We will make this intuition rigorous by showing that there exists a positive $\epsilon = \epsilon(\lambda)$ such that, for all $\mathcal{F}_\ell$-measurable functions f_ℓ and all $\mathcal{F}_{\rm r}$-measurable functions $f_{\rm r}$, we have

$$\sup_{F \in \mathcal{F}_{\rm r}} \left| \widetilde{\mu}(f_\ell \mid F) - \widetilde{\mu}(f_\ell) \right| \le n^{-\epsilon} \|f_\ell\|_1 \quad \text{and} \quad \sup_{F \in \mathcal{F}_\ell} \left| \widetilde{\mu}(f_{\rm r} \mid F) - \widetilde{\mu}(f_{\rm r}) \right| \le n^{-\epsilon} \|f_{\rm r}\|_1, \tag{4.1}$$

where $\widetilde{\mu}(f \mid F)$ stands for the expectation of f given the event F and we use $\|f\|_1$ to denote the L^1 norm $\|f\|_1 = \sum_{\sigma \in \widetilde{\Omega}} \widetilde{\mu}(\sigma) |f(\sigma)|$.

The high-level argument for (4.1) is the following. Fix any valid collection of edges with midpoints in $\Lambda_\ell \setminus \Lambda_{\rm r}$, that is, a collection of edges that do not contain lattice points in their interior and that do not intersect the interior of other edges of the collection. Such a collection of edges can be viewed as a partial triangulation from $\widetilde{\Omega}$. This defines an event $F \in \mathcal{F}_\ell$. We will construct a coupling of one triangulation σ distributed according to $\widetilde{\mu}(\cdot \mid F)$ and another triangulation σ' distributed according to $\widetilde{\mu}(\cdot)$. We do this by first sampling the edges of σ' whose midpoint is in $\Lambda_\ell \setminus \Lambda_{\rm r}$. Call this event $F' \in \mathcal{F}_\ell$. Since we are restricted to the good ensemble, the edges of F and F' have length at most $C \log n$. Therefore, none of them crosses into the right half of $\Lambda_\ell \cap \Lambda_{\rm r}$. Lemma 2.5 therefore ensures that we may couple the sampling of edges in $\Lambda_{\rm r}$ so that, with probability at least $1 - e^{-\epsilon \log n}$, we put the same top-to-bottom crossing of unit verticals in σ and σ' inside the right half of $\Lambda_\ell \cap \Lambda_{\rm r}$. In particular, this implies that we can couple σ and σ' so that they agree on $\Lambda_{\rm r} \setminus \Lambda_\ell$. This will establish (4.1).

The log-Sobolev inequality. An important ingredient in the proof of Theorem 1.1 is the use of the logarithmic Sobolev inequality for the good ensemble. For any positive function f, let $\widetilde{\mu}(f)$ stand for the expectation of f in the good ensemble, and let

$$\mathrm{Ent}(f) = \widetilde{\mu}\left(f \log\left(\tfrac{f}{\widetilde{\mu}(f)}\right)\right) = \sum_\sigma \widetilde{\mu}(\sigma) f(\sigma) \log\left(\tfrac{f(\sigma)}{\widetilde{\mu}(f)}\right)$$

denote the entropy of f. Also, define

$$\mathcal{E}(f,f) = \frac{1}{2} \sum_{\sigma, \sigma' \in \widetilde{\Omega}} \widetilde{\mu}(\sigma) \rho(\sigma, \sigma') (f(\sigma) - f(\sigma'))^2,$$

where

$$\rho(\sigma, \sigma') = \frac{\lambda^{|\sigma'|}}{\lambda^{|\sigma|} + \lambda^{|\sigma'|}} \mathbf{1}\left(\sigma \sim \sigma'\right).$$

As usual $\sigma \sim \sigma'$ means that σ, σ' differ by a single edge flip. Note that $\rho(\sigma, \sigma') = |\Lambda_{n,k}| p(\sigma, \sigma')$, where p is the transition matrix of the discrete time chain. Thus $\mathcal{E}(f,f)$ can be interpreted as the Dirichlet form of the continuous time Markov chain where every edge of the triangulation independently attempts to flip at rate 1.

Let $c_{\rm S}$ be the *log-Sobolev constant* of this Markov chain, defined as the smallest constant $c > 0$ such that for all functions f one has

$$\mathrm{Ent}(f^2) \le c\, \mathcal{E}(f,f). \tag{4.2}$$

It is known (see e.g. [11, Theorem 2.9]) that $c_{\rm S}$ is related to the mixing and relaxation times via

$$\widetilde{T}_{\rm mix} \le \frac{c_{\rm S}}{4}\left(4 + \log_+ \log \widetilde{\mu}_*^{-1}\right) \quad \text{and} \quad 2\,\widetilde{T}_{\rm rel} \le c_{\rm S} \le \widetilde{T}_{\rm rel}\left(\frac{\log(\widetilde{\mu}_*^{-1})}{1 - 2\widetilde{\mu}_*}\right), \tag{4.3}$$

where $\widetilde{\mu}_* = \min_{\sigma\in\widetilde{\Omega}} \widetilde{\mu}(\sigma)$, and we use $\widetilde{T}_{\rm rel}, \widetilde{T}_{\rm mix}$ to denote the relaxation time and the mixing time of the continuous time chain restricted to the good set. These bounds should be compared with (3.1). In particular, it will be crucial for us to work with the log-Sobolev constant rather than the relaxation time in order to obtain the strong bound on mixing time claimed in Theorem 1.1.

Recursion. We will bound the (restricted) log-Sobolev constant via the so-called bisection method introduced in [10]. Let $\Lambda_\ell, \Lambda_{\rm r}$ and $\mathcal{F}_\ell, \mathcal{F}_{\rm r}$ be as above. Using the decay of correlations in (4.1), the decomposition estimate in [5, Proposition 2.1] implies that for all functions $f : \widetilde{\Omega} \mapsto \mathbb{R}$ we have

$$\mathrm{Ent}(f^2) \le \left(1 + O(n^{-\epsilon})\right) \widetilde{\mu}\left[\mathrm{Ent}(f^2 \mid \mathcal{F}_\ell) + \mathrm{Ent}(f^2 \mid \mathcal{F}_{\rm r})\right] \le \left(1 + O(n^{-\epsilon})\right) 2\, c_{\rm S}^{(1)} \mathcal{E}(f, f), \tag{4.4}$$

where $c_{\rm S}^{(1)}$ is the largest log-Sobolev constant among the systems conditioned on $\mathcal{F}_\ell$ and $\mathcal{F}_{\rm r}$ and the factor 2 comes from the double counting of flips within the region $\Lambda_\ell \cap \Lambda_{\rm r}$. Hence, we obtain that $c_{\rm S} \le (1 + O(n^{-\epsilon}))\, 2c_{\rm S}^{(1)}$. We would then like to recursively apply the same strategy to bound $\mathrm{Ent}(f^2 \mid \mathcal{F}_\ell)$ and $\mathrm{Ent}(f^2 \mid \mathcal{F}_{\rm r})$. Indeed, $\widetilde{\mu}(\cdot \mid \mathcal{F}_{\rm r})$ is a Gibbs measure on triangulations with midpoints in Λ_ℓ, and we may split Λ_ℓ into two intersecting slabs, establish decay of correlations and again use the decomposition above to further reduce the original scale. One caveat is that now we have to take into account the boundary conditions dictated by the conditioning on $\mathcal{F}_{\rm r}$. These consist of constraint edges protruding from the right boundary, with midpoints in $\Lambda_{\rm r} \setminus \Lambda_\ell$. The boundary conditions will not be a major problem since we are in the good ensemble so these edges cannot protrude more than a distance $C \log n$. After j such iterations, we will be considering slabs of size roughly $n2^{-j}$, with edges of size at most $C \log n$ protruding from both the left and right boundaries. It will be convenient to iterate this procedure for $j = j_*$ steps, where $n2^{-j_*}$ is roughly $\log^6 n$, so that protruding boundary edges are still far away from the middle of the slab, which is the crucial region for exploiting the decay of correlations. With this strategy, after j_* iterations we obtain

$$c_{\rm S} \le \left(1 + O(n^{-\epsilon})\right)^{j_*} 2^{j_*} c_{\rm S}^{(j_*)}.$$

Employing the general polynomial bound on the relaxation time of Theorem 3.3 and the relation between $c_{\rm S}$ and $T_{\rm rel}$, we obtain that $c_{\rm S}^{(j_*)}$ is at most $\mathrm{polylog}(n)$ uniformly over all boundary conditions in the good ensemble. The main problem is that the term 2^{j_*} is too large (of order $\frac{n}{\log^6 n}$ by our choice of j_*). As in [10] we overcome this difficulty by *randomizing* the location of the split of $\Lambda_{n,k}$ into Λ_ℓ and $\Lambda_{\rm r}$, and similarly for the other scales. The idea is to first split $\Lambda_{n,k}$ into three disjoint slabs with height k, the left and right slabs with horizontal length $\frac{1}{2}(n - \log^3 n)$, and the middle slab with horizontal length $\log^3 n$. Then we further split the middle slab into smaller slabs (that we call *rectangles*) each with horizontal length $4C \log n$. We choose one such rectangle uniformly at random, and define Λ_ℓ to be the midpoints to the left of this rectangle (including the rectangle) and $\Lambda_{\rm r}$ to be the midpoints to the right of this rectangle (including the rectangle). With this randomization, (4.4) will be improved to

$$\mathrm{Ent}(f^2) \le \left(1 + O(1/\log^2 n)\right) c_{\rm S}^{(1)} \mathcal{E}(f, f),$$

where $\log^2 n$ is roughly the number of rectangles in the middle slab of $\Lambda_{n,k}$. Then, iterating j_* times (with j_* as above) we get

$$c_{\rm S} \le \left(1 + O(j_*/\log^2 n)\right) c_{\rm S}^{(j_*)} = \mathrm{polylog}(n). \tag{4.5}$$

Once we obtain (4.5), using (4.3) we can conclude that the continuous time Markov chain restricted to the good ensemble satisfies $\widetilde{T}_{\rm mix} = \mathrm{polylog}(n)$. From this the desired conclusion for the discrete time Glauber dynamics will follow in a simple way.

We now proceed with the detailed proof of Theorem 1.1.

4.2 The good ensemble

Let $\sigma^0, \sigma^1, \ldots$ be the discrete time Markov chain on triangulations of $\Lambda^0_{n,k}$ with no constraint edges. The first step is to show that after a burn-in time of order n^2, during a very long time interval, the largest edge of the triangulation is of order at most $\log n$. Let $C = C(\lambda)$ be a large enough constant, and define

$$\widetilde{\Omega} = \Big\{ \sigma \in \Omega \colon |\sigma_x| \le C \log n \text{ for all } x \in \Lambda_{n,k} \Big\}. \tag{4.6}$$

The set $\widetilde{\Omega}$ represents the good ensemble. The next lemma will allow us to analyze the Markov chain restricted to the set $\widetilde{\Omega}$.

Lemma 4.1. *Fix $\lambda \in (0,1)$. There exists a constant $c_1 = c_1(\lambda)$ so that if we set $T = c_1 n^2$ then for all $n \ge k \ge 1$*

$$\mathbb{P}\left(\bigcap_{t=T}^{T+n^2} \left\{ \sigma^t \in \widetilde{\Omega} \right\} \right) \ge 1 - n^{-2}.$$

Proof. For any given $x \in \Lambda_{n,k}$ and any $t \ge c_1 n^2$, Lemma 2.2 gives that

$$\mathbb{P}\left(|\sigma^t_x| > C \log n \right) \le \exp(-c_2 C \log n),$$

for some constant c_2 independent of C and n. Setting C large enough and taking a union bound over all $x \in \Lambda_{n,k}$ and all integers $t \in [T, T + n^2]$ concludes the proof. □

4.3 Decay of correlations

Let $\Gamma \subset \Lambda$ be a *slab* of width w; that is, for some $x \in \mathbb{Z}$,

$$\Gamma = \Lambda_{n,k} \cap [x, x + w] \times [0, k].$$

We assume throughout that $w \ge \frac{1}{2} C^6 \log^6 n$, where C is fixed as in (4.6).

Partition Γ into three slabs, two of width roughly $\frac{1}{2}(w - C^3 \log^3 n)$ and one of width roughly $C^3 \log^3 n$. More precisely, for Γ as above, let

$$\Gamma_1 = \Lambda_{n,k} \cap \left[x, x + \tfrac{w - C^3 \log^3 n}{2}\right] \times [0,k], \quad \Gamma_2 = \Lambda_{n,k} \cap \left(x + \tfrac{w - C^3 \log^3 n}{2}, x + \tfrac{w + C^3 \log^3 n}{2}\right] \times [0,k]$$

$$\text{and} \quad \Gamma_3 = \Lambda_{n,k} \cap \left(x + \tfrac{w + C^3 \log^3 n}{2}, x + w\right] \times [0,k].$$

Partition the middle slab Γ_2 into disjoint slabs $J_1, J_2, \ldots, J_s$ (from left to right) each of width $4C \log n$, with

$$s = \frac{C^3 \log^3 n}{4C \log n} = \frac{C^2 \log^2 n}{4}. \tag{4.7}$$

Let ι be an integer chosen uniformly at random from $\{1, 2, \ldots, s\}$. Finally, define

$$\Gamma_\ell = \Gamma_1 \cup J_1 \cup J_2 \cup \cdots \cup J_\iota \quad \text{and} \quad \Gamma_{\mathrm{r}} = \Gamma_3 \cup J_\iota \cup J_{\iota+1} \cup \cdots \cup J_s. \tag{4.8}$$

Then, Γ_ℓ represents the left portion of Γ, Γ_{r} represents the right portion of Γ, and $\Gamma_\ell \cap \Gamma_{\mathrm{r}} = J_\iota$.

We need to introduce some more notation to be precise about boundary conditions. For any $\sigma \in \widetilde{\Omega}$, $A \subset \Lambda_{n,k}$, if $\sigma = \{\sigma_x,\ x \in \Lambda_{n,k}\}$ then we write σ_A for the set of edges $\{\sigma_x,\ x \in A\}$. If $\xi = \sigma_A$ for some $\sigma \in \widetilde{\Omega}$ and $A \subset \Lambda_{n,k}$ we say that σ *contains* ξ and we call ξ a *partial triangulation* in $\widetilde{\Omega}$. If $A \cap A' = \emptyset$ and $\xi = \sigma_A$, $\xi' = \sigma_{A'}$ for some $\sigma \in \widetilde{\Omega}$, then we define $\xi \cup \xi' = \sigma_{A \cup A'}$.

We use partial triangulations ξ in $\widetilde{\Omega}$ as boundary conditions for a region $B \subset \Gamma$. Fix a partial triangulation ξ. We denote by $A_\xi \subset \Lambda_{n,k}$ the set of midpoints of the edges in ξ.

Let $\widetilde{\Omega}^\xi$ denote the set of full triangulations $\sigma \in \widetilde{\Omega}$ that contain ξ. We define for any $B \subset \Gamma$, and any ξ such that $A_\xi \subset \Lambda_{n,k} \setminus B$,

$$\widetilde{\Omega}^\xi_B = \{\sigma_B : \ \sigma \in \widetilde{\Omega}^\xi\}.$$

For any $\eta_B \in \widetilde{\Omega}^\xi_B$, let

$$\mu^\xi_B(\eta_B) = \frac{\sum_{\sigma \in \widetilde{\Omega}^\xi : \ \sigma_B = \eta_B} \widetilde{\mu}(\sigma)}{\widetilde{\mu}(\widetilde{\Omega}^\xi)},$$

be the induced probability measure over $\widetilde{\Omega}^\xi_B$; recall that $\widetilde{\mu}(\sigma) = \frac{\mu(\sigma)}{\mu(\widetilde{\Omega})}$. In words, μ^ξ_B is the marginal distribution over midpoints B when we impose a boundary condition ξ. If ξ is empty (no boundary condition) we simply write $\widetilde{\Omega}_B$ and μ_B.

Lemma 4.2. *There exists a positive constant $c = c(\lambda, k)$ such that for any partial triangulation ξ with $A_\xi \subset \Lambda_{n,k} \setminus \Gamma$, for all functions $f_\ell, f_{\mathrm{r}} : \widetilde{\Omega} \mapsto \mathbb{R}$ such that f_ℓ depends only on edges with midpoint in $\Gamma_\ell \setminus J_\iota$ and f_{r} depends only on edges with midpoint in $\Gamma_{\mathrm{r}} \setminus J_\iota$, and for any $\sigma_\ell \in \widetilde{\Omega}^\xi_{\Gamma_\ell \setminus J_\iota}$ and $\sigma_{\mathrm{r}} \in \widetilde{\Omega}^\xi_{\Gamma_{\mathrm{r}} \setminus J_\iota}$, we have*

$$\left|\mu^{\xi \cup \sigma_{\mathrm{r}}}_{\Gamma_\ell \setminus J_\iota}(f_\ell) - \mu^\xi_{\Gamma_\ell \setminus J_\iota}(f_\ell)\right| \le \mu^\xi_{\Gamma_\ell \setminus J_\iota}(|f_\ell|) \exp(-c \log n)$$

and

$$\left|\mu^{\xi \cup \sigma_\ell}_{\Gamma_{\mathrm{r}} \setminus J_\iota}(f_{\mathrm{r}}) - \mu^\xi_{\Gamma_{\mathrm{r}} \setminus J_\iota}(f_{\mathrm{r}})\right| \le \mu^\xi_{\Gamma_{\mathrm{r}} \setminus J_\iota}(|f_{\mathrm{r}}|) \exp(-c \log n),$$

where we employ the notation $\mu(f)$ to denote the expected value of f with respect to the measure μ.

Proof. We will establish only the first estimate; the second follows by a symmetrical argument. Since f_ℓ depends only on edges with midpoint in $\Gamma_\ell \setminus J_\iota$, it is enough to show that, for any $\sigma_{\mathrm{r}} \in \widetilde{\Omega}^\xi_{\Gamma_{\mathrm{r}} \setminus J_\iota}$ and any $\tau \in \widetilde{\Omega}^{\xi \cup \sigma_{\mathrm{r}}}_{\Gamma_\ell \setminus J_\iota}$, we have

$$\left|\frac{\mu^{\xi \cup \sigma_{\mathrm{r}}}_{\Gamma_\ell \setminus J_\iota}(\tau)}{\mu^\xi_{\Gamma_\ell \setminus J_\iota}(\tau)} - 1\right| \le \exp(-c_2 C \log n), \tag{4.9}$$

for some positive $c_2 = c_2(\lambda, k)$, where C is the constant in the definition of the width of J_ι. Assuming (4.9), the proof is completed since

$$\begin{aligned}
\left|\mu^{\xi \cup \sigma_{\mathrm{r}}}_{\Gamma_\ell \setminus J_\iota}(f_\ell) - \mu^\xi_{\Gamma_\ell \setminus J_\iota}(f_\ell)\right| &= \left|\sum\nolimits_{\tau \in \widetilde{\Omega}^{\xi \cup \sigma_{\mathrm{r}}}_{\Gamma_\ell \setminus J_\iota}} \left(\mu^{\xi \cup \sigma_{\mathrm{r}}}_{\Gamma_\ell \setminus J_\iota}(\tau) - \mu^\xi_{\Gamma_\ell \setminus J_\iota}(\tau)\right) f_\ell(\tau)\right| \\
&\le \sum\nolimits_{\tau \in \widetilde{\Omega}^{\xi \cup \sigma_{\mathrm{r}}}_{\Gamma_\ell \setminus J_\iota}} \left|\mu^{\xi \cup \sigma_{\mathrm{r}}}_{\Gamma_\ell \setminus J_\iota}(\tau) - \mu^\xi_{\Gamma_\ell \setminus J_\iota}(\tau)\right| \left|f_\ell(\tau)\right| \\
&\le \sum\nolimits_{\tau \in \widetilde{\Omega}^{\xi \cup \sigma_{\mathrm{r}}}_{\Gamma_\ell \setminus J_\iota}} \mu^\xi_{\Gamma_\ell \setminus J_\iota}(\tau) \exp(-c_2 C \log n) \left|f_\ell(\tau)\right|.
\end{aligned}$$

Let η and η' be random triangulations distributed as $\mu^{\xi \cup \sigma_{\mathrm{r}}}_{\Gamma_\ell}$ and $\mu^\xi_{\Gamma_\ell}$, respectively. Let $\mathbb{P}$ denote the following coupling between η and η'; refer to Figure 2. The idea is to sample recursively edges from the pair (η, η') in vertical strips inside J_ι from right to left from a suitable coupling of $\mu^\xi_{J_\iota}$ and $\mu^{\xi \cup \sigma_{\mathrm{r}}}_{J_\iota}$. Here we will use the estimate of Lemma 2.5 to ensure that, with large probability, there is a common top-to-bottom crossing of unit verticals within J_ι. On this event we can safely resample $(\eta_{\Gamma_\ell \setminus J_\iota}, \eta'_{\Gamma_\ell \setminus J_\iota})$ in such a way that $\eta_{\Gamma_\ell \setminus J_\iota} = \eta'_{\Gamma_\ell \setminus J_\iota} = \tau$.

We now present the details. Consider the midpoints of Γ in order of their horizontal coordinate, from largest to smallest (i.e., from right to left in Figure 2). Let v_0 be the leftmost integer horizontal coordinate of points in $\Gamma_{\mathrm{r}} \setminus J_\iota$, and let $V_0 = \xi \cup \sigma_{\mathrm{r}}$ and $V'_0 = \xi$. Now for $i \ge 1$, define v_i, V_i, V'_i inductively as follows. Let $v_i < v_{i-1}$ be the rightmost

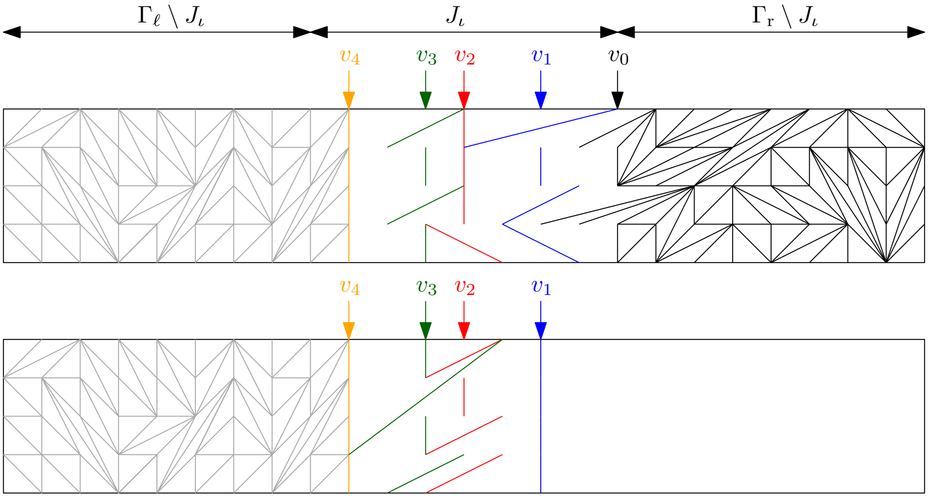

Figure 2: Coupling between $\mu^{\xi\cup\sigma_{\mathrm{r}}}_{\Gamma_\ell}$ (above) and $\mu^{\xi}_{\Gamma_\ell}$ (below). Note that the figure is not to scale: in reality, the middle region J_ι is much smaller than the two outer regions.

integer horizontal coordinate that is not crossed by an edge of $V_0 \cup V_1 \cup V'_1 \cup V_2 \cup V'_2 \cup \cdots \cup V_{i-1} \cup V'_{i-1}$. Using the coupling from Lemma 2.5, sample all edges of η and η' whose midpoints have horizontal coordinate v_i, and denote them by V_i and V'_i, respectively. There are two cases. In the first case, at least one edge of V_i or V'_i is not a unit vertical (as happens with $i = 1, 2$ and 3 in Figure 2). In this case, continue by defining v_{i+1} as described above. If v_{i+1} is a horizontal coordinate in J_ι, sample V_{i+1} and V'_{i+1} as described above and iterate. Otherwise, if v_{i+1} is not in J_ι, stop this procedure and sample the remaining edges of η and η' independently. In the second case, all edges in V_i and V'_i are unit verticals (i.e., they create a top-to-bottom crossing of Γ, as in Figure 2 for $i = 4$). Then stop the procedure above and sample the edges with horizontal coordinate smaller than v_i identically in both η and η' (as depicted by the gray edges in Figure 2), and then sample the remaining edges (that necessarily have midpoints in Γ_{r}) independently in η and η'. Let $I_{\eta,\eta'}$ be the event that η and η' have a common top-to-bottom crossing of unit verticals with midpoint in J_ι.

Let η_ℓ, η'_ℓ be the edges of η, η' with midpoints in $\Gamma_\ell \setminus J_\iota$, and let $\eta_{\mathrm{r}}, \eta'_{\mathrm{r}}$ be the edges of η, η' with midpoints in $\Gamma_{\mathrm{r}} \setminus J_\iota$. Using the above coupling, for any $\tau' \in \widetilde{\Omega}^{\xi}_{\Gamma_\ell\setminus J_\iota}$ we obtain

$$
\begin{aligned}
\mu^{\xi}_{\Gamma_\ell\setminus J_\iota}(\tau') = \mathbb{P}(\eta'_\ell = \tau') &= \sum\nolimits_{\tau\in\widetilde{\Omega}^{\xi\cup\sigma_{\mathrm{r}}}_{\Gamma_\ell\setminus J_\iota}} \mathbb{P}(\eta_\ell = \tau, \eta'_\ell = \tau') \\
&= \mathbb{P}(\eta_\ell = \tau', \eta'_\ell = \tau') + \sum\nolimits_{\tau\in\widetilde{\Omega}^{\xi\cup\sigma_{\mathrm{r}}}_{\Gamma_\ell\setminus J_\iota}\,:\,\tau\neq\tau'} \mathbb{P}(\eta_\ell = \tau, \eta'_\ell = \tau').
\end{aligned}
$$

The first term on the right-hand side above is at most $\mathbb{P}(\eta_\ell = \tau') = \mu^{\xi\cup\sigma_{\mathrm{r}}}_{\Gamma_\ell\setminus J_\iota}(\tau')$. The second term is bounded above by

$$
\mathbb{P}(\eta'_\ell = \tau')\mathbb{P}(\eta'_\ell \neq \eta_\ell \mid \eta'_\ell = \tau') \leq \mathbb{P}(\eta'_\ell = \tau')\mathbb{P}(I^{\mathrm{c}}_{\eta,\eta'} \mid \eta'_\ell = \tau') \leq \mathbb{P}(\eta'_\ell = \tau')\exp(-4cC\log n),
$$

where the first inequality follows since the coupling above gives that $I_{\eta,\eta'}$ implies $\eta'_\ell = \eta_\ell$, and the last step follows from Lemma 2.5. Plugging this into the equation above, and rearranging the terms, we obtain

$$
\mu^{\xi\cup\sigma_{\mathrm{r}}}_{\Gamma_\ell\setminus J_\iota}(\tau') \geq (1 - \exp(-4cC\log n))\,\mu^{\xi}_{\Gamma_\ell\setminus J_\iota}(\tau'),
$$

which holds uniformly over τ' and σ_{r}. Similarly, we write

$$\begin{aligned}\mu^{\xi\cup\sigma_{\mathrm{r}}}_{\Gamma_\ell\setminus J_\iota}(\tau) = \mathbb{P}(\eta_\ell=\tau) &\le \mathbb{P}(\eta'_\ell=\tau) + \mathbb{P}(\eta_\ell=\tau)\mathbb{P}(I^{\mathrm{c}}_{\eta,\eta'} \mid \eta_\ell=\tau)\\ &\le \mu^{\xi}_{\Gamma_\ell\setminus J_\iota}(\tau) + \mu^{\xi\cup\sigma_{\mathrm{r}}}_{\Gamma_\ell\setminus J_\iota}(\tau)\exp(-4cC\log n),\end{aligned}$$

and the proof of (4.9) is completed by rearranging the terms and setting c_2 appropriately. □

4.4 Recursion via bisection

We consider slabs of different scales: we index the scale by j, where $j=0$ corresponds to the full slab $\Lambda_{n,k}$ of width n, while at scale j, we have slabs of width w roughly equal to $n2^{-j}$. The finest scale will be

$$j_* = \min\big\{j\ge 0\colon n2^{-j}\le C^6\log^6 n\big\};$$

in particular, $n2^{-j_*}\ge \frac{1}{2}(C^6\log^6 n)$.

Recall how slabs are split and the definition of ι from the construction of Γ_ℓ and Γ_{r} in the paragraph culminating in (4.8). In this construction, a slab of scale 1 is defined by a choice of $\iota\in\{1,2,\dots,s\}$ and a choice between the left slab

$$\Lambda_{n,k}\cap\big[0,\tfrac{n-C^3\log^3 n}{2}+4\iota C\log n\big]\times[0,k]$$

and the right slab

$$\Lambda_{n,k}\cap\big[\tfrac{n-C^3\log^3 n}{2}+4(\iota-1)C\log n, n\big]\times[0,k].$$

Slabs of higher scale are obtained by inductively choosing a value of ι and one of the two sides for each scale. Hence, there are $(2s)^j$ possible slabs at scale j. Furthermore, setting $W_0=n$, we obtain that any slab of scale j must have width inside the interval

$$W_j=\big[n2^{-j}-jC^3\log^3 n, n2^{-j}+jC^3\log^3 n\big].$$

Consider a given scale $j\in\{0,\dots,j_*\}$, and let $\Gamma=\Gamma_j$ be any of the $(2s)^j$ possible slabs at scale j. Let $w\in W_j$ be the width of Γ. Let $\sigma\in\widetilde{\Omega}$ be an arbitrary triangulation in the good ensemble and set $\xi=\sigma_{\Lambda_{n,k}\setminus\Gamma}\in\widetilde{\Omega}_{\Lambda_{n,k}\setminus\Gamma}$ as a boundary condition for the region Γ. Consider the continuous time Markov chain on $\widetilde{\Omega}^{\xi}_{\Gamma}$ with Dirichlet form

$$\mathcal{E}^{\xi}_{\Gamma}(f,f)=\frac{1}{2}\sum_{\sigma_\Gamma,\sigma'_\Gamma\in\widetilde{\Omega}^{\xi}_{\Gamma}}\mu^{\xi}_{\Gamma}(\sigma_\Gamma)\rho^{\xi}_{\Gamma}(\sigma_\Gamma,\sigma'_\Gamma)(f(\sigma_\Gamma\cup\xi)-f(\sigma'_\Gamma\cup\xi))^2,$$

where $f:\widetilde{\Omega}\mapsto\mathbb{R}$ and

$$\rho^{\xi}_{\Gamma}(\sigma_\Gamma,\sigma'_\Gamma)=\frac{\lambda^{|\sigma'_\Gamma\cup\xi|}}{\lambda^{|\sigma_\Gamma\cup\xi|}+\lambda^{|\sigma'_\Gamma\cup\xi|}}\mathbf{1}\left(\sigma_\Gamma\cup\xi\sim\sigma'_\Gamma\cup\xi\right). \tag{4.10}$$

Let $c_{\mathrm{S}}(\Gamma,\xi)$ denote the log-Sobolev constant defined as the smallest constant $c>0$ such that

$$\mathrm{Ent}^{\xi}_{\Gamma}(f^2)\le c\,\mathcal{E}^{\xi}_{\Gamma}(f,f), \tag{4.11}$$

holds for all functions f, where $\mathrm{Ent}^{\xi}_{\Gamma}(f^2)$ denotes the entropy of f^2 with respect to μ^{ξ}_{Γ}.

Finally we define, for each j,

$$\gamma_j=\sup\big\{c_{\mathrm{S}}(\Gamma,\xi)\colon\Gamma\subset\Lambda_{n,k}\text{ is a slab of width }w\in W_j\text{, and }\xi\in\widetilde{\Omega}_{\Lambda_{n,k}\setminus\Gamma}\big\}.$$

The following lemma summarizes the result of this recursion.

Lemma 4.3. *There exists a positive constant c_2 such that, for any integer $j \in \{0, \dots, j_* - 1\}$,*

$$\gamma_j \leq \left(1 + e^{-c_2 \log n}\right) \left(1 + \tfrac{4}{C^2 \log^2 n}\right) \gamma_{j+1}.$$

Proof. Let Γ be a fixed slab of width $w \in W_j$, and let ξ be a given boundary condition. Let s, ι, Γ_ℓ and Γ_{r} be as described in the paragraph culminating in (4.8). Using Lemma 4.2, the condition in [5, Proposition 2.1] is satisfied with $\epsilon = \exp(-c \log n)$. Then for any function $f \colon \widetilde{\Omega} \mapsto \mathbb{R}$ and for any choice of ι, which defines the way Γ is partitioned, [5, Proposition 2.1] gives

$$\left(1 - \frac{84\epsilon}{(1-\epsilon)^2}\right) \mathrm{Ent}^\xi_\Gamma(f^2) \leq \left(\sum_{\sigma_{\mathrm{r}} \in \widetilde{\Omega}^\xi_{\Gamma_{\mathrm{r}} \setminus J_\iota}} \mu^\xi_{\Gamma_{\mathrm{r}} \setminus J_\iota}(\sigma_{\mathrm{r}}) \mathrm{Ent}^{\xi \cup \sigma_{\mathrm{r}}}_{\Gamma_\ell}(f^2) + \sum_{\sigma_\ell \in \widetilde{\Omega}^\xi_{\Gamma_\ell \setminus J_\iota}} \mu^\xi_{\Gamma_\ell \setminus J_\iota}(\sigma_\ell) \mathrm{Ent}^{\xi \cup \sigma_\ell}_{\Gamma_{\mathrm{r}}}(f^2) \right).$$

Taking expectation over the choice of ι, and rearranging the terms, we obtain a positive constant c_2 such that

$$\mathrm{Ent}^\xi_\Gamma(f^2) \leq \frac{1}{s} \sum_{\iota=1}^{s} \left(1 + e^{-c_2 \log n}\right) \left(\sum_{\sigma_{\mathrm{r}} \in \widetilde{\Omega}^\xi_{\Gamma_{\mathrm{r}} \setminus J_\iota}} \mu^\xi_{\Gamma_{\mathrm{r}} \setminus J_\iota}(\sigma_{\mathrm{r}}) \mathrm{Ent}^{\xi \cup \sigma_{\mathrm{r}}}_{\Gamma_\ell}(f^2) + \sum_{\sigma_\ell \in \widetilde{\Omega}^\xi_{\Gamma_\ell \setminus J_\iota}} \mu^\xi_{\Gamma_\ell \setminus J_\iota}(\sigma_\ell) \mathrm{Ent}^{\xi \cup \sigma_\ell}_{\Gamma_{\mathrm{r}}}(f^2) \right). \tag{4.12}$$

Note that $\mathrm{Ent}^{\xi \cup \sigma_{\mathrm{r}}}_{\Gamma_\ell}(f^2)$ and $\mathrm{Ent}^{\xi \cup \sigma_\ell}_{\Gamma_{\mathrm{r}}}(f^2)$ are entropy functions for slabs on scale $j+1$ given boundary conditions $\xi \cup \sigma_{\mathrm{r}}$ and $\xi \cup \sigma_\ell$, respectively. Therefore, by (4.11) we have

$$\mathrm{Ent}^{\xi \cup \sigma_{\mathrm{r}}}_{\Gamma_\ell}(f^2) \leq c_{\mathrm{S}}(\Gamma_\ell, \xi \cup \sigma_{\mathrm{r}}) \mathcal{E}^{\xi \cup \sigma_{\mathrm{r}}}_{\Gamma_\ell}(f, f) \leq \gamma_{j+1} \mathcal{E}^{\xi \cup \sigma_{\mathrm{r}}}_{\Gamma_\ell}(f, f), \tag{4.13}$$

and similarly for the second term in (4.12). Now we claim that

$$\sum_{\iota=1}^{s} \left(\sum_{\sigma_{\mathrm{r}} \in \widetilde{\Omega}^\xi_{\Gamma_{\mathrm{r}} \setminus J_\iota}} \mu^\xi_{\Gamma_{\mathrm{r}} \setminus J_\iota}(\sigma_{\mathrm{r}}) \mathcal{E}^{\xi \cup \sigma_{\mathrm{r}}}_{\Gamma_\ell}(f, f) + \sum_{\sigma_\ell \in \widetilde{\Omega}^\xi_{\Gamma_\ell \setminus J_\iota}} \mu^\xi_{\Gamma_\ell \setminus J_\iota}(\sigma_\ell) \mathcal{E}^{\xi \cup \sigma_\ell}_{\Gamma_{\mathrm{r}}}(f, f) \right) \leq (1+s) \mathcal{E}^\xi_\Gamma(f, f). \tag{4.14}$$

To prove (4.14) we proceed as follows. Since a given edge σ_x in a triangulation has at most one value $\sigma'_x \neq \sigma_x$ it can flip to, we may write the flip rates (4.10) as

$$\rho^\xi_\Gamma(\sigma_\Gamma, \sigma'_\Gamma) = \sum_{x \in \Gamma} \frac{\lambda^{|\sigma'_x|}}{\lambda^{|\sigma_x|} + \lambda^{|\sigma'_x|}} \mathbf{1}\left(\sigma_\Gamma \cup \xi \sim \sigma'_\Gamma \cup \xi;\ \sigma_x \neq \sigma'_x\right) =: \sum_{x \in \Gamma} \rho^\xi_{x,\Gamma}(\sigma_\Gamma).$$

Therefore,

$$\mathcal{E}^\xi_\Gamma(f, f) = \frac{1}{2} \sum_x \sum_{\sigma_\Gamma \in \widetilde{\Omega}^\xi_\Gamma} \mu^\xi_\Gamma(\sigma_\Gamma) \rho^\xi_{x,\Gamma}(\sigma_\Gamma) (\nabla_x f(\sigma_\Gamma \cup \xi))^2, \tag{4.15}$$

where we use $\nabla_x f$ to denote the difference in values of f before and after the flip at x. It follows that

$$\sum_{\sigma_{\mathrm{r}} \in \widetilde{\Omega}^\xi_{\Gamma_{\mathrm{r}} \setminus J_\iota}} \mu^\xi_{\Gamma_{\mathrm{r}} \setminus J_\iota}(\sigma_{\mathrm{r}}) \mathcal{E}^{\xi \cup \sigma_{\mathrm{r}}}_{\Gamma_\ell}(f, f) = \frac{1}{2} \sum_{x \in \Gamma_\ell} \sum_{\sigma_{\mathrm{r}} \in \widetilde{\Omega}^\xi_{\Gamma_{\mathrm{r}} \setminus J_\iota}} \mu^\xi_{\Gamma_{\mathrm{r}} \setminus J_\iota}(\sigma_{\mathrm{r}}) \sum_{\eta_{\Gamma_\ell} \in \widetilde{\Omega}^{\xi \cup \sigma_{\mathrm{r}}}_{\Gamma_\ell}} \mu^{\xi \cup \sigma_{\mathrm{r}}}_{\Gamma_\ell}(\eta_{\Gamma_\ell}) \rho^{\xi \cup \sigma_{\mathrm{r}}}_{x,\Gamma_\ell}(\eta_{\Gamma_\ell}) (\nabla_x f(\sigma_{\eta_{\Gamma_\ell}} \cup \xi \cup \sigma_{\mathrm{r}}))^2,$$

where, as before, we use the shortcut notation $\sigma_{\rm r} = \sigma_{\Gamma_{\rm r}\setminus J_\iota}$. Using

$$\mu^{\xi}_{\Gamma_{\rm r}\setminus J_\iota}(\sigma_{\rm r})\mu^{\xi\cup\sigma_{\rm r}}_{\Gamma_\ell}(\eta_{\Gamma_\ell})\rho^{\xi\cup\sigma_{\rm r}}_{x,\Gamma_\ell}(\eta_{\Gamma_\ell}) = \mu^{\xi}_{\Gamma}(\eta_{\Gamma_\ell}\cup\sigma_{\rm r})\rho^{\xi}_{x,\Gamma}(\eta_{\Gamma_\ell}\cup\sigma_{\rm r})$$

and rearranging the sum, we obtain

$$\sum_{\sigma_{\rm r}\in\widetilde\Omega^{\xi}_{\Gamma_{\rm r}\setminus J_\iota}} \mu^{\xi}_{\Gamma_{\rm r}\setminus J_\iota}(\sigma_{\rm r})\mathcal{E}^{\xi\cup\sigma_{\rm r}}_{\Gamma_\ell}(f,f) = \frac{1}{2}\sum_{x\in\Gamma_\ell}\sum_{\sigma_\Gamma\in\widetilde\Omega^{\xi}_{\Gamma}} \mu^{\xi}_{\Gamma}(\sigma_\Gamma)\rho^{\xi}_{x,\Gamma}(\sigma_\Gamma)(\nabla_x f(\sigma_\Gamma\cup\xi))^2.$$

A similar expression holds for the second term on the left-hand side of (4.14), and the desired estimate follows from the expression (4.15).

Plugging (4.14) and (4.13) into the bound in (4.12) we have

$$\mathrm{Ent}^{\xi}_{\Gamma}(f^2) \le \left(1+e^{-c_2\log n}\right)\gamma_{j+1}\left(1+\tfrac{1}{s}\right)\mathcal{E}^{\xi}_{\Gamma}(f,f).$$

This establishes that $c_{\rm S}(\Gamma,\xi) \le \left(1+e^{-c_2\log n}\right)\gamma_{j+1}\left(1+\tfrac{1}{s}\right)$. Since this bound does not depend on ξ and the choice of slab Γ at scale j, the proof is completed by using the value of s from (4.7). □

We conclude the proof with the base of the induction.

Lemma 4.4. *There exists a constant $c = c(\lambda,k)$ such that*

$$\gamma_{j_*} \le \log^c n.$$

Proof. Let Γ be a slab at scale j_*, so that the width of Γ is of order $\log^6 n$. Let $\xi\in\widetilde\Omega_{\Lambda_{n,k}\setminus\Gamma}$ be a boundary condition. We note that the argument of Theorem 3.3 can be repeated with no modifications for the chain restricted to the good set $\widetilde\Omega$. Therefore, there exists a constant $c_1 = c_1(\lambda,k)$ independent of Γ and ξ such that the relaxation time of the discrete time chain on Γ with boundary condition ξ is at most $\log^{c_1} n$. Passing to continuous time, we have that $\widetilde T_{\rm rel}(\Gamma,\xi)\le\log^{c_1} n$. Since triangulations in $\widetilde\Omega^{\xi}_{\Gamma}$ have edges of length at most $C\log n$, there exists a constant c_2 such that

$$\min_{\sigma_\Gamma\in\widetilde\Omega^{\xi}_{\Gamma}}\mu^{\xi}_{\Gamma}(\sigma_\Gamma)\ge n^{-c_2},$$

uniformly over all slabs Γ at scale j_* and boundary conditions ξ. Therefore, using the relation between the relaxation time and the log-Sobolev constant from (4.2) we have that

$$c_{\rm S}(\Gamma,\xi)\le\widetilde T_{\rm rel}(\Gamma,\xi)\left(\frac{\log(n^{c_2})}{1/2}\right).$$

Since the bound above is uniform Γ and ξ, this proves the desired bound on γ_{j_*}. □

4.5 Completing the proof

Proof of Theorem 1.1. We start by bounding the mixing time of the discrete time Markov chain on $\widetilde\Omega$. Lemma 4.3 implies that the log-Sobolev constant of the continuous time Markov chain on $\Lambda_{n,k}$ with no boundary condition is at most

$$c_{\rm S}(\Lambda_{n,k})\le\gamma_0\le\left(1+e^{-c_2\log n}\right)^{j_*-1}\left(1+\tfrac{4}{C^2\log^2 n}\right)^{j_*-1}\gamma_{j_*}\le 2\gamma_{j_*},$$

where the last step follows since $j_*\le\log_2 n$. Also, we have that

$$\min_{\sigma\in\widetilde\Omega}\mu(\sigma)\ge\frac{\lambda^{|\Lambda_{n,k}|C\log n}}{(2\lambda)^{|\Lambda_{n,k}|}},$$

where $(2\lambda)^{|\Lambda_{n,k}|}$ comes from Anclin's bound of $2^{|\Lambda_{n,k}|}$ for the number of lattice triangulations [1], and the fact that the total edge length of any triangulation is at least $|\Lambda_{n,k}|$. Therefore, using the relation between the mixing time and log-Sobolev constant in (4.2), we deduce that the mixing time $\widetilde{T}_{\text{mix}}$ of the continuous time Markov chain on $\widetilde{\Omega}$ is bounded above by $c\gamma_{j_*}\log n$. Thus, the mixing time of the discrete chain in $\widetilde{\Omega}$ is at most $|\Lambda_{n,k}|c\gamma_{j_*}\log n$, for some constant c. Using Lemma 4.4 and the fact that $|\Lambda_{n,k}|$ is of order nk, we obtain that the mixing time of the Markov chain restricted to $\widetilde{\Omega}$ is at most $cn\log^c n$, for some new positive constant c (which depends on k and λ).

Now we compare the restricted chain on $\widetilde{\Omega}$ to the original unrestricted chain on $\Omega = \Omega(n,k)$. Let $T_1 = cn\log^c n$ and fix the constant $c > 0$ so that the total variation distance between the restricted chain at time T_1 and the restricted stationary distribution $\widetilde{\mu}$ is at most $1/8$. We obtain the mixing time of the unrestricted chain via the following coupling. Let $T_0 = c_1 n^2$, where c_1 is the constant in Lemma 4.1. Let the unrestricted Markov chain run for $T_0 + T_1$ steps. With probability at least $1 - n^{-2}$, the unrestricted chain never leaves the set $\widetilde{\Omega}$ during the time interval $[T_0, T_0 + T_1]$; therefore, we can couple its steps with those of the restricted chain. This gives that the total variation distance between the unrestricted chain at time $T_0 + T_1$ and the stationary distribution is at most $n^{-2} + 1/8 + \mu(\Omega \setminus \widetilde{\Omega})$. Since $\Omega \setminus \widetilde{\Omega}$ only contains triangulations for which the largest edge is larger than $C\log n$, Lemma 2.2 ensures that $\mu(\Omega \setminus \widetilde{\Omega}) \le n^{-2}$ for large enough C, and therefore the total variation distance between the unrestricted chain at time $T_0 + T_1$ and its stationary distribution is at most $1/4$. This completes the proof of Theorem 1.1. □

References

[1] Emile E. Anclin. An upper bound for the number of planar lattice triangulations. *Journal of Combinatorial Theory, Series A*, 103(2):383–386, August 2003. MR-1996075

[2] Pietro Caputo, Eyal Lubetzky, Fabio Martinelli, Allan Sly, and Fabio Lucio Toninelli. Dynamics of $(2+1)$-dimensional SOS surfaces above a wall: Slow mixing induced by entropic repulsion. *Annals of Probability*, 42(4):1516–1589, 2014. MR-3262485

[3] Pietro Caputo, Fabio Martinelli, Alistair Sinclair, and Alexandre Stauffer. Random lattice triangulations: Structure and algorithms. *Annals of Applied Probability*, 25(3):1650–1685, 2015. Preliminary version appeared in *Proceedings of the 2013 ACM Symposium on Theory of Computing (STOC)*. MR-3325284

[4] Pietro Caputo, Fabio Martinelli, and Fabio Lucio Toninelli. On the approach to equilibrium for a polymer with adsorption and repulsion. *Electronic Journal of Probability*, 13(10):213–258, 2008. MR-2386733

[5] Filippo Cesi. Quasi-factorization of the entropy and logarithmic sobolev inequalities for gibbs random fields. *Probability Theory and Related Fields*, 120:569–584, 2001. MR-1853483

[6] Jesús A. De Loera, Jörg Rambau, and Francisco Santos. *Triangulations*, volume 25 of *Algorithms and Computation in Mathematics*. Springer-Verlag, Berlin, 2010. MR-2743368

[7] Sam Greenberg, Amanda Pascoe, and Dana Randall. Sampling biased lattice configurations using exponential metrics. In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 76–85. SIAM, Philadelphia, PA, 2009. MR-2809307

[8] Volker Kaibel and Günter M. Ziegler. Counting lattice triangulations. In *Surveys in Combinatorics*, volume 307 of *London Mathematical Society Lecture Note Series*, pages 277–307. Cambridge Univ. Press, Cambridge, 2003. MR-2011739

[9] David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov Chains and Mixing Times.* American Mathematical Society, Providence, RI, 2009. MR-2466937

[10] Fabio Martinelli. Lectures on Glauber dynamics for discrete spin models. In *Lectures on Probability Theory and Statistics*, pages 93–191. Springer-Verlag, Berlin, Heidelberg, 2004. MR-1746301

[11] Fabio Martinelli. Relaxation times of markov chains in statistical mechanics and combinatorial structures. In H. Kesten, editor, *Probability on Discrete Structures*. Springer-Verlag, Heidelberg, 2004. MR-2023653

[12] Alistair Sinclair. Improved bounds for mixing rates of Markov chains and multicommodity flow. *Combinatorics, Probability and Computing*, 1(4):351–370, 1992. MR-1211324

[13] Alexandre Stauffer. A Lyapunov function for Glauber dynamics on lattice triangulations, 2015. Preprint at arXiv:1504.07980 [math.PR].