

TIGHTNESS AND TAILS OF THE MAXIMUM IN 3D ISING INTERFACES

BY REZA GHEISSARI* AND EYAL LUBETZKY†

*Courant Institute of Mathematical Sciences, New York University, *reza@cims.nyu.edu; †eyal@courant.nyu.edu*

Consider the 3D Ising model on a box of side length n with minus boundary conditions above the xy -plane and plus boundary conditions below it. At low temperatures, Dobrushin (1972) showed that the interface separating the predominantly plus and predominantly minus regions is localized: its height above a fixed point has exponential tails. Recently, the authors proved a law of large numbers for the maximum height M_n of this interface: for every β large, $M_n / \log n \rightarrow c_\beta$ in probability as $n \rightarrow \infty$.

Here, we show that the laws of the centered maxima $(M_n - \mathbb{E}[M_n])_{n \geq 1}$ are uniformly tight. Moreover, even though this sequence does not converge, we prove that it has uniform upper and lower Gumbel tails (exponential right tails and doubly exponential left tails). Key to the proof is a sharp (up to $O(1)$ precision) understanding of the surface large deviations. This includes, in particular, the shape of a pillar that reaches near-maximum height, even at its base, where the interactions with neighboring pillars are dominant.

CONTENTS

1. Introduction	733
1.1. Proof strategy and outline	736
1.1.1. Proof of Theorem 4	737
1.1.2. Proof of tightness and Gumbel tails for $M_n - \mathbb{E}[M_n]$	738
1.2. Organization	738
2. Preliminaries	738
2.1. Graph notation	739
2.2. The Ising model	739
Domain Markov property	740
FKG inequality	740
Infinite-volume Gibbs measures and DLR condition	740
2.3. The Ising interface with Dobrushin boundary conditions	741
2.4. Walls, ceilings and groups of walls	742
3. Decomposition of tall pillars	745
3.1. Pillars and increments	745
3.2. Preliminary estimates on tall pillars	746
3.3. Tame pillars	747
4. Sharp estimates on the structure of tall pillars	750
4.1. A new base and increment map	751
4.2. Strategy of the map $\Psi_{x,t}$	751
4.3. Properties of $\Psi_{x,t}$	754
4.3.1. Decomposition of the interfaces	755
4.3.2. The excess area of the map	756
4.4. Proof part 1: Interface weights	759
4.5. Proof part 2: Multiplicity	766
Construction of the witness	766
Reconstructing \mathcal{I} from the witness	768
Enumerating over possible witnesses	768
4.6. Proofs of Theorem 4.1 and Proposition 4.2	771

Received September 2019; revised April 2020.

MSC2020 subject classifications. 60K35, 82B20, 82B24.

Key words and phrases. 3D Ising model, low temperature interface, maximum of random surface.

5. A refined submultiplicativity bound	773
6. Tightness and exponential tails of the maximum	779
6.1. Exponential tails for the maximum	781
6.2. The expectation and median of the maximum	782
7. Gumbel tail estimates for the maximum	783
7.1. Coupling of different scales	783
7.2. From multiscale coupling to Gumbel tails	786
7.2.1. Left tail	786
7.2.2. Right tail	788
7.3. Nonconvergence of the centered maximum	789
Acknowledgments	791
References	791

1. Introduction. We study the low-temperature interface of the Ising model in dimensions three and higher. An Ising configuration on a subgraph Λ of \mathbb{Z}^d is an assignment of $\{+1, -1\}$ spins to the (d -dimensional) cells of Λ , denoted $\mathcal{C}(\Lambda)$. The cells are identified with their midpoints, which are the vertices of the dual graph $(\mathbb{Z} + \frac{1}{2})^d$ (when $d = 2$, the cells are the faces of Λ and, when $d = 3$, they are cubes of side length 1), and two cells u, v are considered adjacent ($v \sim w$) if their mid-points are at a Euclidean distance 1. The Ising model on a finite $\Lambda \subset \mathbb{Z}^d$ at inverse-temperature β is the probability measure over configurations $\sigma \in \{\pm 1\}^{\mathcal{C}(\Lambda)}$ with weights

$$\mu_{\Lambda, \beta} \propto \exp[-\beta \mathcal{H}(\sigma)] \quad \text{where } \mathcal{H}(\sigma) = \sum_{v \sim w} \mathbf{1}\{\sigma_v \neq \sigma_w\}.$$

For a set $R \supset \Lambda$ and a configuration η on R , the Ising model on Λ with boundary conditions η , denoted $\mu_{\Lambda, \beta}^\eta$, is the measure $\mu_{R, \beta}$, conditioned on σ coinciding with η on $R \setminus \Lambda$. These definitions extend to infinite subsets $\Lambda \subset \mathbb{Z}^d$ by taking weak limits of measures on (say) finite boxes under suitable boundary conditions.

The Ising model exhibits a rich and extensively studied phase transition on \mathbb{Z}^d ($d \geq 2$): there exists a $\beta_c(d)$ such that, when $\beta < \beta_c(d)$, there is a unique infinite-volume Gibbs measure on \mathbb{Z}^d , whereas when $\beta > \beta_c(d)$, there are multiple distinct infinite-volume measures, for example, those obtained by taking a limit of finite volume measures with plus boundary conditions $\mu_{\mathbb{Z}^d, \beta}^+$ vs. minus boundary conditions $\mu_{\mathbb{Z}^d, \beta}^-$. This can be seen by a classical argument of Peierls, which demonstrates that in a box of side length n with $+$ boundary conditions, there exists a β_0 such that for $\beta > \beta_0$ the minus clusters become “subcritical,” that is, the probability that the origin is part of a connected component (cluster) of at least r minus sites is at most $\exp[-c\beta r]$.

The analysis of this low-temperature phase has then focused on the structure of the interface between predominantly plus and predominantly minus regions. Namely, consider the model on the infinite cylinder

$$\Lambda_n := \{-n, \dots, n\}^{d-1} \times \{-\infty, \dots, \infty\}$$

under *Dobrushin boundary conditions*, which are plus on all cells of \mathbb{Z}^d with negative d th coordinates and minus on all cells of \mathbb{Z}^d with positive d th coordinates: denote the resulting (infinite volume) measure μ_n^\mp (the uniqueness of which, for every $\beta > 0$, follows by a classical coupling argument via the monotonicity of the Ising model in boundary conditions). At low temperatures $\beta > \beta_0$, these boundary conditions impose an *interface* \mathcal{I} , below which the configuration has the features of the plus phase $\mu_{\mathbb{Z}^d}^+$ and above which the measure has the features of the minus phase $\mu_{\mathbb{Z}^d}^-$. For a configuration $\sigma \in \{\pm 1\}^{\mathcal{C}(\Lambda_n)}$, this interface \mathcal{I} is defined by taking the set of all $(d - 1)$ -cells (e.g., edges for $d = 2$ and faces for $d = 3$), separating

disagreeing spins and letting \mathcal{I} be the (maximal) connected component of such $(d - 1)$ -cells that separates the minus cluster of the boundary from the plus cluster of the boundary (see Section 2 for a precise definition).

In two dimensions this interface forms a random decorated curve whose properties as $n \rightarrow \infty$ are, by now, very well understood. For every $\beta > \beta_c(2)$, this interface is *rough*, with typical fluctuations that are $O(\sqrt{n})$. In fact, after a diffusive rescaling the interface is known to converge to a Brownian bridge, and as such its *maximum* height in Λ_n can be seen to also have fluctuations of order \sqrt{n} [11, 12, 18–20, 23, 24].

In dimensions $d \geq 3$, the interface forms a random $(d - 1)$ -dimensional surface whose features are quite different from the two-dimensional case described above. For the ease of exposition, we focus on the important case $d = 3$, where the interface forms a random 2D surface; unless otherwise noted, the new results stated for $d = 3$ extend to $d > 3$ with simple modifications (see Remark 1.2). A landmark result in the study of the 3D Ising interface in μ_n^\mp was the proof by Dobrushin [14] in 1972 that for large β the interface is *rigid*: its typical height fluctuations are $O(1)$ and, in fact, have an exponential tail, for example,

$$\mu_n^\mp(\max\{h : (0, 0, h) \in \mathcal{I}\} \geq r) \leq \exp[-\beta r/3].$$

A consequence of this is the existence of an infinite-volume measure $\mu_{\mathbb{Z}^3}^\mp$ which is not a mixture of $\mu_{\mathbb{Z}^3}^+$ and $\mu_{\mathbb{Z}^3}^-$ (cf., dimension two where all infinite-volume measures are mixtures of $\mu_{\mathbb{Z}^2}^+$ and $\mu_{\mathbb{Z}^2}^-$ [1, 17]).

The present paper studies the maximum height of the 3D Ising interface. Unlike several related models of 2D surfaces whose maximum has been extensively analyzed in recent years, for example, the $(2 + 1)$ D solid-on-solid (SOS) model [5, 7, 8], the discrete Gaussian and $|\nabla\phi|^p$ models [22] and the discrete Gaussian free field (DGFF) [2, 3, 25], to name a few (we refer the reader to [16], Section 1.4, for a more detailed review of this literature), the 3D Ising interface \mathcal{I} is not a height function; it can have *overhangs* intersecting a given column $(x_1, x_2) \times \mathbb{R}$ at multiple heights. Further, subcritical *bubbles* in the plus and minus phases under μ_n^\mp , albeit unseen in \mathcal{I} , do affect its distribution (bubbles and overhangs are precluded from SOS, for instance).

Since the Ising interface is not a height function, we define its maximum height as

$$M_n := \max\{x_3 : (x_1, x_2, x_3) \in \mathcal{I} \text{ for some } (x_1, x_2) \in [-n, n]^2\}.$$

The above bound by Dobrushin on the bulk fluctuations of \mathcal{I} implies, via a union bound, that $M_n \leq (C/\beta) \log n$ with probability $1 - o(1)$. Recently, the authors showed a law of large numbers for this maximum height [16]: there exists β_0 such that every $\beta > \beta_0$,

$$\lim_{n \rightarrow \infty} \frac{M_n}{\log n} = \frac{2}{\alpha} \quad \text{in } \mu_n^\mp\text{-probability,}$$

where, if $v \xleftrightarrow[A]{+} w$ denotes that there is a path of adjacent or diagonally adjacent plus spins between v and w in A (henceforth, we refer to this notion of adjacency as $*$ -connectivity; see Section 2.1 for more details), then

$$(1.1) \quad \alpha := \lim_{h \rightarrow \infty} \frac{\alpha_h}{h} \quad \text{for } \alpha_h = -\log \mu_{\mathbb{Z}^3}^\mp \left(\left(\frac{1}{2}, \frac{1}{2}, \frac{1}{2} \right) \xleftrightarrow[\mathbb{R}^2 \times [0, \infty)]{+} \left(\mathbb{Z} + \frac{1}{2} \right)^2 \times \left\{ h - \frac{1}{2} \right\} \right).$$

The present work continues the analysis of the extrema of the random surface \mathcal{I} given by the interface of low-temperature Ising models in three dimensions and looks at the law of M_n beyond first order asymptotics. We begin by characterizing the mean of M_n in terms of the infinite-volume large deviation quantity α_h and proving the tightness of the centered maximum.

THEOREM 1. *There exist $\beta_0 > 0$ and a sequence $\varepsilon_\beta > 0$ going to 0 as $\beta \rightarrow \infty$, so that, for every $\beta > \beta_0$, the maximum M_n of the interface of the Ising model on Λ_n with Dobrushin boundary conditions satisfies*

$$(1.2) \quad m_n^* - 1 - \varepsilon_\beta \leq \mathbb{E}[M_n] \leq m_n^* + \varepsilon_\beta \quad \text{where } m_n^* = \inf\{h \geq 1 : \alpha_h > 2 \log(2n) - 2\beta\}$$

and

$$\mu_n^\mp(M_n - \mathbb{E}[M_n] \in \cdot) \text{ is a tight sequence.}$$

REMARK 1.1. We also show (cf. Corollary 6.6) that every median m_n of $\mu_n^\mp(M_n \in \cdot)$ has $m_n^* - 1 \leq m_n \leq m_n^*$.

REMARK 1.2. The results extend to the d -dimensional Ising model for any $d \geq 3$, where m_n^* in (1.2) becomes $\inf\{h \geq 1 : \alpha_h > (d - 1) \log(2n) - (d - 1)\beta\}$ and α_h addresses $(\frac{1}{2}, \dots, \frac{1}{2}) \xleftrightarrow[\mathbb{R}^{d-1} \times [0, \infty)]{+} (\mathbb{Z} + \frac{1}{2})^{d-1} \times \{h - \frac{1}{2}\}$.

Having established the tightness of the maximum height of the 3D Ising interface around its mean, one can then ask about the behavior of the limit/subsequential limits of the centered maximum. Using a multiscale argument, we prove the following regarding the limit points of $M_n - \mathbb{E}[M_n]$.

THEOREM 2. *There exists β_0 such that, for every $\beta > \beta_0$, there are $C, \bar{\alpha} > 0$ so that the maximum M_n of the interface of the Ising model on Λ_n with Dobrushin boundary conditions has, for all $r > 0$ and large enough n ,*

$$\begin{aligned} e^{-(\bar{\alpha}r+C)} &\leq \mu_n^\mp(M_n \geq \mathbb{E}[M_n] + r) \leq e^{-(\alpha_r-C)}, \\ e^{-e\bar{\alpha}r+C} &\leq \mu_n^\mp(M_n \leq \mathbb{E}[M_n] - r) \leq e^{-e\alpha_r-C}, \end{aligned}$$

in which, for every $r > 0$, the ratio $\bar{\alpha}r/\alpha_r$ goes to 1 as $\beta \rightarrow \infty$.

PROPOSITION 3. *For no nonrandom sequence (m_n) does $(M_n - m_n)$ converge weakly to a nondegenerate law.*

The above described behavior of uniform Gumbel tails and nonconvergence of its centered maximum matches the behavior exhibited by, for example, the maximum of i.i.d. independent geometric random variables. However, if the Ising interface were instead *tilted* at an angle (say via boundary conditions that are minus above some plane with outward normal $\hat{n} \notin \{\pm e_1, \pm e_2, \pm e_3\}$ and plus below it), a famous open problem is to establish that \mathcal{I} would then be rough even at very low temperatures and resemble the DGFF (see [9] for such a result at zero temperature). It would be interesting to compare the maximum displacement of said tilted interface to that of the DGFF, where the asymptotic behavior of the centered maximum was shown [4] to be tight and, subsequently, found [3, 10] to converge to a randomly-shifted Gumbel distribution.

Unlike the DGFF (and other log-correlated random fields, e.g., BBM)—where the marginal at a site is Gaussian and the difficulty in the analysis of the maximum is due to the logarithmic correlations between sites—in the case of 3D Ising interfaces, obtaining a good understanding of the probability that the interface reaches a height h above a fixed site in $\mathbb{Z}^2 \times \{0\}$ is already a major obstacle.

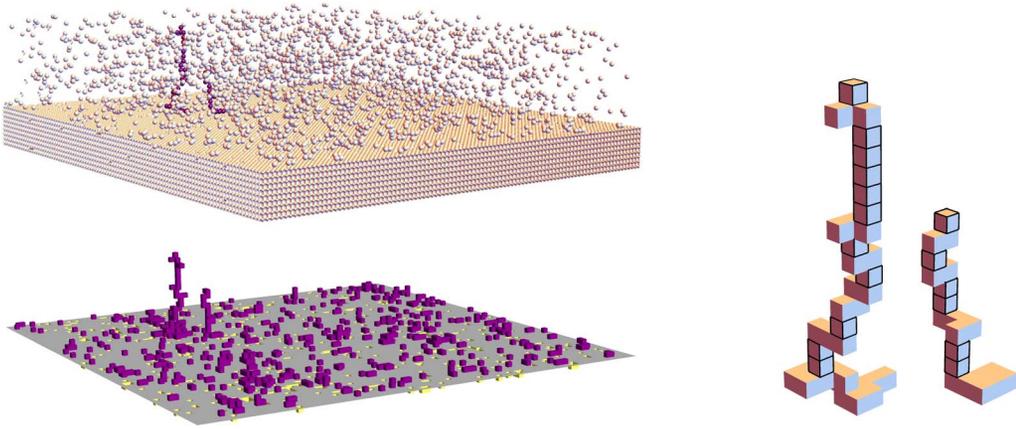


FIG. 1. *Top left: The plus sites of an Ising configuration under μ_n^\mp . Bottom left: From that configuration, one obtains the interface \mathcal{I} by repeatedly flipping the spins in all finite connected components of plus or minus sites and letting \mathcal{I} be the set of faces separating the remaining disagreeing spins. Right: Two neighboring tall pillars of the resulting interface decompose into a base and a sequence of increments delimited by their cut-points (in bold).*

1.1. *Proof strategy and outline.* As in the prequel [16], our analysis of the maximum height of a 3D Ising interface centers on understanding the shape of the interface locally near where it attains atypically large heights. This was formalized via what we refer to as *pillars*: for a face with midpoint $x = (x_1, x_2, 0)$, the pillar of x , denoted \mathcal{P}_x , is obtained from a configuration σ as follows (see Figure 1):

1. Take all finite (+) or (−) clusters (i.e., sites not $*$ -connected to the boundary), simultaneously flip their spins, then repeat this step until no finite clusters remain; the faces separating differing spins in the result comprise the interface \mathcal{I} .
2. Remove from that configuration all sites in the lower half-space $\mathbb{R}^2 \times (-\infty, 0]$.
3. The pillar \mathcal{P}_x is then the (possibly empty) $*$ -connected plus component containing the site with midpoint $x + (0, 0, \frac{1}{2})$, also identified with the set of faces of \mathcal{I} that bound it.

We study the shape of \mathcal{P}_x conditionally on the height of the pillar, $\text{ht}(\mathcal{P}_x)$, exceeding some h . To do so, we define cut-points of \mathcal{P}_x as sites v in \mathcal{P}_x such that no other site in \mathcal{P}_x is at the same height as v . Ordering the cut-points of \mathcal{P}_x in increasing height as $v_1, v_2, \dots, v_{\mathcal{T}+1}$, we decompose \mathcal{P}_x into a sequence of increments $(\mathcal{X}_i)_{i \leq \mathcal{T}}$, where \mathcal{X}_i is the set of sites in \mathcal{P}_x delimited by v_i from below and v_{i+1} from above. We further decompose \mathcal{P}_x into its *base* consisting of the sites in \mathcal{P}_x below v_1 and its *spine* consisting of the sites in \mathcal{P}_x at or above v_1 (due to the difficulty in controlling interactions with nearby pillars, the base was defined differently in [16], namely, it also included a prefix of the spine to mitigate the effect of these interactions).

With these new definitions the results in [16] show that, conditionally on $\text{ht}(\mathcal{P}_x) \geq h$, the pillar \mathcal{P}_x has a base of diameter $O(\log h)$ with an exponential tail beyond that, and all increments above $v_{C \log h}$ have exponential tails on their surface area. It was also shown in [16] that the increment sequence of \mathcal{P}_x , conditioned on having $\mathcal{T} \geq T$ increments, behaves asymptotically as a stationary, weakly mixing sequence; in particular, observables of the increment sequence (e.g., its volume, surface area, and displacement) were shown to obey central limit theorems as $T \rightarrow \infty$. At the level of these central limit theorems and the law of large numbers of M_n , errors of $O(\log h)$ or $O(\log T)$ on the bounds on the size of the base could be sustained.

The following result removes these $O(\log h)$ errors, which we cannot afford in proving Theorems 1–2.

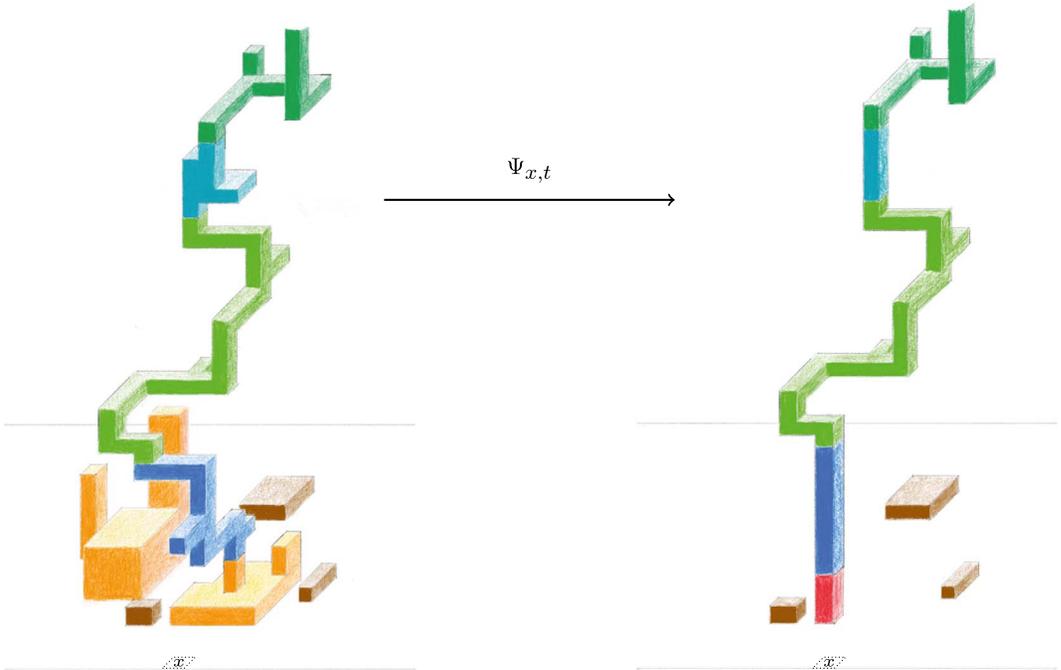


FIG. 2. A map $\Psi_{x,t}$ locally modifies the structure of the interface near the base of a pillar \mathcal{P}_x that reaches a large height $h \gg 1$. The local modifications introduce delicate interactions with nearby fluctuations of the interface (nearby pillars).

THEOREM 4. *There exist $\beta_0, C > 0$ such that, for every $\beta > \beta_0$ and sequences $(h_n), (\Delta_n), (x_n)$ with $h_n \ll \Delta_n$ and $x_n \in \llbracket -n + \Delta_n, n - \Delta_n \rrbracket^2 \times \{0\}$, the pillar \mathcal{P}_x of the interface of the Ising model on Δ_n with Dobrushin boundary conditions, conditional on reaching height h_n , has the following structure:*

1. *The diameter of the base $\mathcal{B}_x := \{w \in \mathcal{P}_x : \text{ht}(w) < \text{ht}(v_1)\}$ is at most r , except with probability $C \exp[-(\beta - C)r]$ for every $1 \leq r \leq \Delta_n$.*
2. *For every t , the surface area of the increment $\mathcal{X}_t := \{w \in \mathcal{P}_x : \text{ht}(v_t) \leq \text{ht}(w) \leq \text{ht}(v_{t+1})\}$ is at most $8 + r$, except with probability $C \exp[-(\beta - C)r]$ for every $1 \leq r \leq \Delta_n$.*

1.1.1. *Proof of Theorem 4.* As in classical Peierls arguments as well as [14, 16], we design a map Ψ on a subset \mathbf{J} of interfaces whose pillars have $\text{ht}(\mathcal{P}_x) \geq h$; the map will show that the subset \mathbf{J} is rare if:

1. $\mu_n^\mp(\mathcal{I})/\mu_n^\mp(\Psi(\mathcal{I}))$ is exponentially small in $\beta(|\mathcal{I}| - |\Psi(\mathcal{I})|)$, while
2. the map has bounded multiplicity, that is, it maps at most C^M elements of \mathbf{J} with $|\mathcal{I}| - |\Psi(\mathcal{I})| = M$ to each element of $\Psi(\mathbf{J})$ (noting $|\mathcal{I}| - |\Psi(\mathcal{I})|$ is the excess energy of the interface \mathcal{I} compared to $\Psi(\mathcal{I})$).

In the context of interfaces of the Ising model, we emphasize that bounding $\mu_n^\mp(\mathcal{I})/\mu_n^\mp(\Psi(\mathcal{I}))$ is complicated by the fact that the Ising measure is not a measure on interfaces but rather on configurations: using cluster expansion, Dobrushin [14] viewed the law $\mu_n^\mp(\mathcal{I} \in \cdot)$ as a perturbation of a measure on interfaces whose weights are proportional to $e^{-\beta|\mathcal{I}|}$, where the perturbation $e^{-\sum_{f \in \mathcal{I}} \mathbf{g}(f, \mathcal{I})}$ captures the subcritical bubbles in the plus and minus phases. Thus, the difficulty in item (1) above reduces to showing that the cumulative effect of $\sum \mathbf{g}(f, \mathcal{I}) - \sum \mathbf{g}(f', \Psi(\mathcal{I}))$ is comparable to the energy gain $|\mathcal{I}| - |\Psi(\mathcal{I})|$.

In Dobrushin’s proof of rigidity, this interaction term was controlled via the decomposition of the interface into *groups of walls* describing the vertical fluctuations of the interface: this

effectively reduced considerations of the interaction terms to horizontal interactions between distinct walls.

In [16], as in this paper, we required estimates conditionally on $\text{ht}(\mathcal{P}_x) \geq h$ and, therefore, had to examine the structure of pillars in a more refined manner than by reducing to their two-dimensional projection. By decomposing pillars into a base and a sequence of increments and accepting $O(\log h)$ errors in the base, we were able to handle the interaction terms by separating out the vertical interactions between shifts of increments from the horizontal interactions induced by deletions of groups of walls.

In both approaches there is an inherent competition between: (1) a desire to delete more walls/straighten additional increments (which simplifies the comparison between the interaction terms and $|\mathcal{I}| - |\Psi(\mathcal{I})|$) and (2) the inability to delete “too much,” at risk of losing control of the multiplicity of the map.

Toward Theorem 4 we need both an exponential tail on the size of the base beyond $O(1)$ and control on increment sizes even at heights that are $O(1)$ (cf., the $O(\log h)$ errors in [16]). As we will see in Section 1.1.2, this is fundamental to establishing Theorems 1–2. To prove Theorem 4, we devise a delicate algorithmic map (Algorithm 1) that iteratively handles the interactions between the horizontal shifts of the increments in the pillar \mathcal{P}_x and the vertical shifts of nearby walls of distinct pillars. See Figure 2 for a visualization of this map and Section 4.2 for a more detailed discussion of the various steps in its construction.

1.1.2. Proof of tightness and Gumbel tails for $M_n - \mathbb{E}[M_n]$. In [16] we used structural results on the shape of pillars \mathcal{P}_x attaining a large height to prove approximate submultiplicativity for the sequence $\mu^\mp(\text{ht}(\mathcal{P}_x) \geq h)$, with a multiplicative error that is $O(e^{\beta \text{diam}(\mathcal{B}_x)^2})$ where \mathcal{B}_x is the base of the pillar \mathcal{P}_x . As the bounds in that paper showed that typically $\text{diam}(\mathcal{B}_x) = O(\log h)$, it follows from Fekete’s Lemma that $\frac{1}{h} \log \mu^\mp(\text{ht}(\mathcal{P}_x) \geq h)$ has a limit α as $h \rightarrow \infty$. Further, deducing in that paper that $M_n / \log n \rightarrow 2/\alpha$, via a second moment argument, relied on the bound on $\text{diam}(\mathcal{B}_x)$, as if $y + (0, 0, \frac{1}{2})$ is interior to \mathcal{B}_x , then $\mathcal{P}_x = \mathcal{P}_y$.

Our proof of tightness of $M_n - \mathbb{E}[M_n]$, therefore, necessitates establishing that $\text{diam}(\mathcal{B}_x) = O(1)$ as well as an $O(1)$ error in the relevant submultiplicativity estimates. Such refined estimates as well as others needed in the second moment argument that is used to establish tightness (such as controlling the increments and tail bounds on the event $\{\text{ht}(\mathcal{P}_x) \geq h\}$ at heights that are $O(1)$), are derived from the new Theorem 4. The Gumbel tails in Theorem 2 are then obtained via a coupling of the maximum M_n to the maximum of $(1 + o(1))(n/L_n)^2$ i.i.d. copies of the maximum at a suitably chosen smaller scale M_{L_n} (see Proposition 7.1), thereby boosting the exponential left tail into a doubly exponential one.

1.2. Organization. Section 2 contains the prerequisite definitions of walls/ceilings in Ising interfaces as per Dobroshin’s framework. Section 3 defines pillars and their decomposition into a base and sequence of increments (refining those of the prequel). Section 4, which is the heart of the proof, defines the map Ψ and proves Theorem 4.1 and Proposition 4.2 which are more detailed versions of Theorem 4 from above. Section 5 proves the refined submultiplicativity estimates (Proposition 5.1 and Corollary 5.2). These are used in Section 6, via a second moment argument, to prove exponential tails—and thus tightness—for the maximum (Proposition 6.1). Section 7 builds the multiscale coupling of the maximum (Proposition 7.1), used to boost the exponential tails into Gumbel tails and prove Theorem 2 as well as Proposition 3.

2. Preliminaries. In this section we formalize the setup of the low-temperature Ising model, define its interface under Dobrushin boundary conditions more precisely and recall the decomposition of this interface into walls and ceilings introduced in [14] to prove rigidity of the interface.

2.1. *Graph notation.* We begin by describing the graph notation we use throughout the paper; though the results in this paper generalize directly to dimensions greater than three, for ease of exposition we present everything in the setup of the three-dimensional integer lattice.

Let \mathbb{Z}^3 be the three-dimensional integer lattice graph with vertex set $\{(v_1, v_2, v_3) \in \mathbb{Z}^3\}$ and edge set identified with the set of nearest-neighbor pairs of vertices $\{v, w\} : d(v, w) = 1\}$, where $d(x, y) = |x - y|$ will always denote the Euclidean distance between two points x, y .

A *face* of \mathbb{Z}^3 is the open set of points in \mathbb{R}^3 bounded by four edges (and four vertices) forming a square of side length one (normal to one of the coordinate axes). A face is *horizontal* if its outward normal vector is $\pm e_3$ and it is *vertical* if its outward normal vector is $\pm e_1$ or $\pm e_2$. A *cell* of \mathbb{Z}^3 is the open set of points bounded by six faces (and eight vertices) forming a cube of side length one. We will frequently identify edges, faces and cells with their midpoints, so that points with two integer and one half-integer coordinate are midpoints of edges, points with one integer and two half-integer coordinates are midpoints of faces and points with three half-integer coordinates are midpoints of cells.

For a set of vertices $\Lambda \subset \mathbb{Z}^3$, we denote by $\mathcal{E}(\Lambda), \mathcal{F}(\Lambda), \mathcal{C}(\Lambda)$ the edges, faces and cells, respectively, whose bounding vertices are all contained in Λ .

Two distinct edges are adjacent if they share a vertex; two distinct faces are adjacent if they share a bounding edge; two distinct cells are adjacent if they share a bounding face. We use the notation “ \sim ” to denote adjacency. A set of faces (resp., edges, cells) F is called connected if, for every $f, f' \in F$, there is a sequence $f_1, \dots, f_m \in F$ such that $f = f_0 \sim f_1 \sim \dots \sim f_m = f'$. We say that two faces (resp., edges, cells) are *connected in Λ* if $F \cap \mathcal{F}(\Lambda)$ is connected (resp., $F \cap \mathcal{E}(\Lambda)$ and $F \cap \mathcal{C}(\Lambda)$ are connected).

Two distinct edges/faces/cells are **-adjacent* if they share a bounding vertex. We define **-connectivity*, analogously to the above definitions for connectivity, w.r.t. the weaker notion of **-adjacency*.

Subsets of \mathbb{Z}^3 . The subsets of \mathbb{Z}^3 that we will primarily consider are boxes or cylinders centered at the origin. Let us denote the centered $(2n + 1) \times (2m + 1) \times (2h + 1)$ box by

$$\Lambda_{n,m,h} = \llbracket -n, n \rrbracket \times \llbracket -m, m \rrbracket \times \llbracket -h, h \rrbracket,$$

where, if $a < b$ are integers, $\llbracket a, b \rrbracket := \{a, a + 1, \dots, b - 1, b\}$. We use Λ_n to denote the infinite cylinder $\Lambda_{n,n,\infty}$.

For any cell-set $C \subset \mathcal{C}(\mathbb{Z}^3)$, its (outer) boundary ∂C is the set of all cells in $\mathcal{C}(\mathbb{Z}^3) \setminus C$ which are adjacent to some cell in C . We use the shorthand $\partial \Lambda_{n,m,h} = \partial(\mathcal{C}(\Lambda_{n,m,h}))$.

Other important subsets we consider are slabs of \mathbb{Z}^3 . For an integer $h \in \mathbb{Z}$, let \mathcal{L}_h be the subgraph of \mathbb{Z}^3 with vertex set $\mathbb{Z}^2 \times \{h\}$ and the resulting face-set. For half-integer $h \in \mathbb{Z} + \frac{1}{2}$, let \mathcal{L}_h consist of the faces and cells of \mathbb{Z}^3 whose midpoints have height h . Let $\mathcal{L}_{>0} = \bigcup_{h>0} \mathcal{L}_h$ be the cell and face-set of the upper half-space, and let $\mathcal{L}_{<0}$ be the cell- and face-set of the lower half-space.

Abusing notation slightly, it will be helpful to use the notation

$$\mathcal{L}_{0,n} = \mathcal{F}(\mathcal{L}_0 \cap \Lambda_n).$$

2.2. *The Ising model.* Since our primary object of study is the interface of the 3D Ising model, it will be convenient to consider the Ising model as an assignment of $\{\pm 1\}$ spins to the vertices of the dual graph $(\mathbb{Z}^3)^*$, identified with the cells of \mathbb{Z}^3 . With this choice the interface will be a connected subset of $\mathcal{F}(\mathbb{Z}^3)$.

An Ising configuration on a subset $\mathcal{C}(\Lambda) \subset \mathcal{C}(\mathbb{Z}^3)$ is an element $\sigma \in \{\pm 1\}^{\mathcal{C}(\Lambda)}$. A boundary condition on $\mathcal{C}(\Lambda)$ is a configuration $\eta \in \{\pm 1\}^{\mathcal{C}(\mathbb{Z}^3)}$. The Ising model at inverse-temperature $\beta > 0$ on $\mathcal{C}(\Lambda)$ with boundary conditions η is the probability measure on $\sigma \in \{\pm 1\}^{\mathcal{C}(\Lambda)}$ given by

$$\mu_{\Lambda,\beta}^\eta(\sigma) = \frac{1}{\mathcal{Z}_{\Lambda,\beta}} \exp[-\beta \mathcal{H}(\sigma)],$$

where

$$\mathcal{H}(\sigma) = - \sum_{\substack{v \sim w \\ v, w \in \mathcal{C}(\Lambda)}} \mathbf{1}\{\sigma_v \neq \sigma_w\} - \sum_{\substack{v \sim w \\ v \in \mathcal{C}(\Lambda), w \in \partial\Lambda}} \mathbf{1}\{\sigma_v \neq \eta_w\}$$

and where the normalizing constant $\mathcal{Z}_{\Lambda, \beta}$, called the *partition function*, is such that $\mu_{\Lambda, \beta}^\eta$ is a probability measure.

We suppress the dependence on β as the choice of β is typically fixed in the context. When $\Lambda = \Lambda_{n, m, h}$, we use the shorthand $\mu_{n, m, h}^\eta = \mu_{\Lambda_{n, m, h}}^\eta$, and when $\Lambda = \Lambda_n = \Lambda_{n, n, \infty}$, we use the shorthand $\mu_n^\eta = \mu_{\Lambda_n, n, \infty}^\eta$.

In this paper we are interested in *Dobrushin boundary conditions*, which are the assignment

$$\eta_v = \begin{cases} -1, & v \in \mathcal{C}(\mathbb{Z}^2 \times \llbracket 0, \infty \rrbracket), \\ +1, & v \in \mathcal{C}(\mathbb{Z}^2 \times \llbracket -\infty, 0 \rrbracket), \end{cases}$$

and we use the shorthand \mp for this choice of η .

Domain Markov property. Observe that the only dependence of the measure $\mu_{\Lambda, \beta}^\eta$ on the boundary conditions η is through the restriction of η to $\partial\Lambda$. This leads to what is known as the domain Markov property: for any two finite subsets $A \subset B \subset \mathcal{C}(\mathbb{Z}^2)$ and every configuration η on $B \setminus A$,

$$\mu_B(\sigma_A \in \cdot \mid \sigma_{B \setminus A} = \eta) = \mu_A^{\eta_{B \setminus A}}(\sigma_A \in \cdot),$$

where σ_A denotes the restriction of σ to the set A .

FKG inequality. The Ising model satisfies an important positive correlation inequality known as the FKG inequality. Consider the natural partial order on configurations $\sigma \in \{\pm 1\}^A$, and suppose f and g are nondecreasing functions in that partial order. Then,

$$\mathbb{E}_{\mu_\Lambda} [f(\sigma)g(\sigma)] \geq \mathbb{E}_{\mu_\Lambda} [f(\sigma)]\mathbb{E}_{\mu_\Lambda} [g(\sigma)],$$

where \mathbb{E}_ν is the expectation with respect to the law ν . A special case of this is when f and g are indicator functions of nondecreasing events.

A recurring example of such an increasing event is $*$ -connectivity via plus cells. For a set $A \subset \mathcal{C}(\Lambda)$, we say $v, w \in A$ are in the same $*$ -connected plus component of A if v and w are $*$ -connected in $\{u \in A : \sigma_u = +1\}$. We use the shorthand $v \xrightarrow{+}_A w$ to denote this event. When $A = \mathcal{C}(\Lambda)$, we omit it from the notation.

Infinite-volume Gibbs measures and DLR condition. If the underlying geometry Λ is an infinite (rather than finite) subset of \mathbb{Z}^3 , the normalizing constant $\mathcal{Z}_{\Lambda, \beta}$ is not finite and the measure $\mu_{\Lambda, \beta}^\eta$ is a priori undefined. Such *infinite-volume Gibbs measures* are instead defined via a consistency relation known as the *DLR conditions*. For an infinite set $\mathcal{C}(\Lambda)$, a measure ν_Λ , defined by its finite-dimensional distributions, satisfies the DLR conditions if, for every finite $A \subset \mathcal{C}(\Lambda)$,

$$\mathbb{E}_{\nu_\Lambda(\sigma_{\mathcal{C}(\Lambda) \setminus A} \in \cdot)} [\nu_\Lambda(\sigma_A \in \cdot \mid \sigma_{\mathcal{C}(\Lambda) \setminus A})] = \nu_\Lambda(\sigma_A \in \cdot).$$

Infinite-volume Gibbs measures need not be unique. For the Ising model on \mathbb{Z}^d , the phase transition of the model is described in terms of the uniqueness/nonuniqueness of the infinite-volume Gibbs measure. In the low-temperature regimes we are interested in, distinct infinite-volume measures are attained by taking weak limits of Ising models on finite boxes with different boundary conditions (e.g., all-plus, all-minus).

We denote the infinite-volume Gibbs measures, obtained by taking limits of $\mu_{n,n,n}^+$ and $\mu_{n,n,n}^-$, as $n \rightarrow \infty$ by $\mu_{\mathbb{Z}^3}^+$ and $\mu_{\mathbb{Z}^3}^-$, respectively. When $\beta > \beta_c$, $\mu_{\mathbb{Z}^3}^+ \neq \mu_{\mathbb{Z}^3}^-$. Dobrushin [13, 14] proved that there exists a β_0 such that, when $\beta > \beta_0$, there exist DLR measures on \mathbb{Z}^3 that are not mixtures of $\mu_{\mathbb{Z}^3}^+$ and $\mu_{\mathbb{Z}^3}^-$, namely, those obtained by taking the limit of $\mu_{n,n,n}^\mp$ as $n \rightarrow \infty$.

2.3. The Ising interface with Dobrushin boundary conditions. The infinite volume measure $\mu_{\mathbb{Z}^3}^\mp$ is characterized by an *interface* separating the minus and plus phases, which look like $\mu_{\mathbb{Z}^3}^-$ and $\mu_{\mathbb{Z}^3}^+$, respectively. Dobrushin showed that this interface is localized—on finite boxes (its height fluctuations above the origin are $O(1)$)—and we are interested in characterizing the law of its maximum height. To that end, we formally define the interface \mathcal{I} separating the plus and minus phases.

DEFINITION 2.1 (Interface). Consider the Ising model with Dobrushin boundary conditions on $\Lambda_{n,m,h}$, that is, $\mu_{n,m,h}^\mp$. For a configuration σ on $\mathcal{C}(\Lambda_{n,m,h})$, define its *interface* $\mathcal{I} = \mathcal{I}(\sigma)$ as follows:

1. Extend the configuration σ to a configuration on all of $\mathcal{C}(\mathbb{Z}^3)$ by taking $\sigma_v = -1$ if $v \in \mathcal{L}_{>0} \setminus \mathcal{C}(\Lambda_{n,m,h})$ and $\sigma_v = +1$ if $v \in \mathcal{L}_{<0} \setminus \mathcal{C}(\Lambda_{n,m,h})$.
2. Let $F(\sigma)$ be the set of faces separating cells of different spins under σ .
3. Call the (maximal) $*$ -connected component of $\mathcal{L}_0 \setminus \mathcal{F}(\Lambda)$ in $F(\sigma)$, the *extended interface*. (This is also the unique infinite $*$ -connected component in $F(\sigma)$.)
4. Let $\mathcal{I} = \mathcal{I}(\sigma)$ be the restriction of the extended interface to $\mathcal{F}(\Lambda)$.

REMARK 2.2. One could use alternative definitions for singling out the interface \mathcal{I} out of the connected sets of faces that separate the minus and plus phases of the boundary, for example, the minimal one or one obtained by a splitting rule. Locally, the difference set between two such definitions would have an exponential tail via a Peierls argument. However, these other choices are not as well tuned to the arguments that follow.

Just as a configuration σ identifies an interface \mathcal{I} , every interface \mathcal{I} identifies a configuration $\sigma(\mathcal{I})$ for which $F(\sigma(\mathcal{I})) = \mathcal{I}$; that is, this configuration is minus everywhere “above” \mathcal{I} and plus “below.” One can obtain this configuration by starting from the boundary sites $\partial\mathcal{C}(\Lambda_{n,m,h})$ and, iteratively, from the boundary inward, assigning spins to the cells of $\Lambda_{n,m,h}$ such that adjacent cells have differing spins if and only if there is a face in \mathcal{I} separating them.

If we call $\mathcal{C}^-(\mathcal{I})$ the set of all minus spins of $\sigma(\mathcal{I})$ and $\mathcal{C}^+(\mathcal{I})$ the set of all plus spins of $\sigma(\mathcal{I})$, we find that each of \mathcal{C}^- and \mathcal{C}^+ is a *single* infinite $*$ -connected set of cells (though we emphasize that each may break up into distinct *nearest-neighbor* connected components $\mathcal{C}_1^-, \dots, \mathcal{C}_r^-$ and $\mathcal{C}_1^+, \dots, \mathcal{C}_s^+$). Recall the following consequence of the definition of the interface above and the domain Markov property of the Ising model.

OBSERVATION 2.3. Conditionally on having an interface \mathcal{I} , the Ising model on $\mathcal{C}(\Lambda_{n,m,h})$ with \pm -boundary conditions is the Ising measure on $\mathcal{C}(\Lambda_{n,m,h}) \setminus \partial\mathcal{I}$ (where $\partial\mathcal{I}$ is the set of all cells that share a bounding vertex with a face of \mathcal{I}) with its induced boundary conditions being \pm on $\partial\Lambda_{n,m,h}$ and being the restriction $\sigma(\mathcal{I})|_{\partial\mathcal{I}}$ on $\partial\mathcal{I}$. This Ising measure is, evidently, a product of Ising measures on the nearest-neighbor connected components $(\mathcal{C}_i^-)_i$ with minus boundary conditions and $(\mathcal{C}_i^+)_i$ with plus boundary conditions.

It follows straightforwardly by Borel–Cantelli that, for every $\beta > 0$, we can take a limit of the measure $\mu_{n,m,h}^\mp$ as $h \rightarrow \infty$ and obtain an infinite-volume measure on the cylinder $\Lambda_{n,m,\infty}$

whose interface is almost surely finite. With this in hand we can move to the Ising interface on $\Lambda_n = \Lambda_{n,n,\infty}$ under μ_n^\mp . As in the preceding works [14, 16], we move from the Ising measure to a measure over interfaces, where the energetic cost of an interface is seen to be given by its cardinality and the lowest energy interface is that coinciding with \mathcal{L}_0 . We therefore define the notion of *excess energy* of one interface with respect to another by

$$m(\mathcal{I}; \mathcal{J}) = |\mathcal{I}| - |\mathcal{J}| = \mathcal{H}(\sigma(\mathcal{J})) - \mathcal{H}(\sigma(\mathcal{I})).$$

Informally, by Observation 2.3, given an interface \mathcal{I} , the measure μ_n^\mp looks like a combination of the measure μ_n^- above \mathcal{I} and μ_n^+ below it. However, the choice of a particular interface modifies these measures above and below the interface as it precludes, say, plus sites that appear under μ_n^- but would be $*$ -adjacent to \mathcal{I} . At low temperatures these plus droplets have exponential tails on their size, and we can sum over their cumulative effect in order to characterize the Ising measure as a Gibbs measure over interfaces with an additional perturbative term.

THEOREM 2.4 ([14], Lemma 1). *Consider the Ising measure $\mu_n = \mu_n^\mp$ on the cylinder $\Lambda_n = \Lambda_{n,n,\infty}$. There exist $\beta_0 > 0$ and a function \mathbf{g} such that, for every $\beta > \beta_0$ and any two interfaces \mathcal{I} and \mathcal{I}' ,*

$$\frac{\mu_n^\mp(\mathcal{I})}{\mu_n^\mp(\mathcal{I}')} = \exp\left[-\beta m(\mathcal{I}, \mathcal{I}') + \left(\sum_{f \in \mathcal{I}} \mathbf{g}(f, \mathcal{I}) - \sum_{f' \in \mathcal{I}'} \mathbf{g}(f', \mathcal{I}')\right)\right],$$

and the function \mathbf{g} satisfies the following for some \bar{c}, \bar{K} independent of β : for all $\mathcal{I}, \mathcal{I}'$ and $f \in \mathcal{I}$ and $f' \in \mathcal{I}'$:

$$(2.1) \quad |\mathbf{g}(f, \mathcal{I})| \leq \bar{K},$$

$$(2.2) \quad |\mathbf{g}(f, \mathcal{I}) - \mathbf{g}(f', \mathcal{I}')| \leq \bar{K} e^{-\bar{c}r(f, \mathcal{I}; f', \mathcal{I}')},$$

where $r(f, \mathcal{I}; f', \mathcal{I}')$ is the largest radius around the origin on which $\mathcal{I} - f$ (\mathcal{I} shifted by the midpoint of the face f) is congruent to $\mathcal{I}' - f'$: that is, to say

$$\mathbf{r}(f, \mathcal{I}; f', \mathcal{I}') = \sup\{r : (\mathcal{I} - f) \cap B_r(0) \equiv (\mathcal{I}' - f') \cap B_r(0)\},$$

where $B_r(0)$ is the ball of radius r around $(0, 0, 0)$ and the congruence relation \equiv is equality as subsets of \mathbb{R}^3 , up to, possibly, reflections and $\pm \frac{\pi}{2}$ rotations in the horizontal plane.

We will use the phrase $\mathbf{r}(f, \mathcal{I}; f', \mathcal{I}')$ is attained by $g \in \mathcal{I}$ (resp., $g' \in \mathcal{I}'$) if g (resp., g') is a face of minimal distance to f (resp., to f') whose presence prevents $\mathbf{r}(f, \mathcal{I}; f', \mathcal{I}')$ from being any larger.

2.4. Walls, ceilings and groups of walls. Dobrushin’s proof of rigidity of the 3D Ising interface used a combinatorial decomposition of the interface to effectively reduce it to a two-dimensional polymer model on \mathcal{L}_0 given by projections of *walls* of \mathcal{I} . We recap the definitions introduced therein in this section and describe the bijection between admissible collections of standard walls and Dobrushin interfaces.

DEFINITION 2.5. For a set of faces or cells A , define its projection $\rho(A) := \{(x_1, x_2, 0) : x = (x_1, x_2, x_3) \in A\}$ so that $\rho(A) \subset \mathcal{L}_0$.

Notice that the projection of a horizontal face is in $\mathcal{F}(\mathcal{L}_0)$ while the projection of a vertical face is in $\mathcal{E}(\mathcal{L}_0)$. For an interface \mathcal{I} and an edge or face $u \in \mathcal{E}(\mathcal{L}_0) \cup \mathcal{F}(\mathcal{L}_0)$, denote by

$$N_\rho(u) := \#\{f \in \mathcal{I} : \rho(f) = u\}.$$

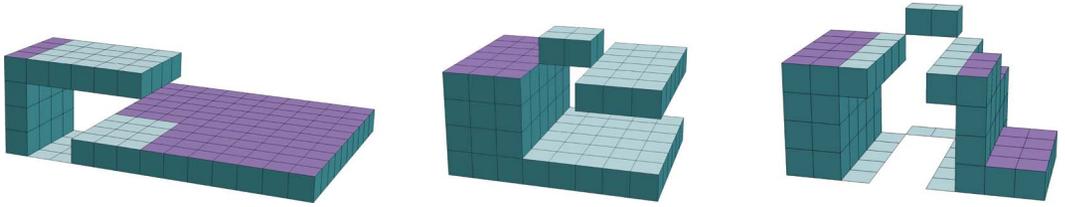


FIG. 3. Three distinct standard walls, with their interior ceiling faces (purple) and wall faces (vertical in teal, horizontal in light blue) as per Definition 2.6. Left example features two distinct (+) components in $\mathbb{R}^2 \times \mathbb{R}_+$ which correspond to two distinct pillars but form a single wall. Middle and right examples feature wall faces that would be disconnected via nearest-neighbor connectivity but form a wall using *-connectivity.

DEFINITION 2.6 (Walls and ceilings). A face $f \in \mathcal{I}$ is a *ceiling face* if f is a horizontal face and $N_\rho(\rho(f)) = 1$. A face $f \in \mathcal{I}$ is a *wall face* if it is not a ceiling face. A *ceiling* is a *-connected set of ceiling faces. By construction, all faces in a ceiling C have the same e_3 coordinate, and we can call that the height of the ceiling $\text{ht}(C)$. A *wall* is a *-connected set of wall faces. Clearly, the projections of distinct walls are disjoint.

See Figure 3 for a depiction of these definitions in subtle scenarios.

DEFINITION 2.7 (Floors and ceilings of a wall). For a wall W , define the complement of its projection

$$\rho(W)^c := (\mathcal{F}(\mathcal{L}_0) \cup \mathcal{E}(\mathcal{L}_0)) \setminus \rho(W),$$

and notice that it splits into an infinite connected component and some finite ones (here, connectivity is seen in \mathbb{R}^2). Any ceiling adjacent to a wall W projects into one of the connected components of $\rho(W)^c$. Call that ceiling that projects into the infinite component of $\rho(W)^c$ the *floor* of W , denoted by $\lfloor W \rfloor$ and collect all other ceilings adjacent to W into $\lceil W \rceil$. For distinct walls W, W' , the sets $\lceil W \rceil$ and $\lceil W' \rceil$ are disjoint.

Importantly, given all the walls of an interface \mathcal{I} , one can reconstruct the full interface by iteratively reading off the heights of the ceilings from the wall collection.

DEFINITION 2.8 (Standard walls). A wall W is a *standard wall* if there exists an interface \mathcal{I}_W such that its only wall is W . As such, a standard wall must have that $\lfloor W \rfloor \subset \mathcal{L}_0$.

For a wall W , we define its *standardization* (called *drift* in [14]) as its vertical shift by $-\text{ht}(\lfloor W \rfloor)$ and denote it by $\theta_{\text{ST}}W$. For any wall W , its standardization is a standard wall.

REMARK 2.9 (Indexing of walls). We can index the walls of \mathcal{I} as follows: assign an arbitrary ordering to the faces of \mathcal{L}_0 . Index a wall W by the minimal face of \mathcal{L}_0 that is interior to W and incident to $\rho(W)$. Clearly, for any admissible collection of standard walls, the indices of distinct walls are distinct.

Since the projections of walls are distinct, the projection of a wall W' is a subset of $\rho(W)^c$: if it is a subset of one of the finite components of $\rho(W)^c$, we say that W' is *nested in* W and write $W' \Subset W$. In this way, to every $W' \Subset W$, we can identify a ceiling in $\lceil W \rceil$ which is the one projecting into that same finite component of $\rho(W)^c$.

DEFINITION 2.10 (Nesting of walls). We say $u \in \mathcal{E}(\mathcal{L}_0) \cup \mathcal{F}(\mathcal{L}_0)$ is *interior to* a wall W if it is not in the infinite component of $\rho(W)^c$. We say that $f \in \mathcal{I}$ is interior to W if $\rho(f)$ is interior to W .

Then, for any $u \in \mathcal{E}(\mathcal{L}_0) \cup \mathcal{F}(\mathcal{L}_0)$, we define its *nested sequence of walls* \mathfrak{W}_u as the collection of all walls to which u is interior. By the definition above, this forms a nested collection $W_1^u \Subset W_2^u \Subset \dots$.

We say a collection of standard walls is *admissible* if their projections are distinct. The following lemma shows a bijection between admissible collections of standard walls and interfaces (also, see [16], Lemma 2.12, for more details).

LEMMA 2.11 (The standard wall representation of \mathcal{I}). *There is a 1-1 correspondence between admissible collections of standard walls and interfaces. Namely, to obtain the standard wall representation of an interface \mathcal{I} , take the union of the standardizations of all its walls. From an admissible collection of standard walls, recover an interface as follows:*

1. Iteratively, for every standardization of a wall $\theta_{\text{ST}}W$:
 - If $W' \Subset \theta_{\text{ST}}W$ and is identified with ceiling $C \in [\theta_{\text{ST}}W]$, then shift W' by $\text{ht}(C)$.
2. From this wall collection, fill in the ceiling faces to obtain the interface \mathcal{I} .

Using the standard wall representation defined above, we note the following important observation.

OBSERVATION 2.12. Consider interfaces \mathcal{I} and \mathcal{J} , such that the standard wall representation of \mathcal{I} contains that of \mathcal{J} (and, additionally, has the standardizations $\mathbf{W} = W_1, \dots, W_r$). By the construction in Lemma 2.11, there is a 1-1 map between the faces of $\mathcal{I} \setminus \mathbf{W}$ and the faces of $\mathcal{J} \setminus \mathbf{H}$ where \mathbf{H} is the set of faces in \mathcal{J} projecting into $\rho(\mathbf{W})$. Moreover, this bijection can be encoded into a map $f \mapsto \theta_{\downarrow} f$ that only consists of vertical shifts and such that all faces projecting into the same component of $\rho(\mathbf{W})^c$ undergo the same vertical shift.

DEFINITION 2.13. For a wall W , define its excess area as

$$m(W) := m(\mathcal{I}_{\theta_{\text{ST}}W}; \mathcal{L}_{0,n}) = |W| - |\mathcal{F}(\rho(W))|,$$

and notice that this always satisfies

$$(2.3) \quad m(W) \geq \frac{1}{2}|W| \quad \text{and} \quad m(W) \geq |\mathcal{E}(\rho(W))| + |\mathcal{F}(\rho(W))|.$$

Notice that for an interface \mathcal{I} with standard wall collection $(W_z)_{z \in \mathcal{L}_{0,n}}$, we have $m(\mathcal{I}; \mathcal{L}_{0,n}) = \sum m(W_z)$.

DEFINITION 2.14 (Closeness and groups of walls). We say that two walls W and W' are close if there exist $u \in \rho(W)$, $u' \in \rho(W')$ such that

$$|u - u'| \leq \sqrt{N_{\rho}(u)} + \sqrt{N_{\rho}(u')}.$$

A collection of walls $F = \bigcup_i W_i$ is a *group of walls* if every wall in F is close to another wall in F and no wall not in F is close to a wall in F . For a nested sequence of walls $\mathfrak{W}_u = W_1^u \cup W_2^u \cup \dots$, this allows us to collect the union of all its groups of walls into

$$\mathfrak{F}_u := \bigcup_i F_i^u \quad \text{where } F_i^u \text{ is the group of walls containing } W_i^u.$$

For collections of walls, for example, groups of walls, nested sequences, define their excess area as the sum of the excess areas of the constituent walls.

Groups of walls are indexed by the minimal index of their constituent walls. However, notice that we do not employ a unique labeling procedure for nested sequences of walls or their groups of walls of nested sequences of walls; if u, u' are both interior to W , then $\mathfrak{W}_u \cap \mathfrak{W}_{u'} \neq \emptyset$.

3. Decomposition of tall pillars.

3.1. *Pillars and increments.* In [16], the authors introduced pillars and their decomposition into an increment sequence in order to understand the large deviations of the interface \mathcal{I} (e.g., its structure at points where it attains atypically large heights). In this section we recall these definitions, though we note crucially that the division of pillars into spines and bases has been modified from the prequel; in [16] we absorbed imprecisions of $O(\log \log n)$, which we cannot afford when proving tightness of the maximum.

DEFINITION 3.1 (Pillar). For an interface \mathcal{I} and a face $x \in \mathcal{F}(\mathcal{L}_0)$, we define the pillar \mathcal{P}_x as follows: consider the Ising configuration $\sigma(\mathcal{I})$, and let $\sigma(\mathcal{P}_x)$ be the (possibly empty) $*$ -connected plus component of the cell with midpoint $x + (0, 0, \frac{1}{2})$ in the upper half-space $\mathcal{L}_{>0}$. The face-set \mathcal{P}_x is then the set of bounding faces of $\sigma(\mathcal{P}_x)$ in $\mathcal{L}_{>0}$.

The following relation between pillars and nested sequences of walls is important.

OBSERVATION 3.2. The walls of the pillar \mathcal{P}_x are contained in the nested sequence of walls $\mathfrak{W}_x = \bigcup_i W_i^x$ together with all walls nested in some W_i^x . Namely, if \mathcal{I} and \mathcal{J} agree on \mathfrak{W}_x and on all walls nested in walls of \mathfrak{W}_x , then $\mathcal{P}_x^{\mathcal{I}} = \mathcal{P}_x^{\mathcal{J}}$. Therefore, if $f \in \mathcal{P}_x$, there exists W such that both f and x are interior to W .

DEFINITION 3.3 (Cut-points). A half-integer $h \in \{\frac{1}{2}, \frac{3}{2}, \dots\}$ is a *cut-height* of \mathcal{P}_x if the intersection $\sigma(\mathcal{P}_x) \cap \mathbb{R}^2 \times [h - \frac{1}{2}, h + \frac{1}{2}]$ consists of a single cell. In that case, that cell (identified with its midpoint $v \in (\mathbb{Z} + \frac{1}{2})^3$) is a cut-point of \mathcal{P}_x . We enumerate the cut-points of \mathcal{P}_x in order of increasing height as v_1, v_2, \dots .

DEFINITION 3.4 (Spine and base). The spine of \mathcal{P}_x , denoted \mathcal{S}_x , is the set of cells in $\sigma(\mathcal{P}_x)$ (resp., faces in \mathcal{P}_x) intersecting $\mathbb{R}^2 \cap [\text{ht}(v_1), \infty)$. The base \mathcal{B}_x of \mathcal{P}_x is the set of cells in $\sigma(\mathcal{P}_x) \setminus \mathcal{S}_x$ (resp., faces in $\mathcal{P}_x \setminus \mathcal{F}(\mathcal{S}_x)$).

REMARK 3.5. We draw attention to the fact that our decomposition of the pillar into a spine and base differs from that used in [16]. There, the beginning of the spine was marked not by v_1 but by a random $v_{\tau_{\text{sp}}}$: the first cut-point to, informally, have height greater than all other pillars in a radius of $R \propto \text{ht}(\mathcal{P}_x)$. This was tailored to the fact that we could sustain errors that were logarithmic in the height of the pillar.

DEFINITION 3.6 (Increments). We decompose a spine \mathcal{S}_x with cut-points v_1, v_2, \dots into its constituent increments. If there are at least $\mathcal{T} + 1 \geq 2$ cut-points, for every $i \leq \mathcal{T}$, define the i th increment as

$$\mathcal{X}_i = \mathcal{S}_x \cap \left(\mathbb{R}^2 \times \left[\text{ht}(v_i) - \frac{1}{2}, \text{ht}(v_{i+1}) + \frac{1}{2} \right] \right),$$

so that the i th increment is the subset of \mathcal{S}_x delimited from below by v_i and from above by v_{i+1} and there are exactly \mathcal{T} increments. (If there are fewer than two cut-points, we say that $\mathcal{T} = 0$.)

Besides the increments, the spine additionally may have a *remainder* $\mathcal{X}_{>\mathcal{T}}$, which we define as the set of faces intersecting $\mathbb{R}^2 \times [\text{ht}(v_{\mathcal{T}+1}), \infty)$. For readability, for a spine \mathcal{S}_x with increment sequence $\mathcal{X}_1, \dots, \mathcal{X}_{\mathcal{T}}, \mathcal{X}_{>\mathcal{T}}$, we use the notation $\mathcal{X}_{\mathcal{T}+1} := \mathcal{X}_{>\mathcal{T}}$ so that we can consistently index over increments and the remainder.

Abusing notation, we may view increments not as subsets of an interface, but as finite $*$ -connected set of cells with at least two cells and whose only cut-points are its bottom-most and top-most cells (modulo lattice translations, achieved by, say, rooting them at the origin). Call the set of all such increments \mathfrak{X} . The face-set of such an increment consists of all its bounding faces, except its bottom-most and top-most horizontal ones. A remainder increment is defined similarly, but its only cut-point is its bottom-most cell.

LEMMA 3.7. *There is a 1-1 correspondence between triplets of v_1 , a sequence of increments $(X_1, \dots, X_T) \in \mathfrak{X}^T$ and a remainder $X_{>T}$, and possible spines of T increments with first cut-point at v_1 .*

Indeed, this follows by identifying the bottom cut-point of X_1 with v_1 and sequentially translating the increments in the increment sequence to identify their bottom cut-point with the top cut-point of the previous increment. For more details, see [16], Section 3.

The simplest increment is what we call the *trivial increment* X_\emptyset , consisting of two vertically consecutive cells, one on top of the other (resp., its eight bounding vertical faces). In proofs where we show that increments have exponential tails, the maps we apply *trivialize* an increment \mathcal{X}_j by replacing it in the increment sequence of \mathcal{S}_x by $\text{ht}(v_{j+1}) - \text{ht}(v_j)$ consecutive trivial increments. Excess areas of increments will be defined w.r.t. this trivialization scheme. Namely, for an increment \mathcal{X}_i ($1 \leq i \leq \mathcal{T}$), define $m(X)$ as

$$(3.1) \quad m(\mathcal{X}_i) = |\mathcal{F}(\mathcal{X}_i)| - 4(\text{ht}(v_{i+1}) - \text{ht}(v_i) + 1)$$

(recall that $\mathcal{F}(X)$ does not include the top most and bottom most faces bounding X). For the remainder increment $\mathcal{X}_{\mathcal{T}+1} = \mathcal{X}_{>T}$, where $v_{\mathcal{T}+2}$ does not exist, this can be defined consistently by arbitrarily setting $\text{ht}(v_{\mathcal{T}+2}) := \text{ht}(\mathcal{P}_x) - \frac{1}{2}$. With these definitions we notice that if $\mathcal{X}_i \neq X_\emptyset$, then

$$m(\mathcal{X}_i) \geq 2(\text{ht}(v_{i+1}) - \text{ht}(v_i) - 1) \vee 2 \quad \text{and} \quad |\mathcal{F}(\mathcal{X}_i)| \leq 3m(\mathcal{X}_i) + 4,$$

since the intersection of \mathcal{X}_i with any height which is not a cut-height has at least six faces vs. four faces in a trivial increment (a nontrivial increment X that has height 1 satisfies $m(X) = 2$ and $|\mathcal{F}(X)| = 10$).

For a spine \mathcal{S}_x and a fixed T , we define its excess area with respect to the reference increment sequence of T trivial increments by

$$m(\mathcal{S}_x) = \left(\sum_{i \leq \mathcal{T}} m(\mathcal{X}_i) \right) + m(\mathcal{X}_{>\mathcal{T}}).$$

We can define an excess area of the base of a pillar as being with respect to the pillar of the same height and no base: for a pillar \mathcal{P}_x with base \mathcal{B}_x and first cut-point v_1 , define

$$m(\mathcal{B}_x) := |\mathcal{F}(\mathcal{B}_x)| - |\mathcal{F}(\rho(\mathcal{B}_x))| - 4 \left(\text{ht}(v_1) - \frac{1}{2} \right).$$

For an $x \in \mathcal{L}_{0,n}$, collect the interfaces with \mathcal{P}_x having at least T increments and at least height H in

$$\mathbf{I}_{x,T,H} = \{ \mathcal{I} : \mathcal{T} \geq T \text{ and } \text{ht}(\mathcal{P}_x) \geq H \}.$$

3.2. *Preliminary estimates on tall pillars.* In this section we recap some results, which can be deduced from Dobrushin’s proof of rigidity [14] and simple modifications around that argument, together with the definitions of pillars and increments; see [16] for short proofs of these.

PROPOSITION 3.8. *There exists C and β_0 such that for every $\beta > \beta_0$, for every $x \in \mathcal{L}_{0,n}$ and every $r \geq 1$,*

$$\mu_n^\mp(\mathfrak{m}(W_x) \geq r) \leq e^{-(\beta-C)r} \quad \text{and} \quad \mu_n^\mp(\mathfrak{m}(\mathfrak{F}_x) \geq r) \leq e^{-(\beta-C)r}.$$

As a consequence of Observation 3.2, the proposition implies the following.

COROLLARY 3.9. *There exists C and β_0 such that for every $\beta > \beta_0$, for every $x \in \mathcal{L}_{0,n}$ and every $r \geq 1$,*

$$\mu_n^\mp(\max\{h : (x_1, x_2, h) \in \mathcal{I}\} \geq r) \leq e^{-4(\beta-C)r} \quad \text{and} \quad \mu_n^\mp(\text{ht}(\mathcal{P}_x) \geq r) \leq e^{-4(\beta-C)r}.$$

In fact, by a simple application of the FKG inequality and forcing argument, we obtain a corresponding lower bound, yielding the following.

PROPOSITION 3.10 (see [16], Proposition 2.29). *There exists $C > 0$ and a sequence ε_β vanishing as $\beta \rightarrow \infty$ such that, for every $\beta > \beta_0$, $x \in \mathcal{L}_{0,n}$ and $h \geq 1$,*

$$(1 - \varepsilon_\beta)e^{-(4\beta+e^{-4\beta})h} \leq \mu_n^\mp(\text{ht}(\mathcal{P}_x) \geq h) \leq e^{-4(\beta-C)h}.$$

3.3. *Tame pillars.* In this section we consider the set of all pillars that have at least T increments and reach a height H . We show that a subset of them, which we call *tame*, have large probability, and from there on in Section 4 we restrict attention to tame pillars on which our future maps will be well defined.

Notice that, if $H \leq T$, the event $\{\text{ht}(\mathcal{P}_x) \geq H\}$ is vacuous, so we take $T < H$.

DEFINITION 3.11. For a given $x \in \mathcal{L}_{0,n}$ and $T < H$, we say that an interface $\mathcal{I} \in \mathbf{I}_{x,T,H}$ is *tame* if \mathcal{I} is in

$$\bar{\mathbf{I}}_{x,T,H} := \left\{ \mathcal{I} \in \mathbf{I}_{x,T,H} : \text{diam}(\mathcal{B}_x) + \frac{1}{4}\mathfrak{m}(\mathcal{S}_x) < d(x, \partial\Lambda_n) \right\}.$$

We observe geometrically that, for a pillar \mathcal{P}_x ,

$$(3.2) \quad \max_{f \in \mathcal{P}_x} d(x, \rho(f)) \leq \text{diam}(\mathcal{B}_x) + \frac{1}{4}\mathfrak{m}(\mathcal{S}_x).$$

Notice that this is less restrictive than the corresponding definition of tameness from [16], as it is the minimal requirement for our (more robust) map in Section 4 to be well defined.

PROPOSITION 3.12. *There exists $C > 0$ such that, for every $\beta > \beta_0$ and every $x \in \mathcal{L}_{0,n}$ and T, H satisfying $0 \leq T < H \ll r$, we have*

$$\mu_n^\mp\left(\text{diam}(\mathcal{B}_x) + \frac{1}{4}\mathfrak{m}(\mathcal{S}_x) \geq r \mid \mathbf{I}_{x,T,H}\right) \leq Ce^{-2(\beta-C)r},$$

and, in particular, taking $r = d(x, \partial\Lambda_n)$,

$$\mu_n^\mp(\bar{\mathbf{I}}_{x,T,H}^c \mid \mathbf{I}_{x,T,H}) \leq Ce^{-2(\beta-C)d(x, \partial\Lambda_n)}.$$

PROOF. This can be read off from a combination of various preliminary bounds in [14, 16]. For the sake of completeness and as an indication of the structure of the proofs in Section 4, we present a full proof using a map $\Phi_{x,T,H}$ which deletes the pillar \mathcal{P}_x and replaces it by a column of $H - 1$ trivial increments.

Namely, let $\Phi_{x,T,H} : \mathbf{I}_{x,T,H} \rightarrow \mathbf{I}_{x,T,H}$ be the following map. First, denote by

$$\mathfrak{F}_{[x]} = \mathfrak{F}_x \cup \bigcup_{f \in \partial_0 x} \mathfrak{F}_f$$

(where $\partial_0 x$ are the four faces of \mathcal{L}_0 adjacent to x). Then, from an interface \mathcal{I} we obtain $\Phi_{x,T,H}(\mathcal{I})$ as the interface with the following standard wall representation:

1. Remove the standardizations of $\mathfrak{F}_{[x]} \cup \mathfrak{F}_{\rho(v_1)}$ from the standard wall representation of \mathcal{I} .
2. Add the standard wall consisting of the bounding faces of a stack of $H - 1$ trivial increments above x (i.e., the cells with midpoints $\{x + (0, 0, n + \frac{1}{2}) : n = 1, \dots, H\}$).

(In the exceptional case $\mathcal{S} = 0$ and v_1 doesn't exist, interpret v_1 as any site of $\sigma(\mathcal{P}_x)$ whose height is $\text{ht}(\mathcal{P}_x) - \frac{1}{2}$.) It is straightforward that $\Phi_{x,T,H}(\mathcal{I})$ is a valid interface in $\mathbf{I}_{x,T,H}$, as we deleted all walls containing x or its adjacent faces in their interior in step (1), so that adding the wall in step (2) preserves the admissibility of the standard wall collection. The pillar $\mathcal{P}_x^{\Phi_{x,T,H}(\mathcal{I})}$ of the resulting interface clearly consists only of the wall added in step (2), and, therefore, it has $H - 1 \geq T$ trivial increments and reaches height H ; hence, $\Phi_{x,T,H}(\mathcal{I}) \in \mathbf{I}_{x,T,H}$.

Notice that, for every $\mathcal{I} \in \mathbf{I}_{x,T,H}$,

$$m(\mathcal{I}; \Phi_{x,T,H}(\mathcal{I})) = m(\mathfrak{F}_{[x]} \cup \mathfrak{F}_{\rho(v_1)}) - 4H \geq 2 \text{diam}(\mathcal{B}_x) + \frac{1}{2} m(\mathcal{S}_x).$$

As such, it suffices for us to show the bound

$$(3.3) \quad \mu_n^\mp(m(\mathcal{I}; \Phi_{x,T,H}(\mathcal{I})) \geq 2r \mid \mathbf{I}_{x,T,H}) \leq e^{-2(\beta-C)r}.$$

We first consider how $\Phi_{x,T,H}$ transforms weights of interfaces. Namely, we claim that the map sends interfaces of low probability to ones of higher probability: there exists $C > 0$ such that, for every $\mathcal{I} \in \mathbf{I}_{x,T,H}$ having $m(\mathcal{I}; \Phi_{x,T,H}(\mathcal{I})) \geq H$,

$$(3.4) \quad \left| \log \frac{\mu_n^\mp(\mathcal{I})}{\mu_n^\mp(\Phi_{x,T,H}(\mathcal{I}))} + \beta m(\mathcal{I}; \Phi_{x,T,H}(\mathcal{I})) \right| \leq C m(\mathcal{I}; \Phi_{x,T,H}(\mathcal{I})).$$

For ease of notation, let $\mathcal{J} = \Phi_{x,T,H}(\mathcal{I})$. We split the set of faces in \mathcal{I} and \mathcal{J} into the following:

- **W**: The faces in $\mathfrak{F}_{[x]} \cup \mathfrak{F}_{\rho(v_1)}$ in \mathcal{I} .
- **B**: All other faces in \mathcal{I} (consisting of all ceiling faces of \mathcal{I} along with all wall faces besides **W**).
- **H**: The set of faces in \mathcal{J} whose projection is in $\mathcal{F}(\rho(\mathbf{W}))$.
- $W_x^\mathcal{J}$: The set of faces in \mathcal{J} from the wall added in step (2) of $\Phi_{x,T,H}$.
- $\theta_\dagger \mathbf{B}$: All other faces in \mathcal{J} .

By Lemma 2.11, there is a 1-1 correspondence between $\theta_\dagger \mathbf{B}$ and faces in **B** given by the vertical shifts induced by ceilings of deleted/added walls from the standard wall representation: encode this 1-1 correspondence into $f \mapsto \theta_\dagger f$. With this splitting in hand, by Theorem 2.4 we need to bound

$$\begin{aligned} \left| \sum_{f \in \mathcal{I}} \mathbf{g}(f, \mathcal{I}) - \sum_{f' \in \mathcal{J}} \mathbf{g}(f', \mathcal{J}) \right| &\leq \sum_{f \in \mathbf{W}} |\mathbf{g}(f, \mathcal{I})| + \sum_{f' \in \mathbf{H}} |\mathbf{g}(f', \mathcal{J})| + \sum_{f' \in W_x^\mathcal{J}} |\mathbf{g}(f', \mathcal{J})| \\ &\quad + \sum_{f \in \mathbf{B}} |\mathbf{g}(f, \mathcal{I}) - \mathbf{g}(\theta_\dagger f, \mathcal{J})|. \end{aligned}$$

The first term is, at most, $\bar{K} |\mathbf{W}| \leq 2\bar{K} m(\mathbf{W})$ by (2.1) and (2.3). The second term is, similarly, at most, $\bar{K} [|\mathcal{E}(\rho(\mathbf{W}))| + |\mathcal{F}(\rho(\mathbf{W}))|] \leq \bar{K} m(\mathbf{W})$ and the third term is, at most, $4\bar{K} H$. The last term satisfies

$$\sum_{f \in \mathbf{B}} |\mathbf{g}(f, \mathcal{I}) - \mathbf{g}(\theta_{\dagger} f, \mathcal{J})| \leq \sum_{f \in \mathbf{B}} \bar{K} e^{-\bar{c}r(f, \mathcal{I}; \theta_{\dagger} \mathcal{J})}.$$

Since the distance between two faces is at least the distance between their projections and the radius $r(f, \mathcal{I}; \theta_{\dagger} f, \mathcal{J})$ must be attained by a wall face, we see that the right-hand side is in turn, at most,

$$\sum_{f \in \mathbf{B}} \sum_{u \in \rho(\mathbf{W} \cup W_x^{\mathcal{J}})} e^{-\bar{c}d(\rho(f), u)} = \sum_{u' \in \rho(\mathbf{W} \cup W_x^{\mathcal{J}})^c} \sum_{u \in \rho(\mathbf{W} \cup W_x^{\mathcal{J}})} N_{\rho}(u') e^{-\bar{c}d(u, u')}.$$

By definition of groups of walls, for every $u' \in \rho(\mathbf{W} \cup W_x^{\mathcal{J}})^c$, we have $N_{\rho}(u') \leq |u - u'|^2 + 1$ (if u' is the projection of a ceiling face, $N_{\rho}(u') = 1$) and, therefore, the right-hand side above is, at most, $\bar{K} [2m(\mathbf{W}) + 4H]$.

Altogether, this implies that, for some C , we have the bound

$$\left| \sum_{f \in \mathcal{I}} \mathbf{g}(f, \mathcal{I}) - \sum_{f' \in \mathcal{J}} \mathbf{g}(f', \mathcal{J}) \right| \leq C [m(\mathbf{W}) + 4H]$$

which implies the bound of (3.4) for a different C as long as $m(\mathcal{I}; \mathcal{J}) \geq H$, say.

On the other hand, let us bound the *multiplicity* of the map $\Phi_{x, T, H}$. Namely, we bound the number of elements in the pre-image $\{\mathcal{I} \in \Phi_{x, T, H}^{-1}(\mathcal{J}) : m(\mathcal{I}; \mathcal{J}) = M\}$ by some C_{Φ}^M uniformly over $\mathcal{J} \in \Phi_{x, T, H}(\mathbf{I}_{x, T, H})$. To do so, we associate to each possible such \mathcal{I} , a $*$ -connected face subset of \mathbb{Z}^3 rooted at x , together with a coloring of those faces by $\{\text{BLUE}, \text{RED}\}$, and bound the number of possible such so-called *witnesses*, from which together with \mathcal{J} we can reconstruct \mathcal{I} . Our witness will consist of the following:

1. Take the standardizations of all walls in $\tilde{\mathfrak{F}}_{[x]} \cup \tilde{\mathfrak{F}}_{\rho(v_1)}$. Color all these faces BLUE.
2. For every $u \in \rho(\mathbf{W})$, add all faces in \mathcal{L}_0 a distance at most $\sqrt{N_{\rho}(u)}$ from u and color them RED.
3. Connect (via a shortest path of faces in \mathcal{L}_0) x to W_1^x and W_i^x to W_{i+1}^x for all i . Do the same for $\rho(v_1)$. Also, add the face at x , and connect x to $\rho(v_1)$. Color those faces added RED.

That this forms a $*$ -connected face subset follows from the definition of closeness of walls. One can easily recover \mathcal{I} from \mathcal{J} and the witness by taking the standard wall representation of \mathcal{J} , removing from it $W_x^{\mathcal{J}}$ and adding in all BLUE faces of our witness to obtain the standard wall representation of \mathcal{I} .

The number of BLUE faces in a witness corresponding to an interface \mathcal{I} with $m(\mathcal{I}; \mathcal{J}) = M$ is, at most, $|\tilde{\mathfrak{F}}_{[x]} \cup \tilde{\mathfrak{F}}_{\rho(v_1)}| \leq 2m(\mathbf{W})$. The number of RED faces added in step (2) of the witness construction is, by definition of closeness and groups of walls, at most,

$$\sum_{u \in \rho(\mathbf{W})} N_{\rho}(u) \leq |\mathbf{W}| \leq 2m(\mathbf{W}).$$

Finally, by Observation 3.2, there is some wall that is deleted to which both x and $\rho(v_1)$ are interior. The number of RED faces added in step (3) of the witness construction is, therefore, at most,

$$3m(\mathbf{W}) + \sum_i m(W_i^x) + \sum_i m(W_i^{\rho(v_1)}) \leq 5m(\mathbf{W}).$$

The number of possible witnesses corresponding to interfaces \mathcal{I} with $m(\mathcal{I}; \mathcal{J}) = M$ is then, at most, the number of possible rooted face subsets of \mathbb{Z}^3 with at most $10M$ faces, multiplied by the number of possible $\{\text{BLUE}, \text{RED}\}$ colorings of those faces. Recall the following combinatorial fact (see, e.g., [14]).

FACT 3.13. There exists a universal constant $s > 0$ such that the number of $*$ -connected face-subsets rooted at (incident to) a fixed vertex, edge or face of \mathbb{Z}^3 , consisting of, at most, $M \geq 1$ faces, is, at most, s^M .

With the above fact in hand, we see that there are, at most, s^{10M} choices for the face subset of the witness and an additional multiplicative 2^{10M} for the number of possible colorings of those faces.

Combining this multiplicity bound with (3.4), we can deduce (3.3) as follows: for $r \geq H$,

$$\begin{aligned} &\mu_n^\mp(\mathfrak{m}(\mathcal{I}; \Phi_{x,T}(\mathcal{I})) \geq r, \mathbf{I}_{x,T,H}) \\ &\leq \sum_{M \geq r} \sum_{\mathcal{J} \in \Phi_{x,T,H}(\mathbf{I}_{x,T,H})} \mu_n^\mp(\mathcal{J}) \sum_{\mathcal{I} \in \Phi_{x,T,H}^{-1}(\mathcal{J}): \mathfrak{m}(\mathcal{I}; \mathcal{J})=M} e^{-(\beta-C)M} \\ &\leq \sum_{M \geq r} \sum_{\mathcal{J} \in \Phi_{x,T,H}(\mathbf{I}_{x,T,H})} e^{-(\beta-C-\log C_\Phi)M} \mu_n^\mp(\mathcal{J}). \end{aligned}$$

Since $\Phi_{x,T,H}(\mathbf{I}_{x,T,H}) \subset \mathbf{I}_{x,T,H}$, we have for some other $C > 0$

$$\mu_n^\mp(\mathfrak{m}(\mathcal{I}; \Phi_{x,T}(\mathcal{I})) \geq r, \mathbf{I}_{x,T,H}) \leq C e^{-(\beta-C)r} \mu_n^\mp(\mathbf{I}_{x,T,H}),$$

from which dividing by $\mu_n^\mp(\mathbf{I}_{x,T,H})$, we obtain (3.3). \square

4. Sharp estimates on the structure of tall pillars. In this section, we obtain estimates on the structure of tall pillars (conditionally on $\mathbf{I}_{x,T,H}$) up to $O(1)$ precision. This is a prerequisite to obtaining tightness of the maximum M_n via a second moment method, as the size of the base \mathcal{B}_x contains the positive correlations between the events $\{\text{ht}(\mathcal{P}_x) \geq h\}$ and $\{\text{ht}(\mathcal{P}_y) \geq h\}$: for example, if the base \mathcal{B}_x contains y in its interior, the events are fully correlated.

THEOREM 4.1. *There exist $\beta_0, C > 0$ such that for every $\beta > \beta_0$, every $x \in \mathcal{L}_{0,n}$ and T, H satisfying $0 \leq T < H \ll d(x, \partial \Lambda_n)$, the following holds:*

(a) Base estimate: For every $1 \leq r \leq d(x, \partial \Lambda_n)$,

$$\mu_n^\mp\left(\left|v_1 - \left(x + \left(0, 0, \frac{1}{2}\right)\right)\right| \geq r \mid \mathbf{I}_{x,T,H}\right) \leq C e^{-2(\beta-C)r},$$

and, in fact,

$$\mu_n^\mp(\mathfrak{m}(\mathcal{B}_x) \geq r \mid \mathbf{I}_{x,T,H}) \leq C e^{-(\beta-C)r}.$$

(b) Increment estimate: For every $t \geq 1$ and every $1 \leq r \leq d(x, \partial \Lambda_n)$,

$$\mu_n^\mp(\mathfrak{m}(\mathcal{X}_t) \geq r \mid \mathbf{I}_{x,T,H}) \leq C e^{-(\beta-c)r}.$$

PROPOSITION 4.2. *There exist $\beta_0, C > 0$ such that for every $\beta > \beta_0$, every $x \in \mathcal{L}_{0,n}$ and T, H satisfying $0 \leq T < H \ll d(x, \partial \Lambda_n)$, the following holds. For every half-integer $\ell \geq \frac{1}{2}$ and $1 \leq r \leq d(x, \partial \Lambda_n)$,*

$$\mu_n^\mp(|\mathcal{F}(\mathcal{P}_x \cap \mathcal{L}_\ell)| \geq 4 + r \mid \mathbf{I}_{x,T,H}) \leq C e^{-(\beta-C)r}.$$

Recall from the Introduction that in [16] the authors proved a bound of $O(\log H)$ on $\text{diam}(\mathcal{B}_x)$ and the exponential tails on increment sizes were restricted to those with index above $O(\log H)$. In that work, affording an $O(\log H)$ error, the interactions between the horizontal shifts of the spine under the map were decoupled from the base and nearby pillars

because, with high probability, no other pillars in the shadow of the pillar \mathcal{P}_x reach a height larger than $O(\log H)$.

At heights that are $O(1)$, we need to deal directly and simultaneously with the interactions between vertical shifts (arising from deletions of groups of walls as in [14]) and the horizontal shifts arising from trivializing increments and shifting the spine appropriately. This induces substantial complications. After defining a map $\Psi_{x,t} : \bar{\mathbf{I}}_{x,T,H} \rightarrow \mathbf{I}_{x,T,H}$, in Section 4.2 we give a reader’s guide to the various difficulties encountered in construction of this map and justify the necessity of its various steps.

4.1. *A new base and increment map.* In this section we define a new map that shrinks the base of \mathcal{P}_x and trivializes the t th increment of the pillar. The map is significantly more involved than the maps in [16], as it deals directly with the interactions between the horizontal shifts of \mathcal{S}_x with the walls near its base which the spine may get close to or hit.

For an increment j , denote the centered *trivialization* of \mathcal{X}_j by

$$\Theta_{\emptyset} \mathcal{X}_j = \left[-\frac{1}{2}, \frac{1}{2} \right]^2 \times \left[\text{ht}(v_j) - \frac{1}{2}, \text{ht}(v_{j+1}) + \frac{1}{2} \right].$$

Denote by $W_y^{\mathcal{K}}$ the wall indexed by y in the interface \mathcal{K} . If $(\tilde{W}_z)_{z \in \mathcal{L}_{0,n}}$ is the standard wall representation of \mathcal{K} , let

$$\Theta_{\dagger}^{\mathcal{K}} \tilde{W}_y = \{ W_y^{\mathcal{K}'} \cup \lceil W_y^{\mathcal{K}'} \rceil : \mathcal{K}' \text{ has standard wall representation } (\tilde{W}_z) \text{ or } (\tilde{W}_z \setminus \tilde{\mathfrak{F}}_{y'}) \}$$

for some y' ;

namely, $\Theta_{\dagger}^{\mathcal{K}} \tilde{W}_y$ is the set of all possible vertical shifts induced on \tilde{W}_y via Lemma 2.11 by deleting the group of walls of a nested sequence of walls.

Finally, for some $x \in \mathcal{L}_{0,n}$ identified with its midpoint, an interface \mathcal{I} and two shift vectors $\omega_1, \omega_2 \in \mathbb{Z}^2 \times \{0\}$, denote by (\tilde{W}_z) the standard wall representation of $\mathcal{I} \setminus \mathcal{S}_x$, and for every wall \tilde{W}_y define

$$\mathfrak{D}_x(\tilde{W}_y, j, \omega_1, \omega_2) = d(\Theta_{\dagger}^{\mathcal{I} \setminus \mathcal{S}_x} \tilde{W}_y, \mathcal{X}_j \cup \theta_{\rho(x+\omega_1+\omega_2)} \mathcal{X}_j \cup \theta_{\rho(x+\omega_2)} \Theta_{\emptyset} \mathcal{X}_j),$$

where for $\omega \in \mathbb{R}^3$ and $A \subset \mathbb{R}^3$, $\theta_{\omega} A$ is the shift of A by the vector ω , that is, $\theta_{\omega} A = A + \omega$.

DEFINITION 4.3. For $x \in \mathcal{L}_{0,n}$, every $t \geq 1$ and every $0 \leq T < H$, define the map $\Psi_{x,t} : \bar{\mathbf{I}}_{x,T,H} \rightarrow \mathbf{I}_{x,T,H}$, as specified in Algorithm 1.

REMARK 4.4. In the exceptional case $T = 0$, when we are applying $\Psi_{x,t}$ to interfaces having pillars \mathcal{P}_x with $\mathcal{T} = 0$ increments (so that they have either zero or one cut-points), we interpret the steps in $\Psi_{x,t}$ in the following way for it to be well defined. If $\mathcal{T} = 0$ but v_1 exists, then recalling that $\mathcal{X}_1 := \mathcal{X}_{>0}$, the remainder will be trivialized, and the rest of the map is applied as is. If $\mathcal{T} = 0$ and \mathcal{P}_x has no cut-points, then take an arbitrary face of \mathcal{P}_x having height $\text{ht}(\mathcal{P}_x)$ to stand-in as “ v_1 ,” and steps 4–6 will be vacuous.

Notice, more generally, that if $t > \mathcal{T}$, step 5 would be vacuous, but the map is still well defined and our results hold by interpreting $m(\mathcal{X}_t) = 0$ if $\mathcal{X}_t = \emptyset$.

4.2. *Strategy of the map $\Psi_{x,t}$.* We now motivate the different steps in the map $\Psi_{x,t}$ and describe why each one is important to the trade-off described in Section 1.1.1 between control of the interaction terms and the multiplicity of the map. Let us recall in more detail the maps introduced in the prequel [16] on pillars that reach height H and used there to establish a bound of $O(\log H)$ on the height of the base and exponential tails on the increments above that height. Let $v_{\tau_{\text{sp}}}$ be a special cut-point index of the spine, marking the first increment

Algorithm 1: The map $\Psi_{x,t}$

1 Let $\{\tilde{W}_y : y \in \mathcal{L}_{0,n}\}$ be the standard wall representation of the interface $\mathcal{I} \setminus \mathcal{S}_x$. Also, let \mathcal{O}_{v_1} be the nested sequence of walls of v_1 , so that $\theta_{\text{ST}}\mathcal{O}_{v_1} = \tilde{\mathfrak{W}}_{v_1}$.

// Base modification

2 Mark $[x] = \{x\} \cup \partial_0 x$ and $\rho(v_1)$ for deletion (where $\partial_0 x$ denotes the four faces in \mathcal{L}_0 adjacent to x).

3 if the interface with standard wall representation $\tilde{\mathfrak{W}}_{v_1}$ has a cut-height, then

- └ Let h^\dagger be the height of the highest such cut-height.
- └ Let y^\dagger be the index of a wall that intersects $(\mathcal{P}_x \setminus \mathcal{O}_{v_1}) \cap \mathcal{L}_{h^\dagger}$ and mark y^\dagger for deletion.

// Spine modification (A): the 1st increment

4 Set $s_1 \leftarrow 0$ and $y_A^* \leftarrow \emptyset$.

for $j = 1$ to $\mathcal{T} + 1$ do

- └ Let $s \leftarrow s_j$ and $s_{j+1} \leftarrow s_j$.
- └ if $m(\mathcal{X}_j) \geq j - 1$ then // (A1)
 - └ Let $s_{j+1} \leftarrow j$.
- └ if $\mathcal{D}_x(\tilde{W}_y, j, -v_{s+1}, 0) \leq m(\tilde{W}_y)$ for some y then // (A2)
 - └ Let $s_{j+1} \leftarrow j$ and mark for deletion every y for which (A2) holds.
- └ if $\mathcal{D}_x(\tilde{W}_y, j, -v_{s+1}, 0) \leq (j - 1)/2$ for some y then // (A3)
 - └ Let $s_{j+1} \leftarrow j$ and let y_A^* be the minimal index y for which (A3) holds.

Let $j^* \leftarrow s_{\mathcal{T}+2}$ and mark y_A^* for deletion.

// Spine modification (B): the t -th increment

5 if $t > j^*$ then

- └ Set $s_t \leftarrow t - 1$ and $y_B^* \leftarrow \emptyset$.
- └ for $k = t$ to $\mathcal{T} + 1$ do
 - └ Let $s \leftarrow s_k$ and $s_{k+1} \leftarrow s_k$.
 - └ if $m(\mathcal{X}_k) \geq k - t$ then // (B1)
 - └ Let $s_{k+1} \leftarrow k$.
 - └ if $\mathcal{D}_x(\tilde{W}_y, j, -v_{s+1}, v_t - v_{j^*+1}) \leq m(\tilde{W}_y)$ for some y then // (B2)
 - └ Let $s_{k+1} \leftarrow k$ and mark for deletion every y for which (B2) holds.
 - └ if $\mathcal{D}_x(\tilde{W}_y, j, -v_{s+1}, v_t - v_{j^*+1}) \leq (k - t)/2$ for some y then // (B3)
 - └ Let $s_{k+1} \leftarrow k$ and let y_B^* be the minimal index y for which (B3) holds.
- └ Let $k^* \leftarrow s_{\mathcal{T}+2}$ and mark y_B^* for deletion.

else

- └ Let $k^* \leftarrow j^*$.

6 foreach index $y \in \mathcal{L}_{0,n}$ marked for deletion do delete $\tilde{\mathfrak{F}}_y$ from the standard wall representation (\tilde{W}_y) .

7 Add a standard wall $W_x^{\mathcal{J}}$ consisting of $\text{ht}(v_1) - \frac{1}{2}$ trivial increments above x .

8 Let \mathcal{K} be the (unique) interface with the resulting standard wall representation.

9 Denoting by $(\mathcal{X}_i)_{i \geq 1}$ the increment sequence of \mathcal{S}_x , set

$$S \leftarrow \begin{cases} \underbrace{(X_\emptyset, X_\emptyset, \dots, X_\emptyset)}_{\text{ht}(v_{j^*+1}) - \text{ht}(v_1)}, \mathcal{X}_{j^*+1}, \dots, \mathcal{X}_{t-1}, \underbrace{X_\emptyset, X_\emptyset, \dots, X_\emptyset}_{\text{ht}(v_{k^*+1}) - \text{ht}(v_t)}, \mathcal{X}_{k^*+1}, \dots) & \text{if } t > j^*, \\ \underbrace{(X_\emptyset, X_\emptyset, \dots, X_\emptyset)}_{\text{ht}(v_{j^*+1}) - \text{ht}(v_1)}, \mathcal{X}_{j^*+1}, \dots) & \text{if } t \leq j^*. \end{cases}$$

10 Obtain $\Psi_{x,t}(\mathcal{I})$ by appending the spine with increment sequence S to \mathcal{K} at $x + (0, 0, \text{ht}(v_1))$.

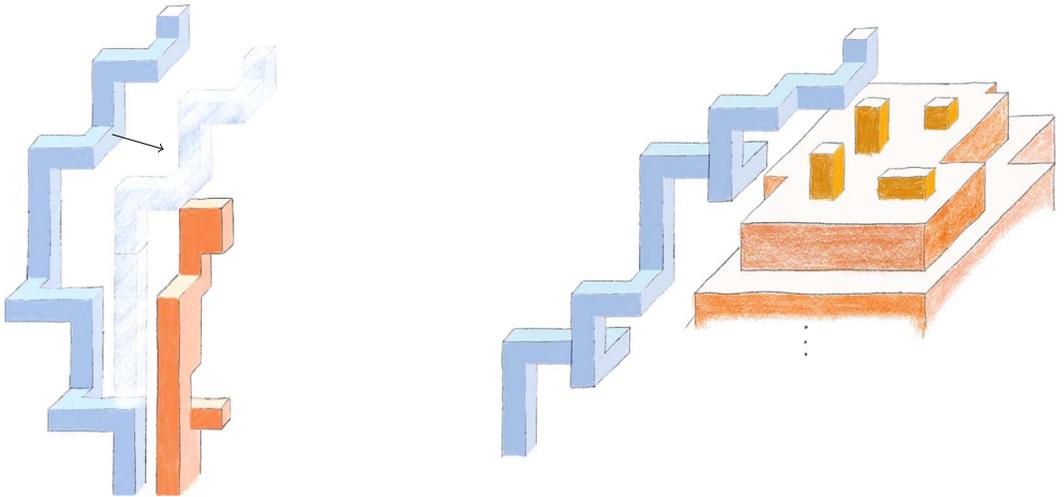


FIG. 4. After trivializing increments $\mathcal{X}_1, \mathcal{X}_4$, the shifted spine intersects an adjacent wall (orange). This constitutes an (A2) violation. Right: The walls in yellow-brown are small (and, therefore, do not violate (A2) but still interact strongly with the increment sequence; these will violate criterion (A3).

whose height is larger than all other pillars in a ball of radius CT about x . For $t > \tau_{\text{SP}}$, the map that proved an exponential tail on the t th increment would simply “trivialize” \mathcal{X}_t and \mathcal{X}_j ($j > t$) in the increment sequence of \mathcal{S}_x if $m(\mathcal{X}_j) \geq e^{\tilde{c}(j-t)/2}$. In bounding the interactions by (2.2), this competed with $e^{\tilde{c}(j-\tau_{\text{SP}})} \leq e^{-\tilde{c}(j-t)}$ (because the horizontal shift of the portion of the spine above v_{t+1} keeps the distance between \mathcal{X}_j and $\mathcal{I} \setminus \mathcal{P}_x$ at least $j - \tau_{\text{SP}}$), and summing these terms over j was $O(1)$.

However, for $t < \tau_{\text{SP}}$, we have no control on the distance between the new spine and walls in $\mathcal{I} \setminus \mathcal{P}_x$; in fact, horizontal shifts of increments $j \in \llbracket t + 1, \tau_{\text{SP}} \rrbracket$ could even *hit* a neighboring wall of $\mathcal{I} \setminus \mathcal{S}_x$ (Figure 4, left), and the map would not yield a valid interface. We thus have to consider the full geometry of these interactions as walls \tilde{W}_y undergo vertical shifts and nearby increments simultaneously undergo horizontal shifts.

With these difficulties laid out, we discuss the various steps in the definition of $\Psi_{x,t}$ and the different scenarios they are designed to address. The base modification (steps 2–3) here is very similar to that used in [16]: it marks the nested sequence $\tilde{\mathfrak{M}}_{[x]}$ for deletion so that the modified spine \mathcal{S} can later be placed above $W_x^{\mathcal{J}}$ at x to form the new pillar $\mathcal{P}_x^{\mathcal{J}}$; the additional deletion of $\tilde{\mathfrak{M}}_{\rho(v_1)}$ and $\tilde{\mathfrak{M}}_{y^\dagger}$ is to exploit the fact that \mathcal{P}_x has no cut-points below v_1 and ensure that the gain in energy $m(\mathcal{I}; \Psi(\mathcal{I}))$ is larger than $\text{ht}(v_1) - \frac{1}{2}$.

The spine modification is substantially more involved. Note that the modifications in (A) (Step 4) and those in (B) (Step 5) are essentially the same, with the latter applied at the t th increment so that we can prove the exponential tail on $m(\mathcal{X}_t)$ simultaneously with the exponential tail on $m(\mathcal{B}_x)$. Thus, let us only discuss the steps in the former (the spine modification (A) at the first increment):

(A1) aims to control interactions between the horizontal shifts of the increments within the spine itself. Unlike [16], where the corresponding modification used a threshold of $m(\mathcal{X}_j) \geq e^{\tilde{c}(j-t)/2}$, which was in a sense the “most lenient” criterion for trivializing increments, it is important that here we use a “strictest possible” criterion only allowing a linear growth of the excess areas of increments.

(A2) ensures that after $\Psi_{x,t}$ is applied, any walls that were hit by horizontal shifts of the spine are deleted. This is achieved via a soft threshold comparing the distance between various horizontal shifts of \mathcal{X}_i to a wall W_y (the relevant quantity in bounding $|\mathbf{g}(f, \mathcal{I}) -$

$\mathbf{g}(f', \Psi_{x,t}(\mathcal{I}))$ via (2.2)), with the excess area $m(W_y)$. Notice that the threshold cannot be done with respect to $m(\mathcal{W}_y)$ instead of $m(W_y)$ because one large wall nesting many smaller walls only counts once toward $m(\mathcal{I}; \Psi_{x,t}(\mathcal{I}))$. (See Figure 4, left.)

(A3) addresses the additional scenario in which many distinct walls of small excess area W_1, W_2, \dots are nested in some W' , and the spine draws close to W_1, W_2, \dots without violating (A2). In this situation, only one *highest* nested sequence of walls violating this criterion is deleted in addition to the trivialization of the increments; otherwise, we would again be overcounting the nesting wall. (See Figure 4, right.)

Finally, it is crucial that the horizontal shifts ω_1, ω_2 considered in $\mathfrak{D}(\tilde{W}_y, j, \omega_1, \omega_2)$ in (A2)–(A3) are to be determined in an algorithmic manner. Namely, we want to ensure that, if an increment is *not* trivialized, its horizontal shift in $\Psi_{x,t}(\mathcal{I})$ did not violate any of the criteria above; the horizontal shift vector with which this needs to be checked is determined by the last increment to have violated one of the trivialization criteria.

4.3. *Properties of $\Psi_{x,t}$.* In this section we begin by showing that the map $\Psi_{x,t}$ is well defined on tame interfaces. We then give a decomposition of the interfaces \mathcal{I} and $\Psi_{x,t}(\mathcal{I})$ and prove some simple inequalities on the excess area $m(\mathcal{I}; \mathcal{J})$.

PROPOSITION 4.5 (Well-definedness of map $\Psi_{x,t}$). *For every $T < H$ and every $\mathcal{I} \in \bar{\mathbf{I}}_{x,T,H}$, the interface $\Psi_{x,t}(\mathcal{I})$ is well defined and is an element of $\mathbf{I}_{x,T,H}$.*

PROOF. First, we claim that the standard wall representation obtained after step 7 is admissible. This is because, after $\tilde{\mathfrak{F}}_{[x]} = \tilde{\mathfrak{F}}_x \cup \cup_{f \in \partial_0 x} \tilde{\mathfrak{F}}_f$ is deleted in step 6, the wall $W_x^{\mathcal{J}}$ has disjoint projection from all remaining standard walls. We next must ensure that, when adding the modified \mathcal{S} in step 10 to \mathcal{K} , it does not intersect any part of the preexisting interface or $\partial\Lambda_n$.

For this, notice that if $t > j^*$, then \mathcal{S} is exactly

$$\begin{aligned} & \bigcup_{1 \leq j \leq j^*} \theta_{\rho(x)} \Theta_{\emptyset} \mathcal{X}_j \cup \bigcup_{j^* < i < t} \theta_{\rho(x-v_{j^*+1})} \mathcal{X}_i \cup \bigcup_{t \leq k \leq k^*} \theta_{\rho(x+v_t-v_{j^*+1})} \Theta_{\emptyset} \mathcal{X}_k \\ & \cup \bigcup_{i > k^*+1} \theta_{\rho(x-v_{k^*+1}+v_t-v_{j^*+1})} \theta \mathcal{X}_i, \end{aligned}$$

and if $t \leq j^*$, then \mathcal{S} is exactly

$$\bigcup_{1 \leq j \leq j^*} \theta_{\rho(x)} \Theta_{\emptyset} \mathcal{X}_j \cup \bigcup_{i > j^*} \theta_{\rho(x-v_{j^*+1})} \mathcal{X}_i.$$

Now, make the following observation regarding the sequence of shifts observed while running $\Psi_{x,t}$.

OBSERVATION 4.6. The sequence (s_i) has $s_{i+1} \neq s_i$ if and only if one of criteria (A1), (A2), (A3) or (B1), (B2), (B3) were attained for i , in which case $s_{i+1} = i$. Consequently, $s_i = j^*$ for every $j^* < i < t$, and $s_i = k^*$ for every $i > k^*$.

Thus, whether $t > j^*$ or $t \leq j^*$, the shifts and trivializations comprising \mathcal{S} are considered in the criteria

$$\mathfrak{D}_x(\tilde{W}_y, j, -v_{s_j+1}, 0) \leq (j-1)/2 \quad \text{and} \quad \mathfrak{D}_x(\tilde{W}_y, j, -v_{s_j+1}, v_t - v_{j^*+1}) \leq (k-t)/2.$$

Also, by definition every face f of $\mathcal{K} \setminus (W_x^{\mathcal{J}} \cup \mathcal{L}_0)$ is in $\Theta_{\dagger}^{\mathcal{I} \setminus \mathcal{S}_x} \tilde{W}_y$ for some y . As such, if \mathcal{S} intersects some preexisting part of $\mathcal{K} \setminus W_x^{\mathcal{J}}$, there would have been some pair y, j such

that the above distance \mathfrak{D}_x would be zero; in that case, that y would have been marked for deletion, and the corresponding face in \mathcal{K} would be in \mathcal{L}_0 yielding a contradiction.

In order to see that the addition of \mathcal{S} does not hit $\partial\Lambda_n$, we use the definition of tameness. In particular, the horizontal displacement of the spine \mathcal{S} from x is always bounded above by

$$\text{diam}(\mathcal{B}_x) + \frac{1}{4}m(\mathcal{S}_x) < d(x, \partial\Lambda_n),$$

where the inequality was by definition of $\mathcal{I} \in \bar{\mathbf{I}}_{x,T,H}$ and $\Psi_{x,t}(\mathcal{I})$ is a valid interface.

Finally, we observe that the resulting interface is in $\mathbf{I}_{x,T,H}$. Notice that the resulting pillar of x in $\Psi_{x,t}(\mathcal{I})$ consists of $W_x^{\mathcal{J}} \cup \mathcal{S}$; on the one hand, this has at least T increments since trivializing increments only increases the total number of increments and, on the other hand, it has the same height as \mathcal{P}_x by construction. \square

4.3.1. *Decomposition of the interfaces.* Fix any interface $\mathcal{I} \in \bar{\mathbf{I}}_{x,T,H}$, and, for ease of notation, let $\mathcal{J} = \Psi_{x,t}(\mathcal{I})$. We begin by partitioning the faces of \mathcal{I} and \mathcal{J} into their constituent parts, as dictated by the map $\Psi_{x,t}$. This partitioning will govern the pairings of $\mathbf{g}(f, \mathcal{I})$ with $\mathbf{g}(f', \mathcal{J})$ when applying (2.2).

Let $\mathbf{Y} \subset \mathcal{L}_0$ be the set of indices of walls in $\mathcal{I} \setminus \mathcal{S}_x$ that were marked for deletion. Let $\mathbf{D} \subset \mathcal{L}_0$ be the indices of walls that were deleted (i.e., walls in $\bigcup\{\tilde{\mathfrak{F}}_y : y \in \mathbf{Y}\}$). Split up the faces of \mathcal{I} as follows:

$\mathbf{X}_A^{\mathcal{I}}$	$\bigcup_{1 \leq j \leq j^*} \mathcal{F}(\mathcal{X}_j)$	Increments between v_1 and v_{j^*+1}
$\mathbf{X}_1^{\mathcal{I}}$	$\bigcup_{j^*+1 \leq j \leq t-1} \mathcal{F}(\mathcal{X}_j)$	If $t > j^*$, increments between v_{j^*+1} and v_t
$\mathbf{X}_B^{\mathcal{I}}$	$\bigcup_{t \leq k \leq k^*} \mathcal{F}(\mathcal{X}_k)$	If $t > j^*$, increments between v_t and v_{k^*+1}
$\mathbf{X}_2^{\mathcal{I}}$	$\bigcup_{k \geq k^*} \mathcal{F}(\mathcal{X}_k)$	Increments above v_{k^*+1}
\mathbf{W}	$\bigcup_{z \in \mathbf{D}} \tilde{W}_z$	All walls that were deleted
\mathbf{B}	$\mathcal{I} \setminus (\mathbf{X}_A^{\mathcal{I}} \cup \mathbf{X}_1^{\mathcal{I}} \cup \mathbf{X}_B^{\mathcal{I}} \cup \mathbf{X}_2^{\mathcal{I}} \cup \mathbf{W})$	The remaining set of faces in \mathcal{I}

where \mathbf{B} splits further into

\mathbf{C}_1	$\bigcup_{y \in \mathbf{Y}} \lceil \tilde{W}_y \rceil$	Ceilings of walls marked for deletion
\mathbf{C}_2	$\bigcup_{z \in \mathbf{D} \setminus \mathbf{Y}} \lceil \tilde{W}_z \rceil \cup \bigcup_{z \notin \mathbf{D}} \tilde{W}_z \cup \lceil \tilde{W}_z \rceil$	Ceilings of walls that were not marked, along with all nondeleted walls and their ceilings
\mathbf{F}	$(\mathbf{B} \cap \mathcal{L}_0) \setminus \bigcup_z \tilde{W}_z \cup \lceil \tilde{W}_z \rceil$	Faces in \mathbf{B} that are only floors of walls

See Figure 5 for a depiction of this splitting.

We next partition the faces of \mathcal{J} . Let us first introduce a few pieces of notation. Denote by $\mathcal{P}_x^{\mathcal{J}}$ the pillar of x in $\mathcal{J} = \Psi_{x,t}(\mathcal{I})$; observe that, by construction, in \mathcal{J} the spine $\mathcal{S}_x^{\mathcal{J}}$ is all of $\mathcal{P}_x^{\mathcal{J}}$. For a given \mathcal{I} , x , t , define the shift map $\tilde{\theta}$ as the horizontal shift on the increments $\mathcal{X}_{j^*+1}, \dots, \mathcal{X}_{t-1}$ and $\mathcal{X}_{k^*+1}, \dots$ from $\Psi_{x,t}$. Namely, for $f \in \mathbf{X}_1^{\mathcal{I}} \cup \mathbf{X}_2^{\mathcal{I}}$, let

$$(4.1) \quad \tilde{\theta} f = \begin{cases} \theta_{\rho(x-v_{j^*+1})} f & \text{if } f \in \mathbf{X}_1^{\mathcal{I}} \text{ or if } f \in \mathbf{X}_2^{\mathcal{I}} \text{ and } t \leq j^*, \\ \theta_{\rho(x-v_{k^*+1+v_t-v_{j^*+1}})} f & \text{if } f \in \mathbf{X}_2^{\mathcal{I}} \text{ and } t > j^*. \end{cases}$$

We can also define the map θ_{\ddagger} on faces in \mathbf{B} that vertically shifts faces of \mathbf{B} to obtain corresponding faces of \mathcal{J} , as dictated by the bijection Lemma 2.11 and the removal of the walls

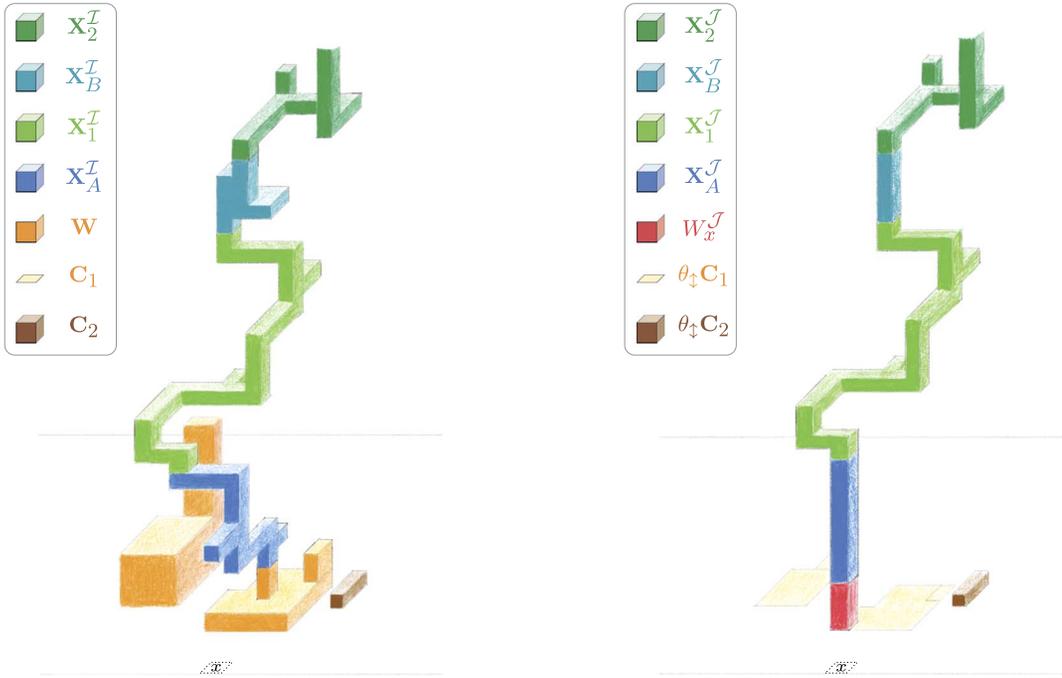


FIG. 5. The decomposition of the interfaces \mathcal{I} (left) and \mathcal{J} (right) into their constituent parts.

in \mathbf{W} . With these, let:

$\mathbf{X}_A^{\mathcal{J}}$	$\Theta_{\emptyset} \mathbf{X}_A^{\mathcal{I}}$	Increments of \mathcal{J} between $\text{ht}(v_1)$ and $\text{ht}(v_{j^*+1})$
$\mathbf{X}_1^{\mathcal{J}}$	$\tilde{\theta} \mathbf{X}_1^{\mathcal{I}}$	Increments of \mathcal{J} between $\text{ht}(v_{j^*+1})$ and $\text{ht}(v_t)$
$\mathbf{X}_B^{\mathcal{J}}$	$\theta_{\rho(v_t - v_{j^*+1})} \Theta_{\emptyset} \mathbf{X}_B^{\mathcal{I}}$	Increments of \mathcal{J} between $\text{ht}(v_t)$ and $\text{ht}(v_{k^*+1})$
$\mathbf{X}_2^{\mathcal{J}}$	$\tilde{\theta} \mathbf{X}_2^{\mathcal{I}}$	Increments of \mathcal{J} above $\text{ht}(v_{k^*+1})$
$W_x^{\mathcal{J}}$	$\mathcal{F}(x + [-\frac{1}{2}, \frac{1}{2}]^2 \times [0, \text{ht}(v_1) - \frac{1}{2}])$	New faces added in step 9 of $\Psi_{x,t}$
$\theta_{\dagger} \mathbf{B}$	$\theta_{\dagger} \mathbf{C}_1 \cup \theta_{\dagger} \mathbf{C}_2 \cup \mathbf{F}$	Vertical translations of \mathbf{B} due to the deletion of walls \mathbf{W}
\mathbf{H}	$\bigcup \{f \in \mathcal{J} \setminus \mathcal{P}_x^{\mathcal{J}} : \rho(f) \in \mathcal{F}(\rho(\mathbf{W}))\}$	Faces added to “fill in” the rest of the interface

Refer to Figure 5 for a depiction of this splitting.

4.3.2. *The excess area of the map.* With the above decomposition in hand, note that the change in energy between \mathcal{I} and $\mathcal{J} = \Psi_{x,t}(\mathcal{I})$ is given by

$$(4.2) \quad m(\mathcal{I}; \mathcal{J}) = \sum_{z \in \mathbf{D}} m(\tilde{W}_z) + \sum_{j=1}^{j^*} m(\mathcal{X}_j) + \sum_{k=t}^{k^*} m(\mathcal{X}_k) \mathbf{1}\{t > j^*\} - |W_x^{\mathcal{J}}|.$$

The following inequalities regarding $m(\mathcal{I}; \mathcal{J})$ will be used repeatedly.

CLAIM 4.7. For every $\mathcal{I} \in \bar{\mathbf{I}}_{x,T,H}$, denoting by $\mathcal{J} = \Psi_{x,t}(\mathcal{I})$, we have

$$(4.3) \quad |W_x^{\mathcal{J}}| \leq \frac{2}{3} m(\tilde{\mathcal{W}}_{v_1} \cup \tilde{\mathcal{W}}_{y^{\dagger}})$$

and, in particular,

$$(4.4) \quad |W_x^{\mathcal{J}}| \leq 2m(\mathcal{I}; \mathcal{J}) \quad \text{and} \quad m(\mathbf{W}) \leq 3m(\mathcal{I}; \mathcal{J}).$$

We also have

$$(4.5) \quad j^* - 1 \leq 6m(\mathcal{I}; \mathcal{J}) \quad \text{and} \quad k^* - t \leq 6m(\mathcal{I}; \mathcal{J}).$$

Before getting to the proof of Claim 4.7, we need some simple geometric observations. Recall that for a face-set W , its height is given by $\text{ht}(W) = \max\{x_3 : (x_1, x_2, x_3) \in W\}$. Observe that $\text{ht}(W) = \text{ht}(W \cup \lceil W \rceil)$.

FACT 4.8. For every $j \geq 1$ and every index $y \in \mathcal{L}_{0,n}$,

$$j - 1 \leq \min_{W \in \Theta_{\ddagger}^{\mathcal{I} \setminus \mathcal{S}_x} \tilde{W}_y} [d_{\ddagger}(W \cup \lceil W \rceil, \mathcal{X}_j) + \text{ht}(W)],$$

where $d_{\ddagger}(A, B) = \min\{|x_3 - y_3| : x \in A, y \in B\}$. In particular, for every ω_1 and ω_2 ,

$$j - 1 \leq \mathfrak{D}_x(\tilde{W}_y, j, \omega_1, \omega_2) + m(\tilde{\mathfrak{F}}_y).$$

PROOF. The first claim follows from the triangle inequality and $\min\{x_3 : (x_1, x_2, x_3) \in \mathcal{X}_j\} \geq j - 1$. Then, for every ω_1, ω_2 ,

$$\begin{aligned} & \min_{W \in \Theta_{\ddagger}^{\mathcal{I} \setminus \mathcal{S}_x} W_y} d_{\ddagger}(W \cup \lceil W \rceil, \mathcal{X}_j) \\ &= \min_{W \in \Theta_{\ddagger}^{\mathcal{I} \setminus \mathcal{S}_x} \tilde{W}_y} d_{\ddagger}(W \cup \lceil W \rceil, \mathcal{X}_j \cup \theta_{\rho(x+\omega_1+\omega_2)} \mathcal{X}_j \cup \theta_{\rho(x+\omega_2)} \Theta_{\emptyset} \mathcal{X}_j) \\ &\leq \min_{W \in \Theta_{\ddagger}^{\mathcal{I} \setminus \mathcal{S}_x} \tilde{W}_y} d(W \cup \lceil W \rceil, \mathcal{X}_j \cup \theta_{\rho(x+\omega_1+\omega_2)} \mathcal{X}_j \cup \theta_{\rho(x+\omega_2)} \Theta_{\emptyset} \mathcal{X}_j) \\ &= \mathfrak{D}_x(\tilde{W}_y, j, \omega_1, \omega_2). \end{aligned}$$

The proof concludes from the observation that $\max_{W \in \Theta_{\ddagger}^{\mathcal{I} \setminus \mathcal{S}_x} \tilde{W}_y} \text{ht}(W) \leq m(\tilde{\mathfrak{F}}_y)$. \square

CLAIM 4.9. Let $\{W_z\}_{z \in \mathcal{L}_{0,n}}$ be the collection of walls corresponding to some interface \mathcal{I} . There exists some $z_0 \in \mathcal{L}_{0,n}$ such that every cut-point v of \mathcal{I} must belong to W_{z_0} (its four vertical bounding faces are in W_{z_0}).

Consequently, for an interface \mathcal{I} with pillar \mathcal{P}_x , the face set of \mathcal{S}_x consists of exactly one wall, together with at most one ceiling face projecting into $\rho(v_1)$.

PROOF. Suppose by way of contradiction that the interface \mathcal{I} has two cut-points w_1, w_2 with $0 < \text{ht}(w_1) < \text{ht}(w_2)$ such that the walls containing the bounding faces of w_1 and w_2 , namely, $W_{\rho(w_1)}$ and $W_{\rho(w_2)}$ are distinct. If both $W_{\rho(w_1)}, W_{\rho(w_2)}$ are standard, then $W_{\rho(w_i)}$ intersect each of $\mathcal{L}_{\frac{1}{2}}, \dots, \mathcal{L}_{\text{ht}(w_i)}$ in at least one cell (i.e., if $\mathcal{I}_{W_{\rho(w_i)}}$ is the interface whose only wall is $W_{\rho(w_i)}$, then $\sigma(\mathcal{I}_{W_{\rho(w_i)}})$ intersects every height between $\frac{1}{2}$ and $\text{ht}(w_i)$ in at least one cell). As a consequence, no height between $\frac{1}{2}$ and $\text{ht}(w_1)$ can be a cut-height of \mathcal{I} . If at least one of $W_{\rho(w_i)}$ is not standard, then it is identified with a ceiling $\mathcal{C}_i \subset \lceil W \rceil$ for some other wall $W \ni W_{\rho(w_i)}$. Call the tallest of those ceilings \mathcal{C} with height $\text{ht}(\mathcal{C})$. By definition of ceiling faces, we have the following observation.

OBSERVATION 4.10. Every cell w sharing a projection with a face $f \in \mathcal{C}$ and having $\text{ht}(w) \leq \text{ht}(\mathcal{C})$ is in $\sigma(\mathcal{I})$.

At the same time, since $W_{\rho(w_i)} \in W$ and $W_{\rho(w_i)} \neq W$, the ceiling \mathcal{C} has at least eight faces. Therefore, there are no cut-heights below $\text{ht}(\mathcal{C})$, yielding a contradiction if $\text{ht}(\mathcal{C}) \geq \text{ht}(w_1)$. If $\text{ht}(\mathcal{C}) < \text{ht}(w_1)$, $W_{\rho(w_1)} \cap \mathcal{L}_h \neq \emptyset$ and $W_{\rho(w_2)} \cap \mathcal{L}_h \neq \emptyset$ for every $\text{ht}(\mathcal{C}) < h \leq \text{ht}(w_1)$, also yielding a contradiction.

To see the conclusion for the spine of a pillar \mathcal{P}_x , take $\mathcal{I}_{\mathcal{P}}$ to be the interface which is at \mathcal{L}_0 except for the faces of \mathcal{P}_x . Applying the first part of the claim to $\mathcal{I}_{\mathcal{P}}$, we see that all cut-points of \mathcal{S}_x are in the same wall W , and, by definitions of cut-points and the observation above, all other faces of \mathcal{S}_x must in W , except possibly one ceiling face projecting into $\rho(v_1)$. \square

COROLLARY 4.11. *The walls whose standardizations are $\tilde{\mathfrak{W}}_{v_1} \cup \tilde{\mathfrak{W}}_{y^\dagger}$ intersect each of $\mathcal{L}_{\frac{1}{2}}, \dots, \mathcal{L}_{\text{ht}(v_1)-1}$ in at least six faces (i.e., the corresponding interface intersects every such height in at least two cells).*

PROOF. By definition, the walls \mathcal{O}_{v_1} (defined s.t. $\theta_{\text{ST}}\mathcal{O}_{v_1} = \tilde{\mathfrak{W}}_{v_1}$) intersected every height between $h^\dagger + 1$ and $\text{ht}(v_1) - 1$ in at least two cells. Now, consider heights between 0 and h^\dagger .

Since y^\dagger and v_1 are both in \mathcal{P}_x , by Observation 3.2 there must exist a wall W such that y^\dagger, v_1 are both interior to W . Since $\tilde{W}_{\rho(v_1)} \neq \tilde{W}_{y^\dagger}$, there are inner-most ceilings of \mathcal{C}_{y^\dagger} and \mathcal{C}_{v_1} in $[W]$ nesting those respective walls. As such, by the observation above, every height between below $\text{ht}(\mathcal{C}_{v_1}) \vee \text{ht}(\mathcal{C}_{y^\dagger})$ is intersected by at least eight cells. Finally, since the walls whose standardizations are \tilde{W}_{v_1} and \tilde{W}_{y^\dagger} attain height h^\dagger , every height between $\text{ht}(\mathcal{C}_{v_1}) \vee \text{ht}(\mathcal{C}_{y^\dagger})$ and $\text{ht}(h^\dagger)$ is intersected by at least one cell by each of those walls. \square

PROOF OF CLAIM 4.7. By Corollary 4.11 there are no cut-points in $\tilde{\mathfrak{W}}_{v_1} \cup \tilde{\mathfrak{W}}_{y^\dagger}$, and, therefore,

$$|W_x^{\mathcal{J}}| = 4 \left(\text{ht}(v_1) - \frac{1}{2} \right) \leq \frac{2}{3} m(\tilde{\mathfrak{W}}_{v_1} \cup \tilde{\mathfrak{W}}_{y^\dagger}).$$

Hence, $|W_x^{\mathcal{J}}| \leq \frac{2}{3} m(\mathbf{W})$, and (4.4) follows from the fact $m(\mathcal{I}; \mathcal{J}) \geq m(\mathbf{W}) - |W_x^{\mathcal{J}}|$ by (4.2).

Let us now turn to proving the comparisons with $j^* - 1$ and $k^* - t$, namely, (4.5). Suppose $j^* > 1$, as otherwise the inequality is trivial. Since \mathcal{X}_{j^*} was deleted, it was due to one of criteria (A1) or (A2) or (A3):

(A1) In this case, $m(\mathcal{X}_{j^*}) \geq j^* - 1$.

(A2) In this case, $\mathfrak{D}_x(\tilde{W}_y, j^*, -v_{5_j+1}, 0) \leq m(\tilde{W}_y)$ for some $y \in \mathbf{Y}$. By Fact 4.8 applied to j^* ,

$$j^* - 1 \leq \mathfrak{D}_x(\tilde{W}_y, j^*, -v_{5_j+1}, 0) + m(\tilde{\mathfrak{F}}_y) \leq 2m(\tilde{\mathfrak{F}}_y).$$

(A3) In this case, $\mathfrak{D}_x(\tilde{W}_y, j^*, -v_{5_j+1}, 0) \leq (j - 1)/2$ for $y = y_A^*$. By Fact 4.8 applied to j^* and y_A^* ,

$$j^* - 1 \leq \mathfrak{D}_x(\tilde{W}_{y_A^*}, j^*, -v_{5_j+1}, 0) + m(\tilde{\mathfrak{F}}_{y_A^*}) \leq (j^* - 1)/2 + m(\tilde{\mathfrak{F}}_{y_A^*}),$$

so

$$j^* - 1 \leq 2m(\tilde{\mathfrak{F}}_{y_A^*}).$$

In any of these above cases, we have $j^* - 1 \leq 2m(\mathbf{W}) \vee m(\mathcal{X}_{j^*}) \leq 6m(\mathcal{I}; \mathcal{J})$ by (4.4).

If $j^* \geq t$, then we are done. Otherwise, since \mathcal{X}_{k^*} was deleted, it was due to either (B1) or (B2) or (B3):

(B1) In this case, $m(\mathcal{X}_{k^*}) \geq k^* - t$.

(B2) In this case, $\mathcal{D}_x(\tilde{W}_y, k^*, -v_{5k+1}, v_t - v_{j^*+1}) \leq m(\tilde{W}_y)$ for some $y \in \mathbf{Y}$. By Fact 4.8,

$$k^* - t \leq k^* - 1 \leq \mathcal{D}_x(\tilde{W}_y, k^*, -v_{5k+1}, v_t - v_{j^*+1}) + m(\tilde{\mathcal{F}}_y) \leq 2m(\tilde{\mathcal{F}}_y).$$

(B3) In this case, $\mathcal{D}_x(\tilde{W}_y, k^*, -v_{5k+1}, v_t - v_{j^*+1}) \leq (k - t)/2$ for $y = y_B^*$. By Fact 4.8,

$$k^* - t \leq \mathcal{D}_x(\tilde{W}_{y_B^*}, k^*, -v_{5k+1}, v_t - v_{j^*+1}) + m(\tilde{\mathcal{F}}_{y_B^*}) \leq (k^* - t)/2 + m(\tilde{\mathcal{F}}_{y_B^*}),$$

so

$$k^* - t \leq 2m(\tilde{\mathcal{F}}_{y_B^*}).$$

In any of these cases, we have $k^* - t \leq 6m(\mathcal{I}; \mathcal{J})$ by (4.4). \square

4.4. *Proof part 1: Interface weights.* In this section we show that the map $\Psi_{x,t}$ amplifies the weights of interfaces by something exponential in the excess area $m(\mathcal{I}; \Psi_{x,t}(\mathcal{I}))$.

As in the preceding works [14, 16], the difficulty here is ensuring that the cumulative effect of the perturbative terms \mathbf{g} in Theorem 2.4 (capturing interactions between different parts of the interface through subcritical droplets) is comparable to $m(\mathcal{I}; \Psi_{x,t}(\mathcal{I}))$. As described in Section 4.2, this is particularly complicated here, as we cannot reduce the interactions to only their horizontal or to only their vertical parts.

PROPOSITION 4.12. *There exists $C > 0$ and β_0 such that for every $\beta > \beta_0$ the following holds. For every $x \in \mathcal{L}_{0,n}$ and $0 \leq T < H$, for every t and every $\mathcal{I} \in \bar{\mathbf{I}}_{x,T,H}$,*

$$\left| \log \frac{\mu_n^\mp(\mathcal{I})}{\mu_n^\mp(\Psi_{x,t}(\mathcal{I}))} + \beta m(\mathcal{I}; \Psi_{x,t}(\mathcal{I})) \right| \leq C m(\mathcal{I}; \Psi_{x,t}(\mathcal{I})).$$

We first prove a series of preliminary estimates to which we will reduce Proposition 4.12 by pairing faces together according to the decomposition of \mathcal{I} and \mathcal{J} from Section 4.3.1.

CLAIM 4.13. *There exists \bar{C} such that, for every $\mathcal{I} \in \bar{\mathbf{I}}_{x,T,H}$,*

$$\sum_{f \in \mathcal{F}(\mathbb{Z}^3)} \sum_{g \in \mathbf{WUH}} e^{-\bar{c}d(f,g)} \leq \bar{C} m(\mathbf{W}).$$

PROOF. By summing the exponential tail, there exists $C > 0$ such that

$$\begin{aligned} \sum_{g \in \mathbf{WUH}} \sum_{f \in \mathcal{F}(\mathbb{Z}^3)} e^{-\bar{c}d(f,g)} &\leq \sum_{z \in \mathbf{D}} \left(\sum_{g \in \tilde{W}_z} C + \sum_{\substack{g' \in \mathcal{J} \setminus \mathcal{S}_x^{\mathcal{J}} \\ \rho(g') \in \mathcal{F}(\rho(\tilde{W}_z))}} C \right) \\ &\leq \sum_{z \in \mathbf{D}} (C|\tilde{W}_z| + C|\mathcal{F}(\rho(\tilde{W}_z))|), \end{aligned}$$

which by (2.3) is, at most, $3C \sum_{z \in \mathbf{D}} m(\tilde{W}_z) \leq 3Cm(\mathbf{W})$. \square

CLAIM 4.14. *There exists \bar{C} such that, for every $\mathcal{I} \in \bar{\mathbf{I}}_{x,T,H}$,*

$$\sum_{f \in \mathcal{F}(\mathbb{Z}^3)} \sum_{\iota \in \{A,B\}} \sum_{g \in \mathbf{X}_\iota^{\mathcal{I}} \cup \mathbf{X}_\iota^{\mathcal{J}}} e^{-\bar{c}d(f,g)} \leq \bar{C} m(\mathcal{I}; \mathcal{J}),$$

PROOF. Summing over all $f \in \mathcal{F}(\mathbb{Z}^3)$, there exists $C > 0$ such that

$$\sum_{\iota \in \{A, B\}} \sum_{g \in \mathbf{X}_\iota^{\mathcal{I}} \cup \mathbf{X}_\iota^{\mathcal{J}}} \sum_{f \in \mathcal{F}(\mathbb{Z}^3)} e^{-\bar{c}d(f, g)} \leq C \left(\sum_{1 \leq j \leq j^*} [\mathfrak{m}(\mathcal{X}_j) + 8(\text{ht}(v_{j+1}) - \text{ht}(v_j) + 1)] \right. \\ \left. + \sum_{t \leq k \leq k^*} [\mathfrak{m}(\mathcal{X}_k) + 8(\text{ht}(v_{k+1}) - \text{ht}(v_k) + 1)] \right),$$

where, in the first inequality, the factor of eight accounts for four faces from the \mathcal{X}_i (equation (3.1)) and four from $\Theta_\emptyset \mathcal{X}_i$ for each height. The telescopic sums give $8(\text{ht}(v_{j^*+1}) - \text{ht}(v_1))$ and $\text{ht}(v_{j^*+1}) - \text{ht}(v_1) \leq j^* + \frac{1}{2} \sum_{j=1}^{j^*} \mathfrak{m}(\mathcal{X}_j)$ since for each height in $\text{ht}(v_1), \dots, \text{ht}(v_{j^*+1})$ that wasn't a cutpoint, we added an excess area of at least two. Accounting for an extra additive j^* as well as the similar telescoping for $t \leq k \leq k^*$, we see this is, at most,

$$16C(j^* - 1) + 16C(k^* - t) + 5C\mathfrak{m}(\mathcal{I}; \mathcal{J}) + 32C,$$

which, by (4.5) of Claim 4.7, is in turn, at most, $\bar{C}\mathfrak{m}(\mathcal{I}; \mathcal{J})$ for some other constant $\bar{C} > 0$. \square

CLAIM 4.15. *There exists \bar{C} such that, for every $\mathcal{I} \in \bar{\mathbf{I}}_{x, T, H}$,*

$$\sum_{f \in \mathcal{F}(\mathbb{Z}^3)} \sum_{g \in W_x^{\mathcal{J}}} e^{-\bar{c}d(f, g)} \leq \bar{C}\mathfrak{m}(\mathbf{W}).$$

PROOF. Summing over all $f \in \mathcal{F}(\mathbb{Z}^3)$, there exists $C > 0$ such that

$$\sum_{g \in W_x^{\mathcal{J}}} \sum_{f \in \mathcal{F}(\mathbb{Z}^3)} e^{-\bar{c}d(f, g)} \leq C|W_x^{\mathcal{J}}|,$$

which is, at most, $\frac{2}{3}C\mathfrak{m}(\mathbf{W})$ by (4.3). \square

The following series of lemmas bounds the interactions (through the subcritical droplets, via $|\mathbf{g}(f, \mathcal{I}) - \mathbf{g}(f', \mathcal{J})|$) between different subsets of \mathcal{I} and \mathcal{J} . The first of these concerns interactions between horizontal shifts of $\mathbf{X}_1^{\mathcal{I}}$ and $\mathbf{X}_2^{\mathcal{J}}$ with \mathbf{F} .

LEMMA 4.16. *There exists \bar{C} such that, for every $\mathcal{I} \in \bar{\mathbf{I}}_{x, T, H}$,*

$$\sum_{f \in \mathbf{X}_1^{\mathcal{I}} \cup \mathbf{X}_2^{\mathcal{I}}} \sum_{g \in \mathcal{L}_0} e^{-\bar{c}d(f, g)} \leq \bar{C},$$

PROOF. Noticing that for every i , $d(\mathcal{X}_i, \mathcal{L}_0) \geq i$, there exists $C > 0$ such that

$$\sum_{g \in \mathbf{X}_1^{\mathcal{I}} \cup \mathbf{X}_2^{\mathcal{I}}} \sum_{f \in \mathcal{L}_0} e^{-\bar{c}d(f, g)} \leq C \sum_{j^* < i < t} |\mathcal{F}(\mathcal{X}_i)| e^{-\bar{c}i} + C \sum_{i > k^*} |\mathcal{F}(\mathcal{X}_i)| e^{-\bar{c}i} \\ \leq C \sum_{i \in (j^*, t) \cup (k^*, \infty)} [4 + 3\mathfrak{m}(\mathcal{X}_i)] e^{-\bar{c}i}.$$

In turn, using the fact that $\mathfrak{m}(\mathcal{X}_i) \leq i - 1$ for $i \in (j^*, t)$ and $i > k^*$ (criterion (A1)), this is, at most,

$$\sum_{i > j^*} C[1 + 3i] e^{-\bar{c}i} \leq \bar{C} e^{-\bar{c}j^*}.$$

\square

The next lemma helps control horizontal interactions induced by vertical shifts of walls and ceilings in \mathbf{B} :

LEMMA 4.17. *There exists \bar{C} such that, for every $\mathcal{I} \in \bar{\mathbf{I}}_{x,T,H}$,*

$$\sum_{f \in \mathbf{B}} \sum_{u \in \rho(\cup_{z \in \mathbf{D}} \tilde{W}_z)} e^{-\bar{c}d(\rho(f),u)} \leq \bar{C} \mathfrak{m}(\mathbf{W}).$$

PROOF. By definition of $N_\rho(u)$, we have

$$\sum_{f \in \mathbf{B}} \sum_{u \in \rho(\cup_{z \in \mathbf{D}} \tilde{W}_z)} e^{-\bar{c}d(\rho(f),u)} = \sum_{u' \in \rho(\cup_{z \in \mathbf{D}} \tilde{W}_z)^c} \sum_{u \in \rho(\cup_{z \in \mathbf{D}} \tilde{W}_z)} N_\rho(u') e^{-\bar{c}d(u,u')}.$$

Since \mathbf{D} is closed under closeness of walls, for every such u, u' we have $N_\rho(u') \leq |u - u'|^2 + 1$. Thus, there exists a $C > 0$ such that the right-hand side above is in turn, at most,

$$C \sum_{z \in \mathbf{D}} |\mathcal{E}(\rho(\tilde{W}_z))| + |\mathcal{F}(\rho(\tilde{W}_z))| \leq C \sum_{z \in \mathbf{D}} \mathfrak{m}(\tilde{W}_z). \quad \square$$

The following lemma controls the vertical interactions between the shift in $\mathbf{X}_2^{\mathcal{I}}$ relative to faces in $\mathbf{X}_1^{\mathcal{I}}$. In this way, $\vec{\theta} \mathbf{X}_1^{\mathcal{I}} = \mathbf{X}_1^{\mathcal{J}}$ and $\vec{\theta} \mathbf{X}_2^{\mathcal{I}} = \mathbf{X}_2^{\mathcal{J}}$.

LEMMA 4.18. *There exists \bar{C} such that, for every $\mathcal{I} \in \bar{\mathbf{I}}_{x,T,H}$,*

$$\sum_{f \in \mathbf{X}_1^{\mathcal{I}}} \sum_{g \in \mathbf{X}_2^{\mathcal{I}}} e^{-\bar{c}d(f,g)} + e^{-\bar{c}d(\vec{\theta} f, \vec{\theta} g)} \leq \bar{C}.$$

PROOF. Assume $t > j^*$, as otherwise $\mathbf{X}_1^{\mathcal{I}}$ is empty. We can bound the left-hand side above by

$$\sum_{f \in \mathbf{X}_1^{\mathcal{I}}} \sum_{k > k^*} \sum_{g \in \mathcal{R}_k} e^{-\bar{c}d(f,g)} + e^{-\bar{c}d(\vec{\theta} f, \vec{\theta} g)} \leq C \sum_{k > k^*} |\mathcal{F}(\mathcal{R}_k)| e^{-\bar{c}(k-t)}$$

for some $C > 0$. By criterion (B1), for every $k > k^*$, $\mathfrak{m}(\mathcal{R}_k) \leq k - t$ and this is, at most,

$$C \sum_{k > k^*} [4 + 3(k - t)] e^{-\bar{c}(k-t)} \leq \bar{C}. \quad \square$$

The remaining two lemmas are more involved, as they control the interactions between faces in \mathbf{C}_1 and \mathbf{C}_2 (which may shift vertically) with the horizontal shifts of $\mathbf{X}_1^{\mathcal{I}} \cup \mathbf{X}_2^{\mathcal{I}}$. Such terms were not considered in previous works, and they cannot be reduced to either two-dimensional bound via projections or to a one-dimensional bound via height differences. As explained in Section 4.2, these bounds are very sensitive to the particular choices for the deletion criteria, particularly (A1), (B1) and (A3), (B3).

Recall that for all $f \in \mathbf{B}$, $\theta_\dagger f$ was defined as the vertical shift of f induced by removal of the walls in \mathbf{W} per Lemma 2.11. With this in mind, note that $\theta_\dagger f \in \Theta_\dagger^{\mathcal{I} \setminus \mathcal{S}_x} W$ for every $f \in W \cup [W]$.

LEMMA 4.19. *There exists \bar{C} such that, for every $\mathcal{I} \in \bar{\mathbf{I}}_{x,T,H}$,*

$$\sum_{f \in \mathbf{C}_1} \sum_{g \in \mathbf{X}_1^{\mathcal{I}} \cup \mathbf{X}_2^{\mathcal{I}}} e^{-\bar{c}d(f,g)} + e^{-\bar{c}d(\theta_\dagger f, \vec{\theta} g)} \leq \bar{C} \mathfrak{m}(\mathcal{I}; \mathcal{J}).$$

PROOF. Begin by considering $g \in \mathbf{X}_1^{\mathcal{I}}$ (assuming $t > j^*$, as otherwise $\mathbf{X}_1^{\mathcal{I}}$ is empty). There exists $C > 0$ such that

$$\begin{aligned} & \sum_{f \in \mathbf{C}_1} \sum_{g \in \mathbf{X}_1^{\mathcal{I}}} e^{-\bar{c}d(f,g)} + e^{-\bar{c}d(\theta_{\downarrow} f, \vec{\theta} g)} \\ & \leq C \sum_{j^* < i < t} \sum_{y \in \mathbf{Y}} |\mathcal{F}(\mathcal{X}_i)| (e^{-\bar{c}d(\lceil \tilde{W}_y \rceil, \mathcal{X}_i)} + e^{-\bar{c}d(\theta_{\downarrow}(\lceil \tilde{W}_y \rceil), \vec{\theta} \mathcal{X}_i)}) \\ & \leq 2C \sum_{j^* < i < t} \sum_{y \in \mathbf{Y}} |\mathcal{F}(\mathcal{X}_i)| e^{-\bar{c}\mathcal{D}_x(\tilde{W}_y, i, v_{j^*+1}, 0)}, \end{aligned}$$

since $\Theta_{\downarrow}^{\mathcal{I} \setminus \mathcal{S}_x} \tilde{W}_y$ includes $\lceil \tilde{W}_y \rceil$ and $\theta_{\downarrow}(\lceil \tilde{W}_y \rceil)$, and $\vec{\theta} \mathcal{X}_i$ is exactly $\theta_{\rho(x-v_{j^*+1})} \mathcal{X}_i$ which was one of the horizontal translates considered in the definition of \mathcal{D}_x . Further, since \mathcal{X}_i was not deleted, by (A1) and (A3)

$$m(\mathcal{X}_i) < (i - 1) \quad \text{and} \quad \mathcal{D}_x(\tilde{W}_y, i, -v_{j^*+1}, 0) > (i - 1)/2,$$

so that for every $j^* < i < t$ and every $y \in \mathbf{Y}$,

$$m(\mathcal{X}_i) < 2\mathcal{D}_x(\tilde{W}_y, i, -v_{j^*+1}, 0).$$

Using $|\mathcal{F}(\mathcal{X}_i)| \leq 3m(\mathcal{X}_i) + 4$, we have that the above sum is, at most,

$$8C \sum_{j^* < i < t} \sum_{y \in \mathbf{Y}} \mathcal{D}_x(\tilde{W}_y, i, -v_{j^*+1}, 0) e^{-\bar{c}\mathcal{D}_x(\tilde{W}_y, i, -v_{j^*+1}, 0)}.$$

Observe first that for some $C > 0$,

$$(4.6) \quad \sum_{y \in \mathbf{Y}} \mathcal{D}_x(\tilde{W}_y, i, -v_{j^*+1}, 0) e^{-\bar{c}\mathcal{D}_x(\tilde{W}_y, i, -v_{j^*+1}, 0)} < C \quad \text{for each } j^* < i < t.$$

Indeed this follows by writing

$$\begin{aligned} \mathcal{D}_x(\tilde{W}_y, i, -v_{j^*+1}, 0) e^{-\bar{c}\mathcal{D}_x(\tilde{W}_y, i, -v_{j^*+1}, 0)} & \leq d(\Theta_{\downarrow}^{\mathcal{I} \setminus \mathcal{S}_x} \tilde{W}_y, \mathcal{X}_i) e^{-\bar{c}d(\Theta_{\downarrow}^{\mathcal{I} \setminus \mathcal{S}_x} \tilde{W}_y, \mathcal{X}_i)} \\ & \quad + d(\Theta_{\downarrow}^{\mathcal{I} \setminus \mathcal{S}_x} \tilde{W}_y, \vec{\theta} \mathcal{X}_i) e^{-\bar{c}d(\Theta_{\downarrow}^{\mathcal{I} \setminus \mathcal{S}_x} \tilde{W}_y, \vec{\theta} \mathcal{X}_i)} \\ & \quad + d(\Theta_{\downarrow}^{\mathcal{I} \setminus \mathcal{S}_x} \tilde{W}_y, \Theta_{\emptyset} \mathcal{X}_i) e^{-\bar{c}d(\Theta_{\downarrow}^{\mathcal{I} \setminus \mathcal{S}_x} \tilde{W}_y, \Theta_{\emptyset} \mathcal{X}_i)} \end{aligned}$$

and noticing that if $y \neq y'$, then $\Theta_{\downarrow}^{\mathcal{I} \setminus \mathcal{S}_x} \tilde{W}_y \cap \Theta_{\downarrow}^{\mathcal{I} \setminus \mathcal{S}_x} \tilde{W}_{y'} = \emptyset$, so that after summing over $y \in \mathbf{Y}$, each term on the right-hand side contributes a constant. (However, we cannot afford an overall bound of order $t - j^*$ which may not be comparable to $m(\mathcal{I}; \mathcal{J})$.)

Thus, we only use the above bound to deal with increments whose height is, at most, the maximal height of a ceiling of \tilde{W}_y or one of its possible vertical shifts. Namely, let

$$\bar{j} = \min \left\{ j : \text{ht}(v_{j+1}) > \max_{y \in \mathbf{Y}} \max_{f \in \Theta_{\downarrow}^{\mathcal{I} \setminus \mathcal{S}_x} \tilde{W}_y} \text{ht}(f) \right\},$$

and denote by \bar{y} the index of the wall attaining this height. Then, using (4.6),

$$\sum_{j^* < i \leq \bar{j}} \sum_{y \in \mathbf{Y}} \mathcal{D}_x(\tilde{W}_y, i, -v_{j^*+1}, 0) e^{-\bar{c}\mathcal{D}_x(\tilde{W}_y, i, -v_{j^*+1}, 0)} \leq C \bar{j} \leq Cm(\tilde{\mathfrak{F}}_{\bar{y}}).$$

For the remaining increments, for every $y \in \mathbf{Y}$, let

$$d_{\rho}(y, i) := d(\rho(\lceil \tilde{W}_y \rceil), \rho(\mathcal{X}_i \cup \theta_{\rho(x-v_{j^*+1})} \mathcal{X}_i \cup \Theta_{\emptyset} \mathcal{X}_i)),$$

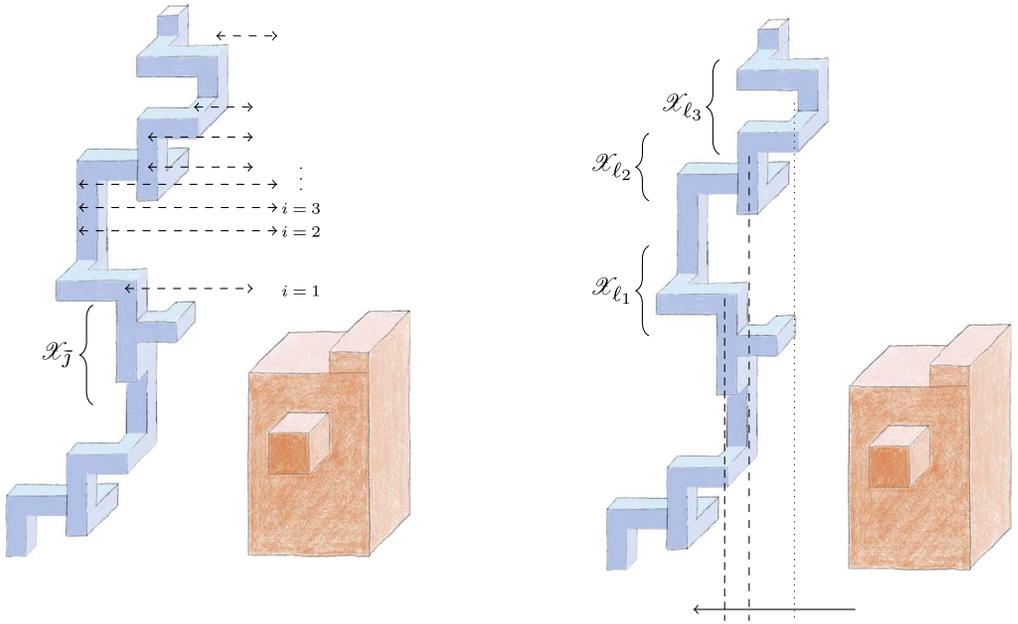


FIG. 6. *Left: The index \bar{j} is determined by the height of the wall $W_{\bar{y}}$ (orange). The distances $d_\rho(\bar{y}, \bar{j} + i)$ are recorded for $i = 1, 2, \dots$. Right: The record times of this distance then form the increment indices $\ell_1 = \bar{j} + 1, \ell_2, \ell_3$.*

and let $\ell_1 = \bar{j} + 1 < \ell_2 < \dots < \ell_r$ be the record times of the function $d_\rho(y, \cdot)$, that is,

$$d_\rho(y, \ell_j) < d_\rho(y, i) \quad \text{for all } \bar{j} < i < \ell_j,$$

$$d_\rho(y, \ell_r) = \min\{d_\rho(y, i) : \bar{j} < i < t\}.$$

(See Figure 6.) Let $\ell_{r+1} := \infty$, and observe that, for every $j = 1, \dots, r$ and every $\ell_j \leq i < \ell_{j+1}$,

$$\mathfrak{D}_x(\tilde{W}_y, i, -v_{j^*+1}, 0) \geq \sqrt{d_\rho(y, \ell_j)^2 + (\text{ht}(v_i) - \text{ht}(v_{\bar{j}+1}))^2} \geq \frac{d_\rho(y, \ell_j) + i - \ell_j}{\sqrt{2}},$$

using the definition of \bar{j} and that it satisfies $\bar{j} < \ell_j$. In particular, there exists $C, C' > 0$ such that

$$\sum_{\ell_j \leq i < \ell_{j+1}} \mathfrak{D}_x(\tilde{W}_y, i, -v_{j^*+1}, 0) e^{-\bar{c}\mathfrak{D}_x(\tilde{W}_y, i, -v_{j^*+1}, 0)} \leq \sum_{\ell_j \leq i < \ell_{j+1}} C e^{-\frac{\bar{c}}{\sqrt{2}}(d_\rho(y, \ell_j) + i - \ell_j)}$$

$$\leq C' e^{-\frac{\bar{c}}{\sqrt{2}}d_\rho(y, \ell_j)}.$$

Summing over $1 \leq j \leq r$ and noticing that $r \leq d_\rho(y, \ell_1)$,

$$\sum_{j=1}^r e^{-\frac{\bar{c}}{\sqrt{2}}d_\rho(y, \ell_j)} = \sum_{j=1}^r e^{-\frac{\bar{c}}{\sqrt{2}}(d_\rho(y, \ell_1) - j)} \leq \bar{C}.$$

Summing over $y \in \mathbf{Y}$, this is, at most, $\sum_{y \in \mathbf{Y}} \bar{C} \leq \bar{C} \sum_{y \in \mathbf{Y}} m(\tilde{W}_y)$.

The treatment of $g \in \mathbf{X}_2^I$ is identical to the above argument, with the sole difference being the values of the horizontal shifts in the definition of \mathfrak{D}_x . \square

LEMMA 4.20. *There exists \bar{C} such that, for every $\mathcal{I} \in \bar{\mathbf{I}}_{x,T,H}$,*

$$\sum_{f \in \mathbf{C}_2} \sum_{g \in \mathbf{X}_1^I \cup \mathbf{X}_2^I} e^{-\bar{c}d(f,g)} + e^{-\bar{c}d(\theta_\dagger f, \vec{\theta} g)} \leq \bar{C}.$$

PROOF. Begin by considering $g \in \mathbf{X}_1^{\mathcal{I}}$. If $y \notin \mathbf{Y}$ is such that $f \in \tilde{W}_y \cup [\tilde{W}_y]$ and $g \in \mathcal{X}_i$ for some $j^* < i < t$, then (A3) implies that

$$(i - 1)/2 < \mathfrak{D}_x(\tilde{W}_y, i, -v_{j^*+1}, 0) \leq d(f, g) \wedge d(\theta_{\dagger} f, \vec{\theta} g),$$

since $\theta_{\dagger} f \in \Theta_{\dagger}^{\mathcal{I} \setminus \mathcal{S}_x} \tilde{W}_y$, and $\vec{\theta} g = \theta_{\rho(x-v_{j^*+1})} g$. Further, since $i > j^*$, by criterion (A1) $m(\mathcal{X}_i) < i - 1$, so

$$\begin{aligned} \sum_{f \in \mathbf{C}_2} \sum_{g \in \mathbf{X}_2^{\mathcal{I}}} e^{-\bar{c}d(f,g)} + e^{-\bar{c}d(\theta_{\dagger} f, \vec{\theta} g)} &\leq 2C \sum_{j^* < i < t} [3m(\mathcal{X}_i) + 4] e^{-\bar{c}(i-1)/2} \\ &\leq 8C \sum_{j^* < i < t} i e^{-\bar{c}(i-1)/2} \leq \frac{\bar{C}}{2} \end{aligned}$$

for some \bar{C} . The treatment of $g \in \mathbf{X}_2^{\mathcal{I}}$ is identical to the above, with the only difference being in the horizontal shift in the definition of \mathfrak{D}_x . \square

PROOF OF PROPOSITION 4.12. By Theorem 2.4 it suffices to show that there exists $C > 0$ such that, for every $\mathcal{I} \in \bar{\mathbf{I}}_{x,T,H}$, if $\mathcal{J} = \Psi_{x,t}(\mathcal{I})$,

$$\left| \sum_{f \in \mathcal{I}} \mathbf{g}(f, \mathcal{I}) - \sum_{f' \in \mathcal{J}} \mathbf{g}(f', \mathcal{J}) \right| \leq C m(\mathcal{I}; \mathcal{J}).$$

Using the partition of the faces of \mathcal{I}, \mathcal{J} in Section 4.3.1, we can expand

$$\begin{aligned} (4.7) \quad &\left| \sum_{f \in \mathcal{I}} \mathbf{g}(f, \mathcal{I}) - \sum_{f' \in \mathcal{J}} \mathbf{g}(f', \mathcal{J}) \right| \\ &\leq \sum_{f \in \mathbf{X}_A^{\mathcal{I}} \cup \mathbf{X}_B^{\mathcal{I}}} |\mathbf{g}(f, \mathcal{I})| + \sum_{f \in \mathbf{W}} |\mathbf{g}(f, \mathcal{I})| \\ (4.8) \quad &+ \sum_{f' \in \mathbf{X}_A^{\mathcal{J}} \cup \mathbf{X}_B^{\mathcal{J}}} |\mathbf{g}(f', \mathcal{J})| + \sum_{f' \in \mathbf{H}} |\mathbf{g}(f', \mathcal{J})| + \sum_{f' \in \mathbf{W}_x^{\mathcal{J}}} |\mathbf{g}(f', \mathcal{J})| \\ (4.9) \quad &+ \sum_{f \in \mathbf{X}_1^{\mathcal{I}} \cup \mathbf{X}_2^{\mathcal{I}}} |\mathbf{g}(f, \mathcal{I}) - \mathbf{g}(\vec{\theta} f, \mathcal{J})| + \sum_{f \in \mathbf{B}} |\mathbf{g}(f, \mathcal{I}) - \mathbf{g}(\theta_{\dagger} f, \mathcal{J})|. \end{aligned}$$

We show that each term on the right-hand side is comparable to $m(\mathcal{I}; \mathcal{J})$. By (2.1) the first sum satisfies

$$\sum_{f \in \mathbf{X}_A^{\mathcal{I}} \cup \mathbf{X}_B^{\mathcal{I}}} |\mathbf{g}(f, \mathcal{I})| \leq \bar{K} |\mathbf{X}_A^{\mathcal{I}} \cup \mathbf{X}_B^{\mathcal{I}}|$$

which is, at most, $\bar{C} \bar{K} m(\mathcal{I}; \mathcal{J})$ by Claim 4.14. The same holds for the first sum in line (4.8) by Claim 4.14. By Claim 4.13 the second sums in (4.7) and (4.8) are bounded in the same way by $\bar{C} \bar{K} m(\mathbf{W})$, which is in turn, at most, $3\bar{C} \bar{K} m(\mathcal{I}; \mathcal{J})$ by (4.4). The third sum in line (4.8) is bounded in this way by $\bar{C} \bar{K} m(\mathbf{W})$ via Claim 4.15, and this is in turn, at most, $3\bar{C} \bar{K} m(\mathcal{I}; \mathcal{J})$ by (4.4).

It remains to consider the two sums in line (4.9), which by (2.2), satisfy

$$\begin{aligned} \sum_{f \in \mathbf{X}_1^{\mathcal{I}} \cup \mathbf{X}_2^{\mathcal{I}}} |\mathbf{g}(f, \mathcal{I}) - \mathbf{g}(\vec{\theta} f, \mathcal{J})| &\leq \sum_{f \in \mathbf{X}_1^{\mathcal{I}} \cup \mathbf{X}_2^{\mathcal{I}}} \bar{K} e^{-\bar{c}r(f, \mathcal{I}; \vec{\theta} f, \mathcal{J})} \quad \text{and} \\ \sum_{f \in \mathbf{B}} |\mathbf{g}(f, \mathcal{I}) - \mathbf{g}(\theta_{\dagger} f, \mathcal{J})| &\leq \sum_{f \in \mathbf{B}} \bar{K} e^{-\bar{c}r(f, \mathcal{I}; \theta_{\dagger} f, \mathcal{J})}. \end{aligned}$$

To evaluate the radius \mathbf{r} , consider the right-hand sides according to the face g attaining $\mathbf{r}(f, \mathcal{I}; f', \mathcal{J})$:

(i) If $g \in \mathbf{X}_A^{\mathcal{I}} \cup \mathbf{X}_B^{\mathcal{I}} \cup \mathbf{X}_A^{\mathcal{J}} \cup \mathbf{X}_B^{\mathcal{J}}$, both these sums are, at most,

$$\sum_{g \in \mathbf{X}_A^{\mathcal{J}} \cup \mathbf{X}_B^{\mathcal{I}} \cup \mathbf{X}_A^{\mathcal{J}} \cup \mathbf{X}_B^{\mathcal{J}}} \bar{K} \left[\sum_{f \in \mathbf{X}_1^{\mathcal{I}} \cup \mathbf{X}_2^{\mathcal{I}}} (e^{-\bar{c}d(f,g)} + e^{-\bar{c}d(\vec{\theta} f, g)}) + \sum_{f \in \mathbf{B}} (e^{-\bar{c}d(f,g)} + e^{-\bar{c}d(\theta_{\dagger} f, g)}) \right].$$

Replacing the sums over f by sums over all $f \in \mathcal{F}(\mathbb{Z}^3)$, Claim 4.14 implies this contributes, at most, $4\bar{C}\bar{K}m(\mathcal{I}; \mathcal{J})$.

(ii) If $g \in \mathbf{W} \cup \mathbf{H} \cup W_x^{\mathcal{J}}$, these sums are, at most,

$$\sum_{g \in \mathbf{W} \cup \mathbf{H} \cup W_x^{\mathcal{J}}} \bar{K} \left[\sum_{f \in \mathbf{X}_1^{\mathcal{I}} \cup \mathbf{X}_2^{\mathcal{I}}} (e^{-\bar{c}d(f,g)} + e^{-\bar{c}d(\vec{\theta} f, g)}) + \sum_{f \in \mathbf{B}} (e^{-\bar{c}d(f,g)} + e^{-\bar{c}d(\theta_{\dagger} f, g)}) \right].$$

Replacing the sums over f by sums over all $f \in \mathcal{F}(\mathbb{Z}^3)$, by Claim 4.13 and Claim 4.15, this contributes, at most, $12\bar{C}\bar{K}m(\mathcal{I}; \mathcal{J})$.

(iii) For $g \in \mathbf{X}_1^{\mathcal{I}} \cup \mathbf{X}_1^{\mathcal{J}} \cup \mathbf{X}_2^{\mathcal{I}} \cup \mathbf{X}_2^{\mathcal{J}}$, let us begin with the first sum ($f \in \mathbf{X}_1^{\mathcal{I}} \cup \mathbf{X}_2^{\mathcal{I}}$). If $f \in \mathbf{X}_1^{\mathcal{I}}$, the radius \mathbf{r} could not have been attained by $g \in \mathbf{X}_1^{\mathcal{I}} \cup \mathbf{X}_1^{\mathcal{J}}$ since all increments in $\mathbf{X}_1^{\mathcal{I}}$ are shifted by the same vector. Then, this reduces to

$$\bar{K} \sum_{f \in \mathbf{X}_1^{\mathcal{I}}} \sum_{g \in \mathbf{X}_2^{\mathcal{I}}} (e^{-\bar{c}d(f,g)} + e^{-\bar{c}d(\vec{\theta} f, \vec{\theta} g)}),$$

which is, at most, $\bar{C}\bar{K}$ by Lemma 4.18. If $f \in \mathbf{X}_2^{\mathcal{I}}$, this reduces to $g \in \mathbf{X}_1^{\mathcal{I}}$ and is handled symmetrically.

Turning to the sum over $f \in \mathbf{B}$, it splits into the following:

$$\bar{K} \sum_{g \in \mathbf{X}_1^{\mathcal{I}} \cup \mathbf{X}_2^{\mathcal{I}}} \sum_{f \in \mathbf{C}_1 \cup \mathbf{C}_2 \cup \mathbf{F}} (e^{-\bar{c}d(f,g)} + e^{-\bar{c}d(\theta_{\dagger} f, \vec{\theta} g)}).$$

The contribution from $f \in \mathbf{C}_1$ is, at most, $\bar{C}\bar{K}m(\mathcal{I}; \mathcal{J})$ by Lemma 4.19; the contribution from $f \in \mathbf{C}_2$ is, at most, $\bar{C}\bar{K}$ by Lemma 4.20; the contribution from $f \in \mathbf{F}$ is, at most, $\bar{C}\bar{K}$ by Lemma 4.16 as $\mathbf{F} \subset \mathcal{L}_0$.

(iv) For $g \in \mathbf{C}_1 \cup \theta_{\dagger} \mathbf{C}_1 \cup \mathbf{C}_2 \cup \theta_{\dagger} \mathbf{C}_2 \cup \mathbf{F}$, the first sum can be expressed as

$$\bar{K} \sum_{f \in \mathbf{X}_1^{\mathcal{I}} \cup \mathbf{X}_2^{\mathcal{I}}} \sum_{g \in \mathbf{C}_1 \cup \mathbf{C}_2 \cup \mathbf{F}} (e^{-\bar{c}d(f,g)} + e^{-\bar{c}d(\vec{\theta} f, \theta_{\dagger} g)}).$$

Up to a change of roles of f and g , this is identical to the term considered in the item above, and its contribution is, therefore, at most, $2\bar{C}\bar{K} + \bar{C}\bar{K}m(\mathcal{I}; \mathcal{J})$ by Lemmas 4.16 and 4.19–4.20.

For the second sum, in which $f \in \mathbf{B}$, note that if the radius $\mathbf{r}(f, \mathcal{I}; \theta_{\dagger} f, \mathcal{J})$ is attained by $g \in \mathbf{C}_1 \cup \mathbf{C}_2 \cup \mathbf{F} \cup \mathbf{W}$ or by θ_{\dagger} of such a g , it must be attained by a face in a wall nested in some wall of \mathbf{W} . Since the distance between two faces is at least the distance between their projections and projections of distinct walls are distinct, the contribution of this term (summed over all possible such g) is, at most,

$$\bar{K} \sum_{f \in \mathbf{B}} \sum_{u \in \rho(\cup_{z \in \mathbf{D}} \tilde{W}_z)} e^{-\bar{c}d(\rho(f), u)},$$

which is, at most, $\bar{C}\bar{K}m(\mathbf{W}) \leq 3\bar{C}\bar{K}m(\mathcal{I}; \mathcal{J})$ by Lemma 4.17 and (4.4).

Altogether, we deduce that all the summands on the right-hand side of (4.7)–(4.9) are bounded by an absolute constant times $m(\mathcal{I}; \mathcal{J})$, implying the desired. \square

4.5. *Proof part 2: Multiplicity.* We next bound the multiplicity of the map $\Psi_{x,t}$ with a fixed excess area M by an exponential in M (independently of β).

PROPOSITION 4.21. *There exists some universal C_Ψ such that, for every $M \geq 1$ and every x, t and $T < H$,*

$$\max_{\mathcal{J} \in \Psi_{x,t}(\bar{\mathbf{I}}_{x,T,H})} |\{\mathcal{I} \in \Psi_{x,t}^{-1}(\mathcal{J}) : m(\mathcal{I}; \mathcal{J}) = M\}| \leq C_\Psi^M.$$

Toward proving Proposition 4.2, we are also interested in a map used to prove an exponential tail on the increment of a pillar that intersects a given height ℓ (as opposed to an increment of a given index). For that purpose, for any pillar \mathcal{P}_x and a half-integer height ℓ , let

$$\tau_\ell = \tau_\ell(\mathcal{P}_x) := \min\{t \geq 1 : \mathcal{X}_t \cap \mathcal{L}_\ell \neq \emptyset\} \vee 1,$$

and define the map $\tilde{\Psi}_{x,\ell}$ as

$$\tilde{\Psi}_{x,\ell}(\mathcal{I}) := \Psi_{x,\tau_\ell(\mathcal{P}_x)}(\mathcal{I}).$$

Clearly, since the bound of Proposition 4.12 is independent of \mathcal{I} and t , the estimate also holds for $\tilde{\Psi}_{x,\ell}(\mathcal{I})$. However, handling the multiplicity is slightly different since interfaces with differing τ_ℓ may be mapped to the same $\mathcal{J} \in \tilde{\Psi}_{x,\ell}(\bar{\mathbf{I}}_{x,T,H})$.

PROPOSITION 4.22. *There exists a universal $\tilde{C}_\Psi > 1$ such that, for every $M \geq 1$ and every x, ℓ and $T < H$,*

$$\max_{\mathcal{J} \in \tilde{\Psi}_{x,\ell}(\bar{\mathbf{I}}_{x,T,H})} |\{\mathcal{I} \in \tilde{\Psi}_{x,\ell}^{-1}(\mathcal{J}) : m(\mathcal{I}; \mathcal{J}) = M\}| \leq \tilde{C}_\Psi^M.$$

We prove Propositions 4.21–4.22 by constructing a witness that (given $\Psi_{x,t}(\mathcal{I})$) is in 1-1 correspondence with the pre-image \mathcal{I} . We then bound the number of all possible such witnesses.

Let us fix any $\mathcal{J} \in \Psi_{x,t}(\bar{\mathbf{I}}_{x,T,H})$ (or, resp., $\mathcal{J} \in \tilde{\Psi}_{x,\ell}(\bar{\mathbf{I}}_{x,T,H})$). We wish to define an injective map $\Xi = \Xi_{\mathcal{J},x,t}$ (resp., $\tilde{\Xi} = \tilde{\Xi}_{\mathcal{J},x,\ell}$ on $\{\mathcal{I} \in \Psi_{x,t}^{-1}(\mathcal{J})\}$) and bound the cardinality of the set $\Xi(\{\mathcal{I} \in \Psi_{x,t}^{-1}(\mathcal{J}) : m(\mathcal{I}; \mathcal{J}) = M\})$ (resp., $\tilde{\Xi}(\{\mathcal{I} \in \Psi_{x,\ell}^{-1}(\mathcal{J}) : m(\mathcal{I}; \mathcal{J}) = M\})$).

Construction of the witness. Fix x and t (resp., ℓ). We describe how, for a given \mathcal{J} and an $\mathcal{I} \in \Psi_{x,t}^{-1}(\mathcal{J})$ (resp., $\mathcal{I} \in \tilde{\Psi}_{x,\ell}^{-1}(\mathcal{J})$), we construct the witness $\Xi_{\mathcal{J},x,t}(\mathcal{I})$ (resp., $\tilde{\Xi}_{\mathcal{J},x,\ell}(\mathcal{I})$). In order to do this in a unified manner, we let

$$\tilde{\Xi}_{\mathcal{J},x,\ell}(\mathcal{I}) = \Xi_{\mathcal{J},x,\tau_\ell(\mathcal{P}_x)}(\mathcal{I}),$$

and then it suffices to describe how to construct $\Xi(\mathcal{I}) = \Xi_{\mathcal{J},x,t}(\mathcal{I})$ for each \mathcal{I} .

Our witness $\Xi(\mathcal{I})$ will consist of six $*$ -connected face-subsets $(\mathcal{F}_i^\gamma)_{i \in \{A,B\}, \gamma \in \{\mathcal{I}, \mathcal{J}, \Psi\}} \subset \mathcal{F}(\mathbb{Z}^3)$, each of which are decorated by coloring its faces BLUE or RED, and associating to each $v \in \mathcal{F}_i^\gamma$, its own individual face-subset $\Upsilon_v \subset \mathcal{F}(\mathcal{F}(\mathbb{Z}^3))$ whose faces are also colored BLUE or RED.

Let us begin by constructing the six $*$ -connected face-subsets \mathcal{F}_i^γ and their colorings. Partition \mathbf{Y} into \mathbf{Y}_\emptyset , and $\mathbf{Y}_A^\mathcal{I} \cup \mathbf{Y}_A^\mathcal{J} \cup \mathbf{Y}_A^\Psi$ along $\mathbf{Y}_B^\mathcal{I} \cup \mathbf{Y}_B^\mathcal{J} \cup \mathbf{Y}_B^\Psi$ as follows:

- Let $\mathbf{Y}_\emptyset = \{v_1, y^\dagger\} \cup [x]$ (recalling that $[x]$ indicates x and its four adjacent faces in $\mathcal{L}_{0,n}$).
- Let $\mathbf{Y}_A^\mathcal{I}$ (resp., $\mathbf{Y}_B^\mathcal{I}$) be the indices of walls that were first marked for deletion due to one of the criteria (A2), (A3) (resp., (B2), (B3)) wherein $\mathfrak{D}_x(\tilde{W}_y, i, \omega_1, \omega_2)$ was attained by $d(\Theta_\dagger^{\mathcal{I}} \setminus \mathcal{S}_x \tilde{W}_y, \mathcal{X}_i)$.

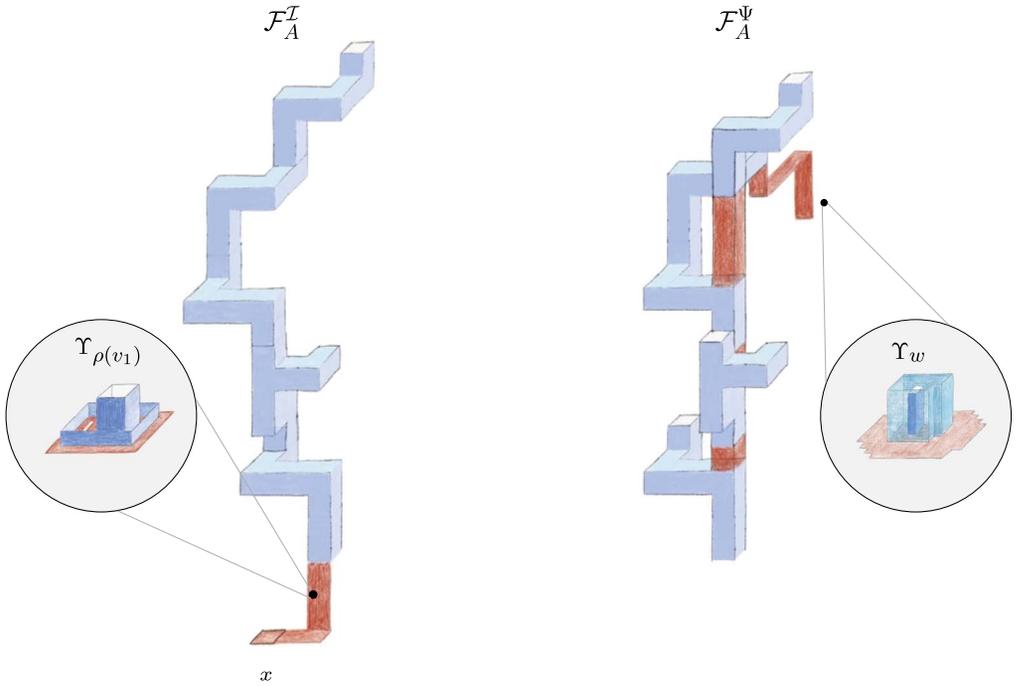


FIG. 7. Two of the six constituent parts of the witness. Left: $\mathcal{F}_A^{\mathcal{I}}$ decorated by {BLUE, RED} and the vertex-decoration $\Upsilon_{\rho(v_1)}$. Right: \mathcal{F}_A^{Ψ} decorated by {BLUE, RED} and the vertex-decoration Υ_w .

- Let $\mathbf{Y}_A^{\mathcal{J}}$ (resp., $\mathbf{Y}_B^{\mathcal{J}}$) be the indices of walls that were first marked for deletion due to (A2) or (A3) (resp., (B2) or (B3)) wherein $\mathcal{D}_x(\tilde{W}_y, i, \omega_1, \omega_2)$ was attained by $d(\Theta_{\dagger}^{\mathcal{I} \setminus \mathcal{S}_x} \tilde{W}_y, \theta_{\rho(x+\omega_2)} \Theta_{\emptyset} \mathcal{X}_i)$.
- Let \mathbf{Y}_A^{Ψ} (resp., \mathbf{Y}_B^{Ψ}) are the indices of walls that were first marked for deletion due to (A2) or (A3) (resp., (B2) or (B3)) wherein $\mathcal{D}_x(\tilde{W}_y, i, \omega_1, \omega_2)$ was attained by $d(\Theta_{\dagger}^{\mathcal{I} \setminus \mathcal{S}_x} \tilde{W}_y, \theta_{\rho(x+\omega_1+\omega_2)} \mathcal{X}_i)$.

When considering the criteria (A2), (A3), (B2), (B3), the Euclidean distance between sets of faces in \mathbb{R}^3 is attained by vertices of \mathbb{Z}^3 , which we will endow with an (arbitrary) lexicographic ordering, giving rise to unique minimizers of the distance. Further, let

$$\mathbf{X}_A^{\Psi} = \bigcup \{ \mathcal{F}(\theta_{\rho(x-v_{s_{j+1}})}) \mathcal{X}_j : 1 \leq j \leq j^* \},$$

$$\mathbf{X}_B^{\Psi} = \bigcup \{ \mathcal{F}(\theta_{\rho(x+v_l-v_{j^*+1}-v_{s_{k+1}})}) \mathcal{X}_k : t \leq k \leq k^* \}.$$

With these definitions in hand, the witness $\Xi(\mathcal{I})$ is constructed as follows:

- (1) The BLUE faces of \mathcal{F}_l^{γ} for $l \in \{A, B\}$ and $\gamma \in \{\mathcal{I}, \mathcal{J}, \Psi\}$ are precisely \mathbf{X}_l^{γ} (calculated via running $\Psi_{x,t}$).
- (2) For each $y \in \mathbf{Y}_l^{\gamma}$ (for $l \in \{A, B\}$ and $\gamma \in \{\mathcal{I}, \mathcal{J}, \Psi\}$):
 - Let $v \in \mathbf{X}_l^{\gamma}$ and $w \in \Theta_{\dagger}^{\mathcal{I} \setminus \mathcal{S}_x} \tilde{W}_y$ be the minimizers of the distance that was violated in the respective deletion criterion (at the first time y was marked for deletion).
 - Add to \mathcal{F}_l^{γ} a shortest (nearest-neighbor) path of RED faces connecting v to w .
- (3) Add to $\mathcal{F}_A^{\mathcal{I}}$ a shortest path of RED faces connecting v_1 and y^{\dagger} and such a path connecting v_1 and x ; include the face x , and color it RED.
- (4) Add to \mathcal{F}_l^{Ψ} for $l \in \{A, B\}$ every face of $\mathbf{X}_l^{\mathcal{J}}$ that is not already present, and color it RED.

(5) Let Σ be the set of all vertices w for which we added a shortest RED path from v to w in step (2) above. Process the vertices in $w \in \Sigma \cup \mathcal{V}(\mathbf{Y}_{\mathcal{B}})$ via some lexicographic order $(w^1, w^2, \dots, w^{|\Sigma \cup \mathcal{V}(\mathbf{Y}_{\mathcal{B}})|})$: for $i = 1, \dots, |\Sigma \cup \mathcal{V}(\mathbf{Y}_{\mathcal{B}})|$:

- If w^i is associated with the wall \tilde{W}_y (i.e., $w^i \in \Theta_{\ddagger}^{\mathcal{I} \setminus \mathcal{S}_x} \tilde{W}_y$), let the BLUE faces of Υ_{w^i} be the set

$$\mathfrak{S}_{w^i} = \tilde{\mathfrak{F}}_y \setminus \bigcup_{j < i} \mathfrak{S}_{w^j}.$$

- For each edge or face $u \in \rho(\mathfrak{S}_{w^i})$, add to Υ_{w^i} the RED set of faces in the ball of radius $N_\rho(u)$ around u in $\mathcal{L}_{0,n}$. Complete Υ_{w^i} by connecting all faces added to it, together with the vertex $\rho(w^i)$, via a RED minimum size spanning tree of faces in $\mathcal{L}_{0,n}$.

(For every other vertex $w \in \mathcal{V}(\mathcal{F}_i^\gamma)$, we let $\Upsilon_w = \emptyset$.)

See Figure 7 for examples of $\mathcal{F}_A^{\mathcal{I}}$ and \mathcal{F}_A^Ψ together with their decorations.

Reconstructing \mathcal{I} from the witness. To see that this indeed yields a “witness” of the pre-image interface \mathcal{I} , we show that, from a witness $\Xi(\mathcal{I})$ and the interface \mathcal{J} , one can reconstruct \mathcal{I} .

LEMMA 4.23. *For every $\mathcal{J} \in \Psi_{x,t}(\bar{\mathbf{I}}_{x,T,H})$ (resp., $\mathcal{J} \in \tilde{\Psi}_{x,\ell}(\bar{\mathbf{I}}_{x,T,H})$), the map $\Xi_{\mathcal{J},x,t}$ (resp., $\tilde{\Xi}_{\mathcal{J},x,\ell}$) is injective on $\Psi_{x,t}^{-1}(\mathcal{J})$ (resp., $\tilde{\Psi}_{x,\ell}^{-1}(\mathcal{J})$).*

PROOF. It suffices to show that, from a given \mathcal{J} and any element of $\Xi_{\mathcal{J},x,t}(\Psi_{x,t}^{-1}(\mathcal{J}))$, we can recover, uniquely, $\mathcal{I} \in \Psi_{x,t}^{-1}(\mathcal{J})$. From a witness in $\Xi_{\mathcal{J},x,t}(\Psi_{x,t}^{-1}(\mathcal{J}))$, we recover \mathcal{I} by reconstructing its spine \mathcal{S}_x together with the standard wall representation of $\mathcal{I} \setminus \mathcal{S}_x$. Given \mathcal{S}_x and $(\tilde{W}_z)_{z \in \mathcal{L}_{0,n}}$, we would obtain \mathcal{I} by first recovering the interface $\mathcal{I} \setminus \mathcal{S}_x$ via Lemma 2.11, then appending to that \mathcal{S}_x .

1. In order to reconstruct the spine \mathcal{S}_x :

- Extract $\mathbf{X}_A^{\mathcal{I}}$ and $\mathbf{X}_B^{\mathcal{I}}$ as exactly the set of BLUE faces of $\mathcal{F}_A^{\mathcal{I}}$ and $\mathcal{F}_B^{\mathcal{I}}$, respectively.
- Extract $\mathbf{X}_1^{\mathcal{J}}$ and $\mathbf{X}_2^{\mathcal{J}}$ by taking (the bounding faces of) all cells in $\mathcal{P}_x^{\mathcal{J}}$, between $\text{ht}(v_{j^*+1})$ and $\text{ht}(v_\ell)$ and above $\text{ht}(v_{k^*+1})$, respectively (these heights are read off from $\mathbf{X}_A^{\mathcal{I}}$ and $\mathbf{X}_B^{\mathcal{I}}$).
- Obtain \mathcal{S}_x by horizontally shifting $\mathbf{X}_1^{\mathcal{J}}$ and $\mathbf{X}_2^{\mathcal{J}}$ so that their bottom cell coincides with the top cell of $\mathbf{X}_A^{\mathcal{I}}$ and $\mathbf{X}_B^{\mathcal{I}}$, respectively.

2. In order to reconstruct the standard wall representation of $\mathcal{I} \setminus \mathcal{S}_x$:

- For every vertex $w \in \mathcal{V}(\mathcal{F}_i^\gamma)$ for $i \in \{A, B\}$ and $\gamma \in \{\mathcal{I}, \mathcal{J}, \Psi\}$, add the faces of \mathfrak{S}_w (exactly the set of BLUE faces of Υ_w).
- Add the standardizations of all walls of $\mathcal{J} \setminus \mathcal{P}_x^{\mathcal{J}}$.

For the corresponding reconstruction from a witness in $\tilde{\Xi}_{\mathcal{J},x,\ell}(\tilde{\Psi}_{x,\ell}^{-1}(\mathcal{J}))$, we recover \mathcal{I} in exactly the same way, noticing that we can read off $\text{ht}(v_{\tau_\ell})$ from $\mathbf{X}_B^{\mathcal{I}}$. \square

Enumerating over possible witnesses. It remains to enumerate over the set of all possible witnesses of interfaces in $\Psi_{x,t}^{-1}(\mathcal{J})$ and $\tilde{\Psi}_{x,\ell}^{-1}(\mathcal{J})$ with excess area $m(\mathcal{I}; \mathcal{J}) = M$ and to show it is, at most, exponential in M .

LEMMA 4.24. *There exists some universal $C_\Psi > 0$ such that, for every $M \geq 1$, x, t , and $\mathcal{J} \in \Psi_{x,t}(\bar{\mathbf{I}}_{x,T,H})$,*

$$|\{\Xi_{\mathcal{J},x,t}(\mathcal{I}) : \mathcal{I} \in \Psi_{x,t}^{-1}(\mathcal{J}), m(\mathcal{I}; \mathcal{J}) = M\}| \leq C_\Psi^M.$$

Similarly, there exists \tilde{C}_Ψ such that, for every $M \geq 1$, x, ℓ and $\mathcal{J} \in \tilde{\Psi}_{x,\ell}(\bar{\mathbf{I}}_{x,T,H})$,

$$|\{\tilde{\Xi}_{\mathcal{J},x,\ell}(\mathcal{I}) : \mathcal{I} \in \tilde{\Psi}_{x,\ell}^{-1}(\mathcal{J}), m(\mathcal{I}; \mathcal{J}) = M\}| \leq \tilde{C}_\Psi^M.$$

Combining the above lemma with Lemma 4.23 immediately implies Propositions 4.21–4.22.

PROOF. Let us prove the bound on the number of possible witnesses corresponding to $\Psi_{x,t}$ and $\tilde{\Psi}_{x,\ell}$, simultaneously, describing in the proof the parts that are different between $\Psi_{x,t}$ and $\tilde{\Psi}_{x,\ell}$.

Fix M, x, t, ℓ and $T < H$ and $\mathcal{J} \in \Psi_{x,t}(\bar{\mathbf{I}}_{x,T,H})$ (resp., $\mathcal{J} \in \Psi_{x,\ell}(\bar{\mathbf{I}}_{x,T,H})$), and consider the number of possible witnesses $\Xi(\mathcal{I})$ for \mathcal{I} satisfying $m(\mathcal{I}; \mathcal{J}) = M$. We decompose this into the number of possible choices of colored face-sets \mathcal{F}_i^γ , and, subsequently, the number of choices of decorations to the vertices of \mathcal{F}_i^γ via the number of choices of colored face-sets $(\Upsilon_w)_w$. Clearly, their product bounds the number of possible choices of witnesses:

Number of faces in \mathcal{F}_i^γ . We first bound the number of faces in each \mathcal{F}_i^γ for $i \in \{A, B\}$ and $\gamma \in \{\mathcal{I}, \mathcal{J}, \Psi\}$.

- (1) The number of BLUE faces in \mathcal{F}_i^γ is exactly $|\mathbf{X}_i^\gamma|$. For $\gamma \in \{\mathcal{I}, \mathcal{J}\}$, this quantity is, at most, $\tilde{C}m(\mathcal{I}; \mathcal{J})$ by Claim 4.14. For $\gamma = \Psi$, we have that

$$|\mathbf{X}_i^\Psi| \leq |\mathbf{X}_i^\mathcal{I}| + |\mathbf{X}_i^\mathcal{J}| \leq 2\tilde{C}m(\mathcal{I}; \mathcal{J}).$$

(Note that \mathbf{X}_i^Ψ consists of shifts of the increments composing $\mathbf{X}_i^\mathcal{I}$; these shifts add an additional cell whenever $s_i \neq s_{i-1}$, due to the additional four faces of the shared cut-point between consecutive increments. We compensate for these via the term $|\mathbf{X}_i^\mathcal{J}|$.)

- (2) For each $w \in \Sigma$ associated with some wall \tilde{W}_y (for some $y \in \mathbf{Y}_i^\gamma$):

- If $y \neq y_i^*$, then the number of RED faces that were added to connect $w \in \Theta_{\downarrow}^{\mathcal{I} \setminus \mathcal{S}_x} \tilde{W}_y$ to $|\mathbf{X}_i^\gamma|$ is at most $2m(\tilde{W}_y)$ (since \tilde{W}_y violated criteria (A2) or (B2)), where the factor of 2 accounts for the transition from Euclidean distance to the graph distance in \mathbb{Z}^3 .
- If $y = y_i^*$, then the number of RED faces that were added to connect $w \in \Theta_{\downarrow}^{\mathcal{I} \setminus \mathcal{S}_x} \tilde{W}_y$ to $|\mathbf{X}_i^\gamma|$ is, at most, $(j^* - 1) \vee (k^* - t)$ (since \tilde{W}_y either violated criteria (A3) or (B3)).

Summing these over all $y \in \mathbf{Y}_i^\gamma$ gives at most $m(\mathbf{W}) + (j^* - 1) + (k^* - t)$ additional RED faces, which is, at most, $15m(\mathcal{I}; \mathcal{J})$ by Claim 4.7.

- (3) The number of RED faces added to connect v_1 to y^\dagger as well as to x is, at most, $m(\tilde{\mathfrak{F}}_{v_1})$, since all of these are part of \mathcal{P}_x , and hence share some nesting wall by Observation 3.2. Thus, the number of such faces is, at most, $2m(\mathbf{W}) \leq 6m(\mathcal{I}; \mathcal{J})$ by Claim 4.7.
- (4) The number of RED faces added to \mathcal{F}_i^Ψ by $\mathbf{X}_i^\mathcal{J}$ is, at most, $\tilde{C}m(\mathcal{I}; \mathcal{J})$ by Claim 4.14.

Altogether, we see that there exists a universal $C > 0$ such that, for every $M \geq 1$,

$$\max_{\mathcal{I} \in \Psi_{x,t}^{-1}(\mathcal{J}) : m(\mathcal{I}; \mathcal{J}) = M} \max_{t \in \{A, B\}, \gamma \in \{\mathcal{I}, \mathcal{J}, \Psi\}} |\mathcal{F}_i^\gamma| \leq CM.$$

Since the constant C above was uniform over the choice of t and the witness constructed by $\tilde{\Xi}_{x,\ell}$ agrees with the witness constructed by $\Xi_{x,t}$ for some $t = t(\mathcal{I})$, they apply equally to the witnesses coming from $\tilde{\Xi}_{x,\ell}$.

Number of possible choices of $\mathcal{F}_i^{\mathcal{Y}}$. We begin by enumerating over the choices of $\mathcal{F}_A^{\mathcal{I}}$: by Fact 3.13, the number of $*$ -connected face sets rooted at x (predetermined) with, at most, CM faces is, at most, s^{CM} ; multiplying this by 2^{CM} for the choices of BLUE and RED colorings of these faces, bounds the total number of possible choices for $\mathcal{F}_A^{\mathcal{I}}$. Notice that the choice of $\mathcal{F}_A^{\mathcal{I}}$ reveals v_1 as its lowest cell that is bounded by (four) BLUE faces and v_{j^*+1} as its highest such cell. This also determines whether $t > j^*$ or not.

Next, if $t > j^*$, we enumerate over the choices of $\mathcal{F}_B^{\mathcal{I}}$. In the case of $\Xi_{x,t}$, its root, v_t , is determined by our above choice of $\mathcal{F}_A^{\mathcal{I}}$ as follows: starting from v , the cut-point of $\mathcal{P}_x^{\mathcal{J}}$ at height $\text{ht}(v_{j^*+1})$, count $t - (j^* + 1)$ extra increments (with t predetermined) to a cut-point v' (marking the top of $\mathbf{X}_1^{\mathcal{J}}$ and the bottom of $\mathbf{X}_B^{\mathcal{J}}$). The root v_t is $\theta_{v'-v}(v_{j^*+1})$. Enumerating over $\mathcal{F}_B^{\mathcal{I}}$ then amounts to another factor of $(2s)^{CM}$.

In the case of $\tilde{\Xi}_{x,\ell}$, we can enumerate over choices of root v_{τ_ℓ} by enumerating over the cut-point v' of \mathcal{J} whose height coincides with $\text{ht}(v_{\tau_\ell})$. In \mathcal{J} the cut-point v' must be within a height of, at most, $m(\mathcal{X}_{\tau_\ell}) \leq M$ from ℓ , so there are at most M such choices of cut-point v' . From that we recover the root v_{τ_ℓ} as $\theta_{v'-v}(v_{j^*+1})$; choosing $\mathcal{F}_B^{\mathcal{I}}$ thus amounts to a factor of $M(2s)^{CM}$.

We bound the number of possible choices for $\mathcal{F}_A^{\mathcal{J}}$ from above by s^{CM} , the number of $*$ -connected sets of CM faces, rooted at $x + (0, 0, \text{ht}(v_1))$ (recall that v_1 can be read off of $\mathcal{F}_A^{\mathcal{I}}$), multiplied by a 2^{CM} factor for coloring those faces by BLUE and RED. If $t > j^*$, the number of choices of $\mathcal{F}_B^{\mathcal{J}}$ is, at most, the number of $*$ -connected sets of CM faces, rooted at the cut-point v' mentioned above (which we read off of \mathcal{J} and $\mathcal{F}_A^{\mathcal{I}}$), along with its coloring by BLUE and RED which combine to, at most, $(2s)^{CM}$.

To enumerate over \mathcal{F}_A^{Ψ} , we first argue that it is a $*$ -connected set of faces. In order to see this, recall that the BLUE faces of \mathcal{F}_A^{Ψ} are comprised of various subsets of shifted increment sequences, which can only become disconnected when there is a change in the shift applied, that is, at increments where $s_{j+1} \neq s_j$. Suppose that $s_{j+1} \neq s_j$ for some $1 \leq j < j^*$; by Observation 4.6 this occurs if and only if $s_{j+1} = j$ and \mathcal{X}_{j+1} is being shifted in \mathbf{X}_A^{Ψ} by

$$\rho(x - v_{s_{j+1}+1}) = \rho(x - v_{j+1}).$$

It suffices to show that for every such j , the cell $\theta_{\rho(x-v_{j+1})}v_{j+1}$ is $*$ -connected to the cell $x + (0, 0, \text{ht}(v_1))$ by faces in \mathcal{F}_A^{Ψ} . This follows since the bounding faces of the cell $\theta_{\rho(x-v_{j+1})}v_{j+1}$ are also in $\mathbf{X}_A^{\mathcal{J}}$ (as a subset of $\theta_{\rho(x)}\Theta_{\emptyset}\mathcal{X}_{j+1}$, to which \mathcal{X}_{j+1} was mapped), $\mathbf{X}_A^{\mathcal{J}}$ connects this cell to $x + (0, 0, \text{ht}(v_1))$ (via trivial increments) and $\mathbf{X}_A^{\mathcal{J}} \subset \mathcal{F}_A^{\Psi}$. As a connected set of, at most, CM faces, colored by BLUE and RED and rooted at $x + (0, 0, \text{ht}(v_1))$ (which is dictated by our choice of $\mathcal{F}_A^{\mathcal{I}}$), there are, at most, $(2s)^{CM}$ choices for \mathcal{F}_A^{Ψ} .

Similarly, if $t > j^*$, to see that \mathcal{F}_B^{Ψ} is $*$ -connected, note that for $t \leq k < k^*$, the shifted increments k and $k + 1$ can only be disconnected if $s_{k+1} \neq s_k$, which occurs if and only if $s_{k+1} = k$, in which case the increment \mathcal{X}_{k+1} will be shifted in \mathbf{X}_B^{Ψ} by

$$\rho(x + v_t - v_{j^*+1} - v_{s_{k+1}+1}) = \rho(x + v_t - v_{j^*+1} - v_{k+1}).$$

The fact that $\mathbf{X}_B^{\mathcal{J}}$ (which, as before, connects $\theta_{\rho(x+v_t-v_{j^*+1}-v_{k+1})}v_{k+1}$ to $x + v_t - v_{j^*+1}$ via trivial increments) is a subset of the faces of \mathcal{F}_B^{Ψ} implies that \mathcal{F}_B^{Ψ} is $*$ -connected, as claimed. As a $*$ -connected set of, at most, CM faces, colored by BLUE and RED and rooted at $x + v_t - v_{j^*+1}$ (read off of our choice of $\mathcal{F}_A^{\mathcal{I}}$ and \mathcal{J}), there are, at most, $(2s)^{CM}$ choices for \mathcal{F}_B^{Ψ} . In the case of $\tilde{\Xi}_{x,\ell}$ where this is $x + v_{\tau_\ell} - v_{j^*+1}$, recall that our choice of $\mathcal{F}_B^{\mathcal{I}}$ picked out the cut-point v' from which we read off v_{τ_ℓ} , so the same bound holds.

Number of possible decorations (Υ_w) to $w \in \mathcal{V}(\mathcal{F}_i^{\mathcal{Y}})$. Let us first count the combined number of faces

$$(4.10) \quad \sum_{\iota \in \{A, B\}, \gamma \in \{\mathcal{I}, \mathcal{J}, \Psi\}} \sum_{w \in \mathcal{V}(\mathcal{F}_i^{\mathcal{Y}})} |\Upsilon_w|.$$

By construction, the sets \mathfrak{S}_w are all disjoint. The number of BLUE faces in total over all Υ_w is, therefore, at most, $\sum_{z \in \mathbf{D}} |W_z| \leq \sum_{z \in \mathbf{D}} 2\mathfrak{m}(W_z) \leq 6\mathfrak{m}(\mathcal{I}; \mathcal{J})$ (as no wall is double counted).

For each w , the number of RED faces added to Υ_w in \mathcal{L}_0 in the ball of radius $\sqrt{N_\rho(u)}$ centered at an edge-or-face $u \in \rho(\mathfrak{S}_w)$ is, at most, using Claim 4.7,

$$\sum_w \sum_{u \in \rho(\mathfrak{S}_w)} N_\rho(u) \leq \sum_w |\mathfrak{S}_w| \leq 2\mathfrak{m}(\mathbf{W}) \leq 6\mathfrak{m}(\mathcal{I}; \mathcal{J}).$$

By induction, every \mathfrak{S}_w consists of the groups of walls of a nested sequence of walls. Indeed, when we allocate \mathfrak{S}_w , if \tilde{W} already decorates some previously processed vertex z , then, necessarily, the entire group of walls of \tilde{W} must also have been allocated to \mathfrak{S}_z , so the remainder is still the group of walls of a nested sequence of walls. For each w , denote this nested sequence of walls by $W_1^w \Subset W_2^w \Subset \dots$.

Within every \mathfrak{S}_w , all close walls are connected, via the additional RED faces in \mathcal{L}_0 in balls of radius $\sqrt{N_\rho(u)}$ for $u \in \rho(\mathfrak{S}_w)$. Each of the $*$ -connected components obtained in this way (whose BLUE faces are precisely a group of walls) corresponds to some W_i^w (say the innermost one it contains). Finally, by definition, $\rho(w)$ is interior to W_1^w . Therefore, to obtain a spanning tree of the face-set, we can include shortest paths of faces from $\rho(w)$ to W_1^w and then from W_i^w to W_{i+1}^w for every i . This adds, at most, $\sum_i |W_i^w| \leq 2\mathfrak{m}(\mathfrak{S}_w)$ many faces, and the minimum spanning tree adds, at most, that many RED faces. Summing over all w , this last contribution (again by (4.4)) is also, at most, $2 \sum_w \mathfrak{m}(\mathfrak{S}_w) \leq 6\mathfrak{m}(\mathcal{I}; \mathcal{J})$.

Altogether, we deduce that the total number of faces (4.10) is, at most, $18\mathfrak{m}(\mathcal{I}; \mathcal{J}) = 18M$.

In order to enumerate over all such possible decorating face-sets, let us first decide how many of the $18M$ faces are allocated to each Υ_w . For every $\iota \in \{A, B\}$ and $\gamma \in \{\mathcal{I}, \mathcal{J}, \Psi\}$, we have that $\mathcal{V}(\mathcal{F}_\iota^\gamma) \leq 4CM$ by our bound on the number of faces in that set; in particular, there are, at most, $24CM$ vertices, between which we wish to partition at most $18M$ decorating faces. The number of such partitions is, at most,

$$\binom{(24C + 18)M - 1}{24CM - 1} \leq 2^{(24C+18)M}.$$

For each such partition, if k_w is the number of decorating faces assigned to w (so that $\sum_w k_w \leq 18M$), then we have s^{k_w} choices for a $*$ -connected decorating face subset rooted at $\rho(w)$ and 2^{k_w} choices of RED and BLUE colors for these faces. Thus, the total number of choices for the decorating colored face subsets corresponding to this partition is, at most, $(2s)^{18M}$.

Multiplying all of the above enumerations yields the desired bounds for some C_Ψ and \tilde{C}_Ψ . □

4.6. Proofs of Theorem 4.1 and Proposition 4.2.

PROOF OF THEOREM 4.1. Recall from Claim 4.7 that, for every t and every \mathcal{I} , we have that

$$\mathfrak{m}(\mathcal{I}; \Psi_{x,t}(\mathcal{I})) \geq 2 \left(\text{ht}(v_1) - \frac{1}{2} \right) + 2d_\rho(v_1, x) \geq 2 \left| v_1 - \left(x + \left(0, 0, \frac{1}{2} \right) \right) \right|$$

as well as $\mathfrak{m}(\mathcal{I}; \Psi_{x,t}(\mathcal{I})) \geq \mathfrak{m}(\mathcal{B}_x)$ and $\mathfrak{m}(\mathcal{I}; \Psi_{x,t}(\mathcal{I})) \geq \mathfrak{m}(\mathcal{X}_t)$ so that it suffices to prove that

$$\mu_n^\mp(\mathfrak{m}(\mathcal{I}; \Psi_{x,t}(\mathcal{I})) \geq r \mid \mathbf{I}_{x,T,H}) \leq C \exp[-(\beta - C)(r \wedge d(x, \partial\Lambda_n))].$$

To see this, recall the definition of $\bar{\mathbf{I}}_{x,T,H}$, and express $\mu_n^\mp(\mathfrak{m}(\mathcal{I}; \Psi_{x,t}(\mathcal{I})) \geq r, \mathbf{I}_{x,T,H})$ as, at most,

$$\mu_n^\mp(\bar{\mathbf{I}}_{x,T,H}^c \mid \mathbf{I}_{x,T,H}) \mu_n^\mp(\mathbf{I}_{x,T,H}) + \mu_n^\mp(\mathfrak{m}(\mathcal{I}; \Psi_{x,t}(\mathcal{I})) \geq r, \bar{\mathbf{I}}_{x,T,H}).$$

The first term above is bounded by $e^{-2(\beta-C)d(x, \partial \Lambda_n)} \mu_n^\mp(\mathbf{I}_{x,T,H})$ by Proposition 3.12, so let us turn to the second term: for every $r \geq 1$,

$$\begin{aligned} \sum_{M \geq r} \sum_{\substack{\mathcal{I} \in \bar{\mathbf{I}}_{x,T,H} \\ m(\mathcal{I}; \Psi_{x,t}(\mathcal{I}))=M}} \mu_n^\mp(\mathcal{I}) &\leq \sum_{M \geq r} \sum_{\substack{\mathcal{I} \in \bar{\mathbf{I}}_{x,T,H} \\ m(\mathcal{I}; \Psi_{x,t}(\mathcal{I}))=M}} e^{-(\beta-C)M} \mu_n^\mp(\Psi_{x,t}(\mathcal{I})) \\ &= \sum_{M \geq r} \sum_{\mathcal{J} \in \Psi_{x,t}(\bar{\mathbf{I}}_{x,T,H})} \mu_n^\mp(\mathcal{J}) \sum_{\substack{\mathcal{I} \in \Psi_{x,t}^{-1}(\mathcal{J}) \\ m(\mathcal{I}; \Psi_{x,t}(\mathcal{I}))=M}} e^{-(\beta-C)M} \\ &\leq \sum_{M \geq r} C_\Psi^M e^{-(\beta-C)M} \mu_n^\mp(\Psi_{x,t}(\bar{\mathbf{I}}_{x,T,H})). \end{aligned}$$

In the first inequality above, we used Proposition 4.12 (and the fact that $\bar{\mathbf{I}}_{x,T,H}$ is in the domain of $\Psi_{x,t}$ by Proposition 4.5), and in the second inequality, we used Proposition 4.21.

Now, noting by Proposition 4.5 that

$$\Psi_{x,t}(\bar{\mathbf{I}}_{x,T,H}) \subset \bar{\mathbf{I}}_{x,T,H}$$

we deduce that

$$\begin{aligned} \mu_n^\mp(m(\mathcal{I}; \Psi_{x,t}(\mathcal{I})) \geq r, \bar{\mathbf{I}}_{x,T,H}) &\leq C e^{-(\beta-C-\log C_\Psi)r} \mu_n^\mp(\bar{\mathbf{I}}_{x,T,H}) \\ &\leq C e^{-(\beta-C-\log C_\Psi)r} \mu_n^\mp(\mathbf{I}_{x,T,H}). \end{aligned}$$

Combining these estimates and dividing through by $\mu_n^\mp(\mathbf{I}_{x,T,H})$ then yields the desired conditional bound. \square

PROOF OF PROPOSITION 4.2. By definition of $\tilde{\Psi}_{x,\ell}$, for every half-integer ℓ , we have

$$m(\mathcal{I}; \tilde{\Psi}_{x,\ell}(\mathcal{I})) \geq |\mathcal{F}(\mathcal{P}_x \cap \mathcal{L}_\ell)| - 4.$$

Indeed, this follows from the fact that if $\mathcal{P}_x \cap \mathcal{L}_\ell$ is not a single cell, either it is part of an increment in the spine, in which case τ_ℓ is the index of that increment and we use $m(\mathcal{I}; \tilde{\Psi}_{x,\ell}(\mathcal{I})) \geq m(\mathcal{X}_{\tau_\ell})$, or it is part of the base, in which case this follows from $m(\mathcal{I}; \tilde{\Psi}_{x,\ell}(\mathcal{I})) = m(\mathcal{I}; \Psi_{x,1}(\mathcal{I})) \geq m(\mathcal{B}_x)$.

As such, it suffices for us to show that, for every $r \geq 1$,

$$\mu_n^\mp(m(\mathcal{I}; \tilde{\Psi}_{x,\ell}(\mathcal{I})) \geq r \mid \mathbf{I}_{x,T,H}) \leq C \exp[-(\beta - C)(r \wedge d(x, \partial \Lambda_n))].$$

Arguing as in the proof of Theorem 4.1, we bound the left-hand side above by Proposition 3.12 by

$$e^{-2(\beta-C)d(x, \partial \Lambda_n)} \mu_n^\mp(\mathbf{I}_{x,T,H}) + \mu_n^\mp(m(\mathcal{I}; \tilde{\Psi}_{x,\ell}(\mathcal{I})) \geq r, \bar{\mathbf{I}}_{x,T,H}).$$

The second term above is then bounded as

$$\begin{aligned} \sum_{M \geq r} \sum_{\substack{\mathcal{I} \in \bar{\mathbf{I}}_{x,T,H} \\ m(\mathcal{I}; \tilde{\Psi}_{x,\ell}(\mathcal{I}))=M}} \mu_n^\mp(\mathcal{I}) &\leq \sum_{M \geq r} \sum_{\mathcal{J} \in \tilde{\Psi}_{x,\ell}(\bar{\mathbf{I}}_{x,T,H})} \mu_n^\mp(\mathcal{J}) \sum_{\substack{\mathcal{I} \in \tilde{\Psi}_{x,\ell}^{-1}(\mathcal{J}) \\ m(\mathcal{I}; \mathcal{J})=M}} e^{-(\beta-C)M} \\ &\leq \sum_{M \geq r} \tilde{C}_\Psi^M e^{-(\beta-C)M} \mu_n^\mp(\tilde{\Psi}_{x,\ell}(\bar{\mathbf{I}}_{x,T,H})). \end{aligned}$$

Then, noting by Proposition 4.5 that for every $\mathcal{I} \in \bar{\mathbf{I}}_{x,T,H}$, for every ℓ , $\Psi_{x,\ell}(\mathcal{I}) \in \bar{\mathbf{I}}_{x,T,H}$ implies that

$$\tilde{\Psi}_{x,\ell}(\bar{\mathbf{I}}_{x,T,H}) \subset \bar{\mathbf{I}}_{x,T,H}.$$

We then deduce that

$$\begin{aligned} \mu_n^\mp(\mathbf{m}(\mathcal{I}; \tilde{\Psi}_{x,\ell}(\mathcal{I})) \geq r, \bar{\mathbf{I}}_{x,T,H}) &\leq C e^{-(\beta-C-\log \tilde{C}_\Psi)r} \mu_n^\mp(\bar{\mathbf{I}}_{x,T,H}) \\ &\leq C e^{-(\beta-C-\log \tilde{C}_\Psi)r} \mu_n^\mp(\mathbf{I}_{x,T,H}), \end{aligned}$$

at which point dividing through by $\mu_n^\mp(\mathbf{I}_{x,T,H})$ implies the desired bound. \square

5. A refined submultiplicativity bound. In order to establish tightness for the centered maximum M_n under μ_n^\mp , we need to replace the approximate submultiplicativity bound on $\mu_n^\mp(\text{ht}(\mathcal{P}_x) \geq h_1 + h_2)$ obtained in [16]—which had an $O(\exp(\log^2(h_1 \vee h_2)))$ multiplicative error term—by one in which the multiplicative error is only $O(1)$. We will in fact show this with an error that is $1 + \varepsilon_\beta$ for some sequence $\varepsilon_\beta > 0$ that vanishes as $\beta \rightarrow \infty$.

Let us denote by A_h^x the event, measurable with respect to the configuration on $\mathcal{C}(\mathbb{Z}^2 \times \llbracket 0, h \rrbracket)$, given by

$$(5.1) \quad A_h^x = \left\{ \sigma : x + \left(0, 0, \frac{1}{2}\right) \xrightarrow[\mathbb{Z}^2 \times \llbracket 0, h \rrbracket]{+} \mathcal{L}_{h-\frac{1}{2}} \right\} \quad \text{for } x \in \mathcal{F}(\mathcal{L}_0),$$

so that, recalling the definition (1.1) of α_h , we have

$$\alpha_h = -\log \mu_{\mathbb{Z}^3}^\mp(A_h^o) \quad \text{where } o = \left(\frac{1}{2}, \frac{1}{2}, 0\right).$$

Further, let

$$(5.2) \quad E_h^x = \{ \sigma : \text{ht}(\mathcal{P}_x) \geq h \} \quad \text{for } x \in \mathcal{F}(\mathcal{L}_0).$$

We showed in [16], equation (6.3), (this will also follow from Claim 5.3 below) that, for n large and every $x \in \mathcal{L}_{0,n}$,

$$(5.3) \quad (1 - \varepsilon_\beta) \mu_n^\mp(A_h^x) \leq \mu_n^\mp(E_h^x) \leq (1 + \varepsilon_\beta) \mu_n^\mp(A_h^x),$$

where $\varepsilon_\beta \rightarrow 0$ as $\beta \rightarrow \infty$; the same also applies under $\mu_{\mathbb{Z}^3}^\mp$. Thus, for another such sequence ε_β ,

$$(5.4) \quad \tilde{\alpha}_h = -\log \mu_{\mathbb{Z}^3}^\mp(E_h^o) \quad \text{satisfies} \quad |\tilde{\alpha}_h - \alpha_h| \leq \varepsilon_\beta,$$

and we recall from Proposition 3.10 that, for some $C > 0$ and every $h \geq 1$,

$$4\beta - C \leq -\frac{1}{h} \log \mu_n^\mp(E_h^o) \leq 4\beta + e^{-4\beta} + \frac{C}{h}.$$

(The existence of $\alpha = \lim_{h \rightarrow \infty} \alpha_h/h$, as established in the prequel [16], implies that $4\beta - C \leq \alpha \leq 4\beta + e^{-4\beta}$. Our next results will rederive the limit and give it a more accurate description; see Corollary 5.2 below.) The inequalities in (5.3) tie the approximate submultiplicativity of $\{\mu_n^\mp(E_h^o)\}_{h \geq 1}$ to that of $\{\mu_{\mathbb{Z}^3}^\mp(A_h^o)\}_{h \geq 1}$ (equivalently, the super-additivity of $\tilde{\alpha}_h$ to that of α_h), which the following result establishes.

PROPOSITION 5.1. *There exists $\beta_0 > 0$ such that, for every $\beta > \beta_0$, every $h = h_1 + h_2$ for $h_1, h_2 \geq 1$, which may depend on n , and every $x, x_1, x_2 \in \mathcal{L}_{0,n}$ such that $d(x, \partial\Lambda_n) \gg h$ and $d(x_i, \partial\Lambda_n) \gg h$ for $i = 1, 2$,*

$$(5.5) \quad \mu_n^\mp(A_h^x) \leq (1 + \varepsilon_\beta) \mu_n^\mp(A_{h_1}^{x_1}) \mu_n^\mp(A_{h_2}^{x_2}),$$

where $\varepsilon_\beta > 0$ vanishes as $\beta \rightarrow \infty$. Consequently, for another such sequence $\varepsilon_\beta > 0$,

$$(5.6) \quad \mu_n^\mp(E_h^x) \leq (1 + \varepsilon_\beta) \mu_n^\mp(E_{h_1}^{x_1}) \mu_n^\mp(E_{h_2}^{x_2}).$$

In light of the preceding inequalities, the above proposition readily implies the following corollary.

COROLLARY 5.2. *There exists β_0 such that, for every $\beta > \beta_0$ and every $h_1, h_2 \geq 1$,*

$$(5.7) \quad \alpha_{h_1} + \alpha_{h_2} - \varepsilon_\beta \leq \alpha_{h_1+h_2} \leq \alpha_{h_1} + (4\beta + e^{-4\beta})h_2$$

for some $\varepsilon_\beta > 0$ that vanishes as $\beta \rightarrow \infty$. In particular, the limit $\alpha = \lim_{h \rightarrow \infty} \alpha_h/h$ ($= \lim_{h \rightarrow \infty} \tilde{\alpha}_h/h$) exists and satisfies $\alpha = \sup_h (\alpha_h - \varepsilon_\beta)/h$.

PROOF. For the left-hand side, fix any h_1, h_2 , take $x, x_1, x_2 = o$ and send $n \rightarrow \infty$ (whence $\mu_n^\mp(A_h^o) \rightarrow \mu_{\mathbb{Z}^3}^\mp(A_h^o)$ is by the weak convergence of μ_n^\mp to $\mu_{\mathbb{Z}^3}^\mp$, as established by Dobrushin, and the analogous fact for E_h^o follows from Corollary 5.6).

Similarly, for the right-hand side it suffices for us to prove, in the setting of Proposition 5.1, that

$$\mu_n^\mp(A_h^o) \geq \mu_n^\mp(A_{h_1}^o) e^{-(4\beta + e^{-4\beta})h_2},$$

and then send $n \rightarrow \infty$. Reveal the entire configuration σ on $\Lambda_n \cap \mathcal{L}_{<h_1}$ under $\mu_n^\mp(\cdot | A_{h_1}^o)$. On the event $A_{h_1}^o$, in the configuration we revealed, $o + (0, 0, \frac{1}{2})$ is connected by plus sites to $\mathcal{L}_{h_1 - \frac{1}{2}}$: call $Y - (0, 0, \frac{1}{2})$ an arbitrary site in the plus cluster of $o + (0, 0, \frac{1}{2})$ at height $h_1 - \frac{1}{2}$. Then, by a simple calculation (see, e.g., the proof in [16] of the similar left-hand of Proposition 3.10), independently of the configuration outside of the set of cells $\{Y + (0, 0, \frac{k}{2}) : k = 1, \dots, h_2\}$, the probability that those h_2 cells are all plus—and, therefore, A_h^o holds—is, at least, $e^{-(4\beta + 4e^{4\beta})h_2}$. \square

We will need the following comparison which will imply the inequality (5.3) that was stated above.

CLAIM 5.3. *There exist $\beta_0 > 0$ and a sequence $\varepsilon_\beta > 0$ vanishing as $\beta \rightarrow \infty$ such that, for all $\beta > \beta_0$, every n and every $x \in \mathcal{L}_{0,n}$,*

$$\mu_n^\mp(E_h^x | A_h^x) \geq 1 - \varepsilon_\beta \quad \text{and} \quad \mu_n^\mp(A_h^x | E_h^x) \geq 1 - \varepsilon_\beta.$$

PROOF. Let us begin with the first inequality: by definition, we have, for every $x \in \mathcal{L}_{0,n}$ and $h \geq 1$,

$$\left(A_h^x \cap \bigcap_{k \geq 1} \left\{ x - \left(0, 0, \frac{k}{2} \right) \in \sigma(\mathcal{I}) \right\} \right) \subset (E_h^x \cap A_h^x).$$

By Proposition 3.8, x is not interior to any wall of \mathcal{I} , and, therefore, $\bigcap_{k \geq 1} \{x - (0, 0, \frac{k}{2}) \in \sigma(\mathcal{I})\}$, except with probability ε_β going to zero as $\beta \rightarrow \infty$. Then, by the FKG inequality we deduce that

$$\mu_n^\mp(E_h^x \cap A_h^x) \geq \mu_n^\mp \left(A_h^x \cap \bigcap_{k \geq 1} \left\{ x - \left(0, 0, \frac{k}{2} \right) \in \sigma(\mathcal{I}) \right\} \right) \geq (1 - \varepsilon_\beta) \mu_n^\mp(A_h^x)$$

implying the desired bound after dividing through by $\mu_n^\mp(A_h^x)$.

Let us turn to the second inequality. We can expose the inner boundary of the plus $*$ -connected component of the cell-set $\sigma(\mathcal{I})$ under the measure $\mu_n^\mp(\cdot | E_h^x)$ as follows: reveal the entire minus $*$ -connected component of the boundary $\partial \Lambda \cap \mathcal{L}_{>0}$ by exposing the minus $*$ -connected component and all $*$ -adjacent (bounding) plus spins—exactly the plus (inner) boundary of the $*$ -connected plus component of $\sigma(\mathcal{I})$ together with (inner) boundaries of

finite plus bubbles in the minus phase. Since $x + (0, 0, \frac{1}{2})$ is plus in $\sigma(\mathcal{I})$ (on the event E_h^x), and, for every \mathcal{I} , the measure of $\mu_n^\mp(\cdot)$ on $\{v : \sigma(\mathcal{I})_v = +1\}$ stochastically dominates that induced by $\mu_{\mathbb{Z}^3}^+$ on $\{v : \sigma(\mathcal{I})_v = +1\}$, we have

$$\begin{aligned} \mu_n^\mp(A_h^x \mid E_h^x) &\geq \inf_{\mathcal{I} \in E_h^x} \mu_{\{v:\sigma(\mathcal{I})_v=+1\}}^+ \left(x + \left(0, 0, \frac{1}{2} \right) \xleftrightarrow[\mathcal{L}_{>0}^+]{\mathcal{S}_x} \right) \\ &\geq \mu_{\mathbb{Z}^3}^+ \left(x + \left(0, 0, \frac{1}{2} \right) \xleftrightarrow[\mathcal{L}_{>0}^+]{\infty} \right); \end{aligned}$$

this is seen to be at least $1 - \varepsilon_\beta$ by the classical Peierls argument. \square

The proof of Proposition 5.1 follows the same argument that was used to establish the weaker approximate submultiplicativity bound in [16], whereas here the error terms can be better tracked and controlled via the improved estimates on the shape of the pillar (in particular, the exponential tail on the size of the base conditioned on the pillar reaching height h vs. the bound in the prequel which had an extra $O(\log h)$ term). For completeness we include the full argument instead of only listing the needed modifications. We begin with recalling several decorrelation estimates for pillars which are needed for the proof.

PROPOSITION 5.4 ([15], [13], Lemma 5, as well as [6], Proposition 2.3). *There exists $C > 0$ such that, for every $\beta > \beta_0$, every $n \leq m$, for any subset $F \subset \mathcal{L}_{0,n}$,*

$$\| \mu_n^\mp((\mathcal{F}_y)_{y \in F} \in \cdot) - \mu_m^\mp((\mathcal{F}_y)_{y \in F} \in \cdot) \|_{\text{TV}} \leq C e^{-d(F, \partial \Lambda_n)/C}.$$

In particular, sending m to ∞ and via the tightness of $(\mathcal{F}_y)_{y \in F}$, this holds if we replace μ_m by $\mu_{\mathbb{Z}^3}^\mp$.

PROPOSITION 5.5 ([13], see also [6], Proposition 2.1). *There exist $\beta_0 > 0$ and $C > 0$ such that, for every $\beta > \beta_0$, every n and every two subsets $F_1, F_2 \subset \mathcal{L}_{0,n}$,*

$$\begin{aligned} &\| \mu_n^\mp((\mathcal{F}_y)_{y \in F_1} \in \cdot, (\mathcal{F}_z)_{z \in F_2} \in \cdot) - \mu_n^\mp((\mathcal{F}_y)_{y \in F_1} \in \cdot) \mu_n^\mp((\mathcal{F}_z)_{z \in F_2} \in \cdot) \|_{\text{TV}} \\ &\leq C e^{-d(F_1, F_2)/C}. \end{aligned}$$

Propositions 5.4–5.5 readily translate to similar estimates on the collections of pillars (see the short proofs of Corollaries 6.4 and 6.6 in [16], addressing the special cases where F, F_1, F_2 were balls of radius r about some fixed faces $x_n, y_n \in \mathcal{L}_{0,n}$ and following from the respective special cases of the above propositions).

COROLLARY 5.6 (see [16], Corollary 6.4). *There exist $\beta_0 > 0$ and $C > 0$ such that, for every $\beta > \beta_0$, every subset $F_n \subset \mathcal{L}_{0,n}$ and every subset $F_m \subset \mathcal{L}_{0,m}$, which is a horizontal translation of F_n ,*

$$\| \mu_n^\mp((\mathcal{P}_y)_{y \in F_n} \in \cdot) - \mu_m^\mp((\mathcal{P}_y)_{y \in F_m} \in \cdot) \|_{\text{TV}} \leq C e^{-(d(F_n, \partial \Lambda_n) \vee d(F_m, \partial \Lambda_m))/C}.$$

COROLLARY 5.7 (see [16], Corollary 6.6). *There exist $\beta_0 > 0$ and some $C > 0$ such that, for every $\beta > \beta_0$, every n and every two subsets $F_1, F_2 \subset \mathcal{L}_{0,n}$,*

$$\| \mu_n^\mp((\mathcal{P}_y)_{y \in F_1} \in \cdot, (\mathcal{P}_y)_{y \in F_2} \in \cdot) - \mu_n^\mp((\mathcal{P}_y)_{y \in F_1} \in \cdot) \mu_n^\mp((\mathcal{P}_y)_{y \in F_2} \in \cdot) \|_{\text{TV}} \leq C e^{-d(F_1, F_2)/C}.$$

PROOF OF PROPOSITION 5.1. Recall that $h = h_1 + h_2$, and suppose without loss of generality that $h_1 \geq h_2$. Define the vertical shift of the event A_h^x , for a given vertex x , as

$$\theta_{h_1} A_{h_2}^x = \left\{ \sigma : x + \left(0, 0, h_1 + \frac{1}{2} \right) \xleftrightarrow[\mathbb{Z}^2 \times \llbracket h_1, h_1 + h_2 \rrbracket]{+} \mathcal{L}_{h_1 + h_2 - \frac{1}{2}} \right\},$$

noting that $\mu_n^{\mp}(\theta_{h_1} A_{h_2}^x) = \mu_n^{\mp, -h_1}(A_{h_2}^x)$, where we denote by $(\mp, -m)$ boundary conditions are those that are plus on $\partial\Lambda \cap \mathcal{L}_{<-m}$ and minus on $\partial\Lambda \cap \mathcal{L}_{>-m}$. Hence, for every n , by monotonicity in boundary conditions,

$$\mu_n^{\mp, -h_1}(A_{h_2}^x) \leq \mu_n^{\mp}(A_{h_2}^x).$$

A naive approach to establishing submultiplicativity would be to expose the plus $*$ -component of $x + (0, 0, \frac{1}{2})$ in the slab $\mathbb{Z}^2 \times [0, h_1]$, wherein the measurable event $A_{h_1}^x$ guarantees that $x + (0, 0, \frac{1}{2})$ is $*$ -connected to $\mathcal{L}_{h_1 - \frac{1}{2}}$, and, as the tip of the $*$ -component at $\mathcal{L}_{h_1 - \frac{1}{2}}$ is now situated in the minus phase, the conditional probability of A_h^x should be at most that of the unconditional $A_{h_2}^x$. However, revealing the plus $*$ -component introduces some positive information (the connection event $A_{h_1}^x$ is increasing) along with negative information (e.g., minus spins along its boundary). We will control this using our new estimates on the shape of the pillar.

Denote by \mathcal{A} the $*$ -connected plus component of $x + (0, 0, \frac{1}{2})$ in $[-n_h, n_h]^2 \times [0, h_1]$ (noting that the event $A_{h_1}^x$ merely says that \mathcal{A} intersects the slab $\mathcal{L}_{h_1 - \frac{1}{2}}$). An important fact, which we will use later on, is that, on the event that $\mathcal{P}_x \neq \emptyset$, this plus $*$ -component \mathcal{A} is a subset of the plus sites in \mathcal{P}_x .

DEFINITION 5.8. Let \mathcal{A} denote the set of possible realizations of \mathcal{A} that satisfy the following properties:

1. The intersection of \mathcal{A} with $\mathcal{L}_{\frac{1}{2}}$ is the single cell whose lower bounding face is x .
2. The intersection of \mathcal{A} with $\mathcal{L}_{h_1 - \frac{1}{2}}$ is a single cell; denote its upper bounding face by $Y \in \mathcal{L}_{h_1}$.
3. We further have $d(x, \rho(Y)) \leq d(x, \partial\Lambda_n)/2$.

Let Γ denote the (neither increasing nor decreasing) event $\{\mathcal{A} \in \mathcal{A}\}$, noting that $\Gamma \subset A_{h_1}^x$.

CLAIM 5.9. In the setting of Proposition 5.1, there exists a sequence ε_β going to zero as $\beta \rightarrow \infty$ such that

$$\mu_n^{\mp}(A_h^x, \Gamma) \leq (1 + \varepsilon_\beta) \mu_n^{\mp}(A_{h_1}^{x_1}) \mu_n^{\mp}(A_{h_2}^{x_2}).$$

CLAIM 5.10. In the setting of Proposition 5.1, there exists a sequence ε_β going to zero as $\beta \rightarrow \infty$ such that

$$\mu_n^{\mp}(A_h^x) \leq (1 + \varepsilon_\beta) \mu_n^{\mp}(A_h^x, \Gamma).$$

PROOF OF CLAIM 5.9. Set

$$r = \min\{d(x_1, \partial\Lambda_n), d(x_2, \partial\Lambda_n), d(x, \partial\Lambda_n)\},$$

and recall that $r \gg h$ by assumption; thus, Corollary 5.6 (for $m = n$) implies that, for each $i = 1, 2$,

$$\|\mu_n^{\mp}(\mathcal{P}_x \in \cdot) - \mu_n^{\mp}(\mathcal{P}_{x_i} \in \cdot)\|_{TV} \leq C \exp(-r/C)$$

for some C independent of β . By relating the events $A_{h_i}^x$ and $E_{h_i}^x$ via (5.3), we then obtain that

$$\begin{aligned} \mu_n^{\mp}(A_{h_i}^x) &\leq (1 + \varepsilon_\beta) \mu_n^{\mp}(E_{h_i}^x) \leq (1 + \varepsilon_\beta) \mu_n^{\mp}(E_{h_i}^{x_i}) + C \exp(-r/C) \\ (5.8) \quad &\leq (1 + \varepsilon_\beta)^2 \mu_n^{\mp}(A_{h_i}^{x_i}) + C \exp(-r/C) \\ &\leq (1 + o(1))(1 + \varepsilon_\beta)^2 \mu_n^{\mp}(A_{h_i}^{x_i}) \end{aligned}$$

using $\mu_n^\mp(A_{h_i}^{x_i}) \geq \exp(-4\beta + e^{-4\beta}h_i)$ and $r \gg h_i$ by our hypothesis. Thus, the claim will follow once we show that, for some other sequence $\varepsilon_\beta > 0$ that vanishes as $\beta \rightarrow \infty$,

$$(5.9) \quad \mu_n^\mp(A_h^x, \Gamma) \leq (1 + \varepsilon_\beta)\mu_n^\mp(A_{h_1}^x)\mu_n^\mp(A_{h_2}^x).$$

Recall that $\Gamma \subset A_{h_1}^x$, whence

$$\mu_n^\mp(A_h^x, \Gamma) = \mu_n^\mp(\Gamma)\mu_n^\mp(A_h^x | \Gamma) \leq \mu_n^\mp(A_{h_1}^x)\mu_n^\mp(A_h^x | \Gamma),$$

so in order to establish (5.9) it remains to show that

$$(5.10) \quad \mu_n^\mp(A_h^x | \Gamma) \leq (1 + \varepsilon_\beta)\mu_n^\mp(A_{h_2}^x).$$

Denoting by \mathbb{E}_Γ expectation w.r.t. $\mu_n^\mp(\cdot | \Gamma)$, whereby \mathcal{A} accepts values in \mathcal{A} , we have that

$$\mu_n^\mp(A_h^x | \Gamma) = \mathbb{E}_\Gamma[\mu_n^\mp(A_h^x | \mathcal{A})] = \mathbb{E}_\Gamma[\mu_n^\mp(\theta_{h_1}A_{h_2}^Y | \mathcal{A})],$$

since Γ stipulates that $Y - (0, 0, \frac{1}{2})$ is the unique cell in the intersection of \mathcal{A} with $\mathcal{L}_{h_1 - \frac{1}{2}}$, whence $Y + (0, 0, \frac{1}{2})$ must then be $*$ -connected to $\mathcal{L}_{h - \frac{1}{2}}$ in $\mathcal{C}(\mathbb{Z}^2 \times \llbracket h_1, h \rrbracket)$ in order for A_h^x to occur.

Reveal $\mathcal{A} \in \mathcal{A}$, the $*$ -connected plus component of $x + (0, 0, \frac{1}{2})$, only in the slab $\mathcal{C}(\llbracket -n, n \rrbracket^2 \times \llbracket 0, h_1 \rrbracket)$. A key observation now is that the boundary spins of the $*$ -component revealed in this manner are plus at $x + (0, 0, \frac{1}{2})$ and $Y - (0, 0, \frac{1}{2})$ and minus on all sites in $\Lambda_n \setminus \mathcal{A}$ that are $*$ -adjacent to \mathcal{A} . Indeed, the boundary conditions on $\partial\Lambda_n$ are all minus between heights $[0, h_1]$, and our definition of \mathcal{A} forces every $*$ -adjacent spin in $\mathcal{L}_{>0} \cap \mathcal{L}_{<h_1}$ to be minus, except at $x + (0, 0, \frac{1}{2})$ and $Y - (0, 0, \frac{1}{2})$ (those are plus as per \mathcal{A}). Denoting by $\partial\mathcal{A}$ these boundary spins and by $(\mp, \partial\mathcal{A})$ their addition to our Dobrushin boundary conditions, the domain Markov property implies that

$$\mathbb{E}_\Gamma[\mu_n^\mp(\theta_{h_1}A_{h_2}^Y | \mathcal{A})] = \mathbb{E}_\Gamma[\mu_n^{\mp, \partial\mathcal{A}}(\theta_{h_1}A_{h_2}^Y)] \leq \sup_{\mathcal{A} \in \mathcal{A}} \mu_n^{\mp, \partial\mathcal{A}}(\theta_{h_1}A_{h_2}^Y).$$

For a fixed $\mathcal{A} \in \mathcal{A}$ (hence, fixed Y), in view of the above fact that $\partial\mathcal{A}$ includes plus spins only at $x + (0, 0, \frac{1}{2})$ and $Y - (0, 0, \frac{1}{2})$, the FKG inequality w.r.t. the Ising measure conditioned on $\sigma_{x+(0,0,\frac{1}{2})} = \sigma_{Y-(0,0,\frac{1}{2})} = +1$ enables us to omit the conditioning on its minus spins and obtain that

$$\begin{aligned} \sup_{\mathcal{A} \in \mathcal{A}} \mu_n^{\mp, \partial\mathcal{A}}(\theta_{h_1}A_h^Y) &\leq \sup_{\mathcal{A} \in \mathcal{A}} \mu_n^\mp(\theta_{h_1}A_h^Y | \sigma_{x+(0,0,\frac{1}{2})} = \sigma_{Y-(0,0,\frac{1}{2})} = +1) \\ &\leq \sup_{\mathcal{A} \in \mathcal{A}} \mu_n^{\mp, -h_1}(A_{h_2}^{\rho(Y)} | \sigma_{x+(0,0,-h_1+\frac{1}{2})} = \sigma_{\rho(Y)-(0,0,\frac{1}{2})} = +1) \end{aligned}$$

by translation. Another application of FKG—now for monotonicity in boundary conditions—allows us to move from $\mu_n^{\mp, -h_1}$ to μ_n^\mp and to conclude that the last expression is, at most,

$$\sup_{\mathcal{A} \in \mathcal{A}} \mu_n^\mp(A_{h_2}^{\rho(Y)} | \sigma_{x+(0,0,-h_1+\frac{1}{2})} = \sigma_{\rho(Y)-(0,0,\frac{1}{2})} = +1) \leq (1 + \varepsilon_\beta)^2 \sup_{\mathcal{A} \in \mathcal{A}} \mu_n^\mp(A_{h_2}^{\rho(Y)}),$$

where the last inequality is justified as follows. For a fixed face $u \in \mathcal{L}_{0,n}$ (here, we would take $u = \rho(Y)$ for a worst-case realization of \mathcal{A}), if we denote $B_x = \{\sigma_{x+(0,0,-h_1+\frac{1}{2})} = +1\}$ and $B_u = \{\sigma_{u-(0,0,\frac{1}{2})} = +1\}$, then $\mu_n^\mp(A_{h_2}^u | B_x, B_u) \leq \mu_n^\mp(A_{h_2}^u) / \mu_n^\mp(B_x, B_u)$; now, $\mu_n^\mp(B_x, B_u) \geq \mu_n^\mp(B_x)\mu_n^\mp(B_u)$ by FKG, thus it remains to show that each of the events B_x and B_u has probability at least $1 - \varepsilon_\beta$ under μ_n^\mp . By the results of Dobrushin (see, e.g., Proposition 3.8), if $z = (z_1, z_2, z_3)$ is a fixed point $z_3 < 0$, then $(z_1, z_2, 0)$ has no walls of the interface \mathcal{I} nesting it except with probability ε_β . In particular, $z \in \sigma(\mathcal{I})$, whence a Peierls argument shows that with probability $1 - \varepsilon_\beta$ its spin is plus.

Finally, for each $\mathcal{A} \in \mathcal{A}$, deterministically, $d(x, \rho(Y)) \leq d(x, \partial\Lambda_n)/2$, so by the triangle inequality

$$d(\rho(Y), \partial\Lambda_n) \geq d(x, \partial\Lambda_n) - d(x, \rho(Y)) \geq d(x, \partial\Lambda_n)/2 \gg h$$

by assumption. Thus, the same argument used to compare x to x_i in (5.8) shows that

$$\mu_n^\mp(A_{h_2}^{\rho(Y)}) \leq (1 + \varepsilon_\beta)\mu_n^\mp(A_{h_2}^x),$$

establishing (5.10) and thus concluding the proof. \square

PROOF OF CLAIM 5.10. Writing

$$\mu_n^\mp(\Gamma, A_h^x \mid E_h^x) \leq \frac{\mu_n^\mp(\Gamma, A_h^x)}{\mu_n^\mp(E_h^x)} = \frac{\mu_n^\mp(A_h^x)}{\mu_n^\mp(E_h^x)} \mu_n^\mp(\Gamma \mid A_h^x) \leq (1 + \varepsilon_\beta)\mu_n^\mp(\Gamma \mid A_h^x),$$

with the last inequality by (5.3), it remains to show $\mu_n^\mp(\Gamma, A_h^x \mid E_h^x) \geq 1 - \varepsilon'_\beta$ for some other sequence $\varepsilon'_\beta > 0$ vanishing as $\beta \rightarrow \infty$, which will altogether imply that $\mu_n^\mp(\Gamma, A_h^x)/\mu_n^\mp(A_h^x) \geq (1 - \varepsilon'_\beta)/(1 + \varepsilon_\beta)$, as required. Using that $\mu_n^\mp((A_h^x)^c \mid E_h^x) \leq \varepsilon_\beta$ by Claim 5.3, we have

$$\mu_n^\mp(\Gamma^c \cup (A_h^x)^c \mid E_h^x) \leq \mu_n^\mp((A_h^x)^c \mid E_h^x) + \mu_n^\mp(\Gamma^c, A_h^x \mid E_h^x) \leq \varepsilon_\beta + \mu_n^\mp(\Gamma^c, A_h^x \mid E_h^x),$$

and it remains to bound the last term in the right-hand by ε_β . Examining the criteria for Γ in Definition 5.8, observe that A_h^x implies (through the inclusion $A_{h_1}^x \supset A_h^x$) that \mathcal{A} intersects both $\mathcal{L}_{\frac{1}{2}}$ and $\mathcal{L}_{h_1 - \frac{1}{2}}$. Hence,

$$\Gamma^c \cap A_h^x = (B_1 \cup B_2 \cup B_3) \cap A_h^x \quad \text{where} \quad \begin{cases} B_1 = \{|\mathcal{C}(\mathcal{A} \cap \mathcal{L}_{\frac{1}{2}})| \geq 2\}, \\ B_2 = \{|\mathcal{C}(\mathcal{A} \cap \mathcal{L}_{h_1 - \frac{1}{2}})| \geq 2\}, \\ B_3 = \left\{d(x, \rho(Y)) > \frac{1}{2}d(x, \partial\Lambda_n)\right\}. \end{cases}$$

Since $E_h^x = \mathbf{I}_{x,0,h}$, the aforementioned fact that \mathcal{A} is a subset of the plus spins in \mathcal{P}_x implies that

$$\mu_n^\mp(B_1 \mid E_h^x) \leq \mu_n^\mp(\mathcal{B}_x \neq \emptyset \mid \mathbf{I}_{x,0,h})$$

and

$$\mu_n^\mp(B_2 \mid E_h^x) \leq \mu_n^\mp(|\mathcal{C}(\mathcal{P}_x \cap \mathcal{L}_{h_1 - \frac{1}{2}})| \geq 2 \mid \mathbf{I}_{x,0,h}).$$

Theorem 4.1 and Proposition 4.2, respectively, show (using the hypothesis $h \ll d(x, \partial\Lambda_n)$) that these two probabilities are, at most, $\exp(-(\beta - C))$.

Finally, the event B_3 implies that $d(x, \rho(y)) > \frac{1}{2}d(x, \partial\Lambda_n)$ for some face $y \in \mathcal{P}_x$. By (3.2), this implies $\text{diam}(\mathcal{B}_x) + \frac{1}{4}\mathfrak{m}(\mathcal{S}_x) > \frac{1}{2}d(x, \partial\Lambda_n)$, whence by Proposition 3.12,

$$\begin{aligned} \mu_n^\mp(B_3 \mid \mathbf{I}_{x,0,h}) &\leq \mu_n^\mp\left(\text{diam}(\mathcal{B}_x) + \frac{1}{4}\mathfrak{m}(\mathcal{S}_x) > \frac{1}{2}d(x, \partial\Lambda_n) \mid \mathbf{I}_{x,0,h}\right) \\ &\leq C e^{-(\beta - C)d(x, \partial\Lambda_n)} = o(1). \end{aligned}$$

Combined, we have that $\mu_n^\mp(\Gamma^c \cap A_h^x) \leq 2 \exp(-(\beta - C)) + o(1)$, as needed. \square

Combining Claims 5.9–5.10 concludes the proof. \square

6. Tightness and exponential tails of the maximum. In this section we prove left and right exponential tails for $M_n - m_n^*$, as stated in the next proposition:

PROPOSITION 6.1. *There exist $\beta_0 > 0$ and a sequence $\varepsilon_\beta > 0$ vanishing as $\beta \rightarrow \infty$ such that the following holds for all $\beta > \beta_0$. Letting $\tilde{\alpha}_h$ be as in (5.4) and m_n^* be as in (1.2), for every $1 \leq \ell \leq \sqrt{\log n}$,*

$$\begin{aligned} \mu_n^\mp(M_n \geq m_n^* + \ell) &\leq (1 + \varepsilon_\beta) \exp(2\beta - \tilde{\alpha}_\ell), \\ \mu_n^\mp(M_n < m_n^* - \ell) &\leq (1 + \varepsilon_\beta) \exp(2\beta - \tilde{\alpha}_\ell). \end{aligned}$$

Before proving this result, we will establish some preliminary estimates. Recalling that m_n^* is the first h such that $\alpha_h > 2 \log(2n) - 2\beta$, the relation between $\alpha_h, \tilde{\alpha}_h$ in (5.4) and the bound $\alpha_{h+1} \leq \alpha_h + 4\beta + \varepsilon_\beta$ by (5.7) together imply that, for the sequence $\tilde{\alpha}_h$, we have

$$(6.1) \quad 2 \log(2n) - 2\beta - \varepsilon_\beta \leq \tilde{\alpha}_{m_n^*} \leq 2 \log(2n) + 2\beta + \varepsilon_\beta.$$

Next, recall that $E_h^x = \{\text{ht}(\mathcal{P}_x) \geq h\}$ (see (5.2)), so that $\{M_n \geq h\} = \bigcup_{x \in \mathcal{L}_{0,n}} E_h^x$; we will separate the analysis of E_h^x for x near and away from $\partial\Lambda_n$ as follows. Define the interior of $\mathcal{L}_{0,n}$,

$$(6.2) \quad \mathcal{L}_{0,n}^- = \mathcal{F}(\llbracket -(n - \log^2 n), n - \log^2 n \rrbracket^2 \times \{0\}),$$

and observe that, by Corollary 5.6, for $x \in \mathcal{L}_{0,n}^-$ we can couple $\mu_n^\mp(\mathcal{P}_x \in \cdot)$ to $\mu_{\mathbb{Z}^3}^\mp(\mathcal{P}_o \in \cdot)$ and find that, for some fixed $c > 0$,

$$(6.3) \quad \mu_n^\mp(E_h^x) = \mu_{\mathbb{Z}^3}^\mp(E_h^o) + O(e^{-c \log^2 n}) = (1 + o(1))e^{-\tilde{\alpha}_h} \quad \text{if } 1 \leq h \ll \log^2 n,$$

absorbing the $O(e^{-c \log^2 n})$ as $\tilde{\alpha}_h \leq \alpha_h + \varepsilon_\beta$ by Corollary 5.2, so $\mu_n^\mp(E_h^o) = e^{-\tilde{\alpha}_h} = e^{-o(\log^2 n)}$ for $h \ll \log^2 n$.

Further, define, for $h \in \mathbb{Z}_+$, the event

$$(6.4) \quad G_h^x = E_h^x \cap \{\mathcal{B}_x = \emptyset\}$$

(noting that this event also implies the event A_h^x from (5.1)), and let

$$Z_h = \sum_{x \in \mathcal{L}_{0,n}^-} \mathbf{1}\{G_h^x\}.$$

CLAIM 6.2. *There exist $\beta_0 > 0$ and a sequence $\varepsilon_\beta > 0$ vanishing as $\beta \rightarrow \infty$ such that, for every $\beta > \beta_0$, the following hold. If $1 \leq h \ll \log^2 n$ then, for every $x \in \mathcal{L}_{0,n}^-$ and large enough n ,*

$$(6.5) \quad \mu_n^\mp(G_h^x) \geq (1 - \varepsilon_\beta) \mu_n^\mp(E_h^x).$$

Consequently, if $1 \leq h = m_n^* - \ell$ for $\ell > 0$, then

$$\mathbb{E}[Z_h] \geq (1 - \varepsilon_\beta) e^{\tilde{\alpha}_\ell - 2\beta}.$$

PROOF. For every $x \in \mathcal{L}_{0,n}^-$, we have

$$\mu_n^\mp((G_h^x)^c \mid E_h^x) = \mu_n^\mp(\mathcal{B}_x \neq \emptyset \mid E_h^x) \leq C e^{-2\beta},$$

by part (a) of Theorem 4.1, where we took $r = 1$ and used that $E_h^x = \mathbf{I}_{x,0,h}$ and $h \ll \log^2 n \leq d(x, \partial\Lambda_n)$ since $x \in \mathcal{L}_{0,n}^-$; this yields (6.5). For the second part of the claim, notice that

$1 \leq h < m_n^* = (\frac{2}{\alpha} + o(1)) \log n$ by Corollary 5.2. By (6.3) (now $h = O(\log n)$), we have $\mu_n^\mp(E_h^x) = (1 + o(1))e^{-\tilde{\alpha}h}$, and the super-additivity in Corollary 5.2 shows that

$$\tilde{\alpha}_h = \tilde{\alpha}_{m_n^* - \ell} \leq \tilde{\alpha}_{m_n^*} - \tilde{\alpha}_\ell + \varepsilon_\beta \leq 2 \log(2n) - \tilde{\alpha}_\ell + 2\beta + 2\varepsilon_\beta,$$

using (6.1) for the last inequality. Combining these, while noting that $|\mathcal{L}_{0,n}^-| = (1 - o(1))|\mathcal{L}_{0,n}| = (4 - o(1))n^2$, we obtain that the expectation of Z_h under μ_n^\mp satisfies

$$\mathbb{E}[Z_h] \geq (1 - o(1))(1 - \varepsilon_\beta)e^{\tilde{\alpha}_\ell - 2\beta - 2\varepsilon_\beta} \geq (1 - \varepsilon'_\beta)e^{\tilde{\alpha}_\ell - 2\beta},$$

for some other sequence $\varepsilon'_\beta > 0$ vanishing as $\beta \rightarrow \infty$. \square

CLAIM 6.3. *There exist $\beta_0 > 0$ and a sequence $\varepsilon_\beta > 0$ vanishing as $\beta \rightarrow \infty$ such that, for every $\beta > \beta_0$, every $1 \leq h \ll \log^2 n$ and every $x \neq y \in \mathcal{L}_{0,n}^-$ such that $d(x, y) \leq \log^2 n$, if n is large enough then*

$$\mu_n^\mp(G_h^y | G_h^x) \leq (1 + \varepsilon_\beta)e^{6\beta} \mu_n^\mp(G_h^y).$$

PROOF. We use a similar revealing procedure to that used in the proof of submultiplicativity above to reveal \mathcal{P}_x , without obtaining too much positive information about \mathcal{P}_y . Let \mathcal{A}_x be the $*$ -connected plus component of $x + (0, 0, \frac{1}{2})$ in $\mathcal{L}_{>0}$, and let \mathcal{A}_y be the $*$ -connected plus component of $y + (0, 0, \frac{1}{2})$ in $\mathcal{L}_{>0}$. If we reveal \mathcal{A}_x on the event G_h^x , we expose the plus $*$ -component \mathcal{A}_x along with all $*$ -adjacent (bounding) minus spins in $\mathcal{L}_{>0}$. On the event G_h^x , whereby the first cut-point of \mathcal{P}_x is $x + (0, 0, \frac{1}{2})$, the exterior boundaries of \mathcal{P}_x and \mathcal{A}_x coincide, and, therefore, the event G_h^x is measurable with respect to the set of sites revealed in this manner. As such, we can express

$$\mu_n^\mp(G_h^y | G_h^x) \leq \sup_{\mathcal{A}_x \in G_h^x} \mu_n^\mp(E_h^y | \mathcal{A}_x).$$

The boundary sites revealed by $\mathcal{A}_x \in G_h^x$ are all minus, except a single plus site at $x + (0, 0, \frac{1}{2})$, and so by the FKG inequality and the fact that E_h^y is an increasing event, this is, at most,

$$\mu_n^\mp(E_h^y | \sigma_{x+(0,0,\frac{1}{2})} = +1) \leq e^{6\beta} \mu_n^\mp(E_h^y).$$

The fact that $\mu_n^\mp(G_h^y) \geq (1 - \varepsilon_\beta)\mu_n^\mp(E_h^y)$ by (6.5) concludes the proof. \square

CLAIM 6.4. *There exists $\beta_0 > 0$ such that, for every $\beta > \beta_0$, there is some $C > 0$ such that, for every $x, y \in \mathcal{L}_{0,n}^-$, we have*

$$|\mu_n^\mp(G_h^x)\mu_n^\mp(G_h^y) - \mu_n^\mp(G_h^x, G_h^y)| \leq Ce^{-d(x,y)/C}.$$

PROOF. Notice that the pair of events G_h^x and G_h^y are measurable with respect to the pair of walls W_x, W_y . This is because the bounding faces of the spine \mathcal{S}_x (resp., \mathcal{S}_y) are all part of the same wall as shown in Claim 4.9, and the wall W_x (resp., W_y) contains the four bounding faces of $x + (0, 0, \frac{1}{2})$ (resp., $y + (0, 0, \frac{1}{2})$). As such, we can bound the difference above as

$$|\mu_n^\mp(G_h^x)\mu_n^\mp(G_h^y) - \mu_n^\mp(G_h^x, G_h^y)| \leq \|\mu_n^\mp(W_x \in \cdot)\mu_n^\mp(W_y \in \cdot) - \mu_n^\mp(W_x \in \cdot, W_y \in \cdot)\|_{TV},$$

which is, at most, $Ce^{-d(x,y)/C}$ by Proposition 5.5. \square

6.1. *Exponential tails for the maximum.* We are now ready to deduce that the centered maximum has left and right exponential tails (and is, therefore, tight).

PROOF OF PROPOSITION 6.1. We begin with the right tail. Letting

$$h = m_n^* + \ell \quad \text{for } 1 \leq \ell \leq (1/\beta) \log n$$

(n.b. we could have taken here $\ell \leq (C/\beta) \log n$ for any absolute constant C), we have

$$\begin{aligned} \mu_n^\mp(M_n \geq m_n^* + \ell) &\leq \sum_{x \in \mathcal{L}_{0,n} \setminus \mathcal{L}_{0,n}^-} \mu_n^\mp(E_h^x) + \sum_{x \in \mathcal{L}_{0,n}^-} \mu_n^\mp(E_h^x) \\ &\leq |\mathcal{L}_{0,n} \setminus \mathcal{L}_{0,n}^-| e^{-(4\beta-C)h} + |\mathcal{L}_{0,n}^-| (1 - o(1)) e^{-\tilde{\alpha}h}, \end{aligned}$$

using Proposition 3.10 for the first sum and (6.3) for the second one. By Corollary 5.2, we have that

$$\tilde{\alpha}h - (4\beta - C)h \leq (\alpha + \varepsilon_\beta - 4\beta + C)h \leq (C + 2\varepsilon_\beta)h \leq \varepsilon'_\beta \log n,$$

where the last inequality used the assumption on ℓ and the facts that $m_n^* = (2/\alpha + o(1)) \log n$ and $\alpha > 4\beta - C$. When combined with the fact that $|\mathcal{L}_{0,n} \setminus \mathcal{L}_{0,n}^-|/|\mathcal{L}_{0,n}| = O(\frac{\log^2 n}{n})$, this implies that

$$\frac{|\mathcal{L}_{0,n} \setminus \mathcal{L}_{0,n}^-| e^{-(4\beta-C)h}}{|\mathcal{L}_{0,n}^-| e^{-\tilde{\alpha}h}} \leq n^{-1+\varepsilon'_\beta+o(1)},$$

which, in light of the first inequality in the proof, shows that, for β large enough (so as to have $\varepsilon'_\beta < 1$),

$$\begin{aligned} \mu_n^\mp(M_n \geq m_n^* + \ell) &\leq (1 + o(1)) |\mathcal{L}_{0,n}^-| e^{-\tilde{\alpha}h} \leq (1 + \varepsilon_\beta) |\mathcal{L}_{0,n}^-| \exp(-\tilde{\alpha}m_n^* - \tilde{\alpha}\ell) \\ &\leq (1 + \varepsilon_\beta) e^{2\beta+\varepsilon_\beta-\tilde{\alpha}\ell} \leq (1 + \varepsilon'_\beta) e^{2\beta-\tilde{\alpha}\ell}, \end{aligned}$$

using Proposition 5.1 in the first line and (6.1) in the second line. This establishes the right tail.

REMARK 6.5. One can extend the right tail bound to hold for all ℓ (as opposed to $\ell \leq (1/\beta) \log n$)—albeit with a suboptimal rate: there exists some $C > 0$ such that

$$(6.6) \quad \mu_n^\mp(M_n \geq m_n^* + \ell) \leq C e^{-(2\beta-C)\ell} \quad \text{for all } \ell > 0.$$

Indeed, consider $\ell > (1/\beta) \log n$ (having already established the desired right tail for smaller values of ℓ). The bound $\mu_n^\mp(E_h^x) \leq C \exp(-(4\beta - C)h)$ in Proposition 3.8 holds uniformly over all $x \in \mathcal{L}_{0,n}$, and so

$$\mu_n^\mp(E_{m_n^*+\ell}^x) \leq \mu_n^\mp(E_\ell^x) \leq C e^{-(4\beta-C)\ell} \leq C n^{-2} e^{-(2\beta-C)\ell},$$

whence

$$\mu_n^\mp(M_n \geq m_n^* + \ell) \leq |\mathcal{L}_{0,n}| C n^{-2} e^{-(2\beta-C)\ell} \leq 4C e^{-(2\beta-C)\ell},$$

as claimed.

Let us now turn to the lower tail for M_n . Let

$$h = m_n^* - \ell \quad \text{for } 1 \leq \ell \leq \sqrt{\log n}.$$

Since $Z_h > 0$ implies E_h^x for some $x \in \mathcal{L}_{0,n}^-$ and, in particular, that $M_n \geq h$, it will suffice to establish an appropriate upper bound on $\mu_n^\mp(Z_h = 0)$, which we will infer from a second moment calculation. Write

$$\begin{aligned} \mathbb{E}[Z_h^2] &= \sum_{x,y \in \mathcal{L}_{0,n}^-} \mu_n^\mp(G_h^x, G_h^y) \\ &\leq \sum_{x \in \mathcal{L}_{0,n}^-} \mu_n^\mp(G_h^x) + \sum_{x \in \mathcal{L}_{0,n}^-} \sum_{y \in \mathcal{L}_{0,n}^- \cap B(x, \log^2 n): y \neq x} \mu_n^\mp(G_h^x) \mu_n^\mp(G_h^y | G_h^x) \\ &\quad + \sum_{x \in \mathcal{L}_{0,n}^-} \sum_{y \in \mathcal{L}_{0,n}^- \setminus B(x, \log^2 n)} (\mu_n^\mp(G_h^x) \mu_n^\mp(G_h^y) + |\mu_n^\mp(G_h^x) \mu_n^\mp(G_h^y) - \mu_n^\mp(G_h^x, G_h^y)|). \end{aligned}$$

Denoting these three summations by Ξ_1, Ξ_2, Ξ_3 (in order), we first observe that Ξ_1 is exactly $\mathbb{E}[Z_h]$. For the second summation we apply Claim 6.3, yielding

$$\Xi_2 \leq 4n^2 \log^4 n (1 + \varepsilon_\beta) e^{6\beta} \sup_{x,y \in \mathcal{L}_{0,n}^-} \mu_n^\mp(G_h^x) \mu_n^\mp(G_h^y) \leq n^{2+o(1)} \sup_{x \in \mathcal{L}_{0,n}^-} \mu_n^\mp(E_h^x)^2.$$

Using (6.3) (here, $h \leq m_n^* = O(\log n)$), we have $\mu_n^\mp(E_h^x) = (1 + o(1))e^{-\tilde{\alpha}_h}$, and $|\tilde{\alpha}_h - \tilde{\alpha}_{m_n^*}| = O(\ell) = O(\sqrt{\log n})$ by Corollary 5.2, whence $e^{-\tilde{\alpha}_h} = n^{-2+o(1)}$ by (6.1). Combined with the last equation,

$$\Xi_2 \leq n^{-2+o(1)} = o(1).$$

Finally, for the last summation,

$$\sum_{x \in \mathcal{L}_{0,n}^-} \sum_{y \in \mathcal{L}_{0,n}^- \setminus B(x, \log^2 n)} \mu_n^\mp(G_h^x) \mu_n^\mp(G_h^y) \leq \mathbb{E}[Z_h]^2,$$

whereas, by Claim 6.4,

$$\sum_{x \in \mathcal{L}_{0,n}^-} \sum_{y \in \mathcal{L}_{0,n}^- \setminus B(x, \log^2 n)} |\mu_n^\mp(G_h^x) \mu_n^\mp(G_h^y) - \mu_n^\mp(G_h^x, G_h^y)| \leq C |\mathcal{L}_{0,n}^-|^2 e^{-\log^2 n / C} = o(1),$$

and we deduce that $\Xi_3 \leq \mathbb{E}[Z_h]^2 + o(1)$. Putting all of these together, we obtain by the Paley–Zygmund inequality that

$$\mu_n^\mp(Z_h > 0) \geq \frac{\mathbb{E}[Z_h]^2}{\mathbb{E}[Z_h^2]} \geq \frac{\mathbb{E}[Z_h]^2}{\mathbb{E}[Z_h]^2 + \mathbb{E}[Z_h] + o(1)}.$$

As $\mathbb{E}[Z_h] \geq (1 - \varepsilon_\beta)e^{\tilde{\alpha}_\ell - 2\beta}$ by Claim 6.2, we see that

$$\mu_n^\mp(Z_h = 0) \leq \frac{1 + o(1)}{\mathbb{E}[Z_h] + 1 + o(1)} \leq (1 + \varepsilon'_\beta) e^{2\beta - \tilde{\alpha}_\ell},$$

and, as $\{Z_h = 0\}$ is implied by $\{M_n < h\}$, this concludes the proof. \square

6.2. *The expectation and median of the maximum.* The following is a straightforward consequence of the results we have established in this section.

COROLLARY 6.6. *There exist $\beta_0 > 0$ and a sequence $\varepsilon_\beta > 0$ vanishing as $\beta \rightarrow \infty$ such that, for all $\beta > \beta_0$, if m_n^* is defined as in (1.2) and m_n is a median of M_n , then $m_n \in [m_n^* - 1, m_n^*]$ and $m_n^* - 1 - \varepsilon_\beta \leq \mathbb{E}[M_n] \leq m_n^* + \varepsilon_\beta$.*

PROOF. For the bounds on the median m_n , by Proposition 6.1 we have that

$$\mu_n^\mp(M_n \geq m_n^* + 1) \leq (1 + \varepsilon)e^{2\beta - \tilde{\alpha}_1} \leq Ce^{-2\beta} \leq \varepsilon'_\beta,$$

using that $\exp(-\tilde{\alpha}_1) = \mu_{\mathbb{Z}^3}^\mp(E_1^o) \leq Ce^{-4\beta}$. This implies that, once β is large enough such that $\varepsilon'_\beta < \frac{1}{2}$, the median satisfies $m_n \leq m_n^*$. Further, by that same proposition, $\mu_n^\mp(M_n \leq m_n^* - 2) \leq (1 + \varepsilon)e^{2\beta - \tilde{\alpha}_1} \leq \varepsilon'_\beta$, whence $m_n \geq m_n^* - 1$.

For the bound on the expectation, note that by Proposition 6.1, as argued above for the median, we have

$$\sum_{\ell=1}^{\sqrt{\log n}} \mu_n^\mp(M_n - m_n^* \geq \ell) < \varepsilon_\beta \quad \text{and} \quad \sum_{\ell=1}^{\sqrt{\log n}} \mu_n^\mp(m_n^* - 1 - M_n \geq \ell) < \varepsilon_\beta.$$

Therefore (denoting by a_+ the positive part of a),

$$\begin{aligned} (6.7) \quad \mathbb{E}[(M_n - m_n^*)_+] &\leq \sum_{\ell \geq 1} \mu_n^\mp(M_n - m_n^* \geq \ell) \\ &\leq \sum_{\ell=1}^{\sqrt{\log n}} \mu_n^\mp(M_n - m_n^* \geq \ell) + \sum_{\ell \geq \sqrt{\log n}} \mu_n^\mp(M_n - m_n^* \geq \ell) \\ (6.8) \quad &\leq \varepsilon_\beta + Ce^{-(2\beta - C)\sqrt{\log n}} \leq \varepsilon'_\beta, \end{aligned}$$

where we used the uniform bound (6.6) on the right tail to obtain the second line. At the same time,

$$\begin{aligned} (6.9) \quad \mathbb{E}[(m_n^* - 1 - M_n)_+] &\leq m_n^* \mu_n^\mp(M_n \leq m_n^* - \sqrt{\log n}) + \sum_{\ell=1}^{\sqrt{\log n}} \mu_n^\mp(m_n^* - 1 - M_n \geq \ell) \\ &\leq \varepsilon_\beta + O(\log n)e^{-c\sqrt{\log n}} \leq \varepsilon'_\beta. \end{aligned}$$

Letting

$$p_n = \mu_n^\mp(M_n < m_n^*),$$

we may express

$$\begin{aligned} \mathbb{E}[M_n \mathbf{1}\{M_n \geq m_n^*\}] &= m_n^*(1 - p_n) + \mathbb{E}[(M_n - m_n^*)_+], \\ \mathbb{E}[M_n \mathbf{1}\{M_n \leq m_n^* - 1\}] &= (m_n^* - 1)p_n - \mathbb{E}[(m_n^* - 1 - M_n)_+] \end{aligned}$$

and deduce from (6.8) and (6.9) that

$$m_n^*(1 - p_n) \leq \mathbb{E}[M_n \mathbf{1}\{M_n \geq m_n^*\}] \leq m_n^*(1 - p_n) + \varepsilon_\beta,$$

whereas

$$(m_n^* - 1)p_n - \varepsilon_\beta \leq \mathbb{E}[M_n \mathbf{1}\{M_n \leq m_n^* - 1\}] \leq (m_n^* - 1)p_n.$$

Combining these, $|\mathbb{E}[M_n] - \xi_n| \leq \varepsilon_\beta$, where $\xi_n = m_n^* - p_n \in [m_n^* - 1, m_n^*]$, as required. \square

7. Gumbel tail estimates for the maximum.

7.1. *Coupling of different scales.* The following proposition compares M_n , the maximum height of the interface \mathcal{I} under μ_n^\mp , to the maxima of i.i.d. copies on boxes of a smaller scale. This will later be used to deduce Gumbel tail bounds for the centered maximum.

PROPOSITION 7.1. *There exists $\beta_0 > 0$ such that the following holds for all $\beta > \beta_0$. Fix $\gamma > 0$, let L_n be a sequence with $n(\log n)^{-\gamma} < L_n < n$ and set $\kappa_n = (\lfloor n/L_n \rfloor)^2$ and $R_n = n \bmod L_n$. Then,*

$$\|\mu_n^\mp(M_n \in \cdot) - \mathbb{P}(\max\{Y_1, \dots, Y_{\kappa_n}\} \in \cdot)\|_{TV} \leq O((R_n/n)^{1/3} + (\log n)^{-10}),$$

where Y_1, \dots, Y_{κ_n} are i.i.d. with law $\mathbb{P}(Y_i \in \cdot) = \mu_{L_n}^\mp(M_{L_n} \in \cdot)$. In particular,

$$\sup_x |\mu_n^\mp(M_n \leq x) - \mu_{L_n}^\mp(M_{L_n} \leq x)^{\kappa_n}| \leq O((R_n/n)^{1/3} + (\log n)^{-10}).$$

We first need the following simple claim, ruling out the improbable scenario where the maximum is attained above a fixed microscopic subset of faces of $\mathcal{L}_{0,n}$.

CLAIM 7.2. *For every $\delta > 0$, there exists $\beta_0 > 0$ such that the following holds for all $\beta > \beta_0$. Let $S_0 \subset \mathcal{L}_{0,n}$ be a deterministic set of faces of size $|S_0| \leq n^{2-\delta}$. Let M_n be the maximum height of \mathcal{I} under μ_n^\mp , and let $M_n^- = \max\{\text{ht}(\mathcal{P}_x) : x \in \mathcal{L}_{0,n} \setminus S_0\}$. Then, for every large enough n , we have $\mu_n^\mp(M_n \neq M_n^-) \leq (\log n)^{-3\beta}$.*

PROOF. Let $h = m_n^* - \log \log n$. We may bound the sought probability by

$$\begin{aligned} \mu_n^\mp(M_n \neq M_n^-) &\leq \mu_n^\mp\left(\max_{x \in S_0} \text{ht}(\mathcal{P}_x) \geq h\right) + \mu_n^\mp(M_n < h) \\ &\leq |S_0| \exp(-(4\beta - C)h) + O(\exp(-\tilde{\alpha}_{\log \log n})), \end{aligned}$$

with the last inequality relying on Proposition 3.10 to bound the first probability and Proposition 6.1 to bound the second one. The last term is, at most, $O((\log n)^{-4\beta-C})$ by the definition (5.4) of $\tilde{\alpha}_h$ and the inequality succeeding it. That same inequality implies that

$$(4\beta - C)h \geq \tilde{\alpha}_h - \varepsilon_\beta h \geq \tilde{\alpha}_{m_n^*} - (\varepsilon_\beta - o(1))h \geq (2 - \varepsilon'_\beta) \log n,$$

where the first inequality used that $\tilde{\alpha}_{m_n^*} - \tilde{\alpha}_h = O(\log \log n) = o(h)$ by Corollary 5.2. In light of this,

$$|S_0| e^{-(4\beta-C)h} \leq n^{2-\delta} n^{-2+\varepsilon'_\beta} \leq n^{-\delta+\varepsilon'_\beta+o(1)} < n^{-\delta/2+o(1)}$$

for β large enough so that $\varepsilon'_\beta < \delta/2$. Hence, combined with the above inequality on $\mu_n^\mp(M_n < h)$, we get that $\mu_n^\mp(M_n \neq M_n^-)$ is $O((\log n)^{-4\beta+C}) < (\log n)^{-3\beta}$ for large enough n provided $\beta > C$. \square

PROOF OF PROPOSITION 7.1. First, consider the case $R_n = 0$. Partition $\mathcal{L}_{0,n}$ into disjoint boxes $\mathcal{B}_1, \dots, \mathcal{B}_{\kappa_n}$, each of side length $2L_n$, and, further, let

$$\mathcal{B}_i^- = \{x \in \mathcal{B}_i : d(x, \partial \mathcal{B}_i) \geq \log^2 n\} \quad (i = 1, \dots, \kappa_n).$$

We will show that M_n is equal to $\max_i \bar{M}_{L_n}^{(i)}$ with high probability, where

$$\bar{M}_{L_n}^{(i)} = \max\{\text{ht}(\mathcal{P}_x) : x \in \mathcal{B}_i^-\},$$

which, in turn, can be coupled (with negligible error) to κ_n i.i.d. copies of $\bar{M}_{L_n}^{(1)}$ under $\mu_{\mathcal{B}_1}^\mp$. Each of those i.i.d. copies will then be coupled to M_{L_n} .

First, since $\kappa_n = O((\log n)^{2\gamma})$, we have

$$\left| \mathcal{L}_{0,n} \setminus \left(\bigcup_{i=1}^{\kappa_n} \mathcal{B}_i^- \right) \right| \leq \kappa_n \cdot 4L_n \log^2 n = O(n(\log n)^{2\gamma+2})$$

and thus infer from Claim 7.2 that (with some room)

$$(7.1) \quad \mu_n^\mp \left(M_n \neq \max_{1 \leq i \leq \kappa_n} \bar{M}_{L_n}^{(i)} \right) < (\log n)^{-10}.$$

Second, as the boxes \mathcal{B}_i^- have pairwise distances at least $\log^2 n$, iterating Corollary 5.7 per box shows that

$$\left\| \mu_n^\mp \left((\mathcal{P}_x)_{x \in \cup_i \mathcal{B}_i^-} \in \cdot \right) - \prod_i \mu_n^\mp \left((\mathcal{P}_x)_{x \in \mathcal{B}_i^-} \in \cdot \right) \right\|_{\text{TV}} = O(\kappa_n e^{-c \log^2 n}) < n^{-10}.$$

Moreover, Corollary 5.6 gives

$$\sum_{i=1}^{\kappa_n} \left\| \mu_n^\mp \left((\mathcal{P}_x)_{x \in \mathcal{B}_i^-} \in \cdot \right) - \mu_{\mathcal{B}_i^-}^\mp \left((\mathcal{P}_x)_{x \in \mathcal{B}_i^-} \in \cdot \right) \right\|_{\text{TV}} = O(\kappa_n e^{-c \log^2 n}) < n^{-10},$$

thus, with probability $1 - O(n^{-10})$ we may couple $(\bar{M}_{L_n}^{(1)}, \dots, \bar{M}_{L_n}^{(\kappa_n)})$ under μ_n^\mp to $(Y_1^-, \dots, Y_{\kappa_n}^-)$ where Y_i^- are i.i.d. distributed according to $\mu_{\mathcal{B}_i^-}^\mp(\bar{M}_{L_n}^{(i)} \in \cdot)$. Letting Y_i be i.i.d. copies of $\mu_{\mathcal{B}_i^-}^\mp(\max_{x \in \mathcal{B}_i^-} \text{ht}(\mathcal{P}_x) \in \cdot)$, we again apply Claim 7.2, this time to $\mu_{\mathcal{B}_i^-}^\mp$ with $S_0 = \mathcal{B}_i \setminus \mathcal{B}_i^-$, so that $|S_0| = O(|L_n| \log^2 n) = L_n^{1+o(1)}$ and

$$\sum_{i=1}^{\kappa_n} \mu_{\mathcal{B}_i^-}^\mp (Y_i \neq Y_i^-) < \kappa_n (\log n)^{-3\beta} < (\log n)^{-10}$$

for β large. Altogether, the total variation distance between the law of M_n and $\max\{Y_1, \dots, Y_{\kappa_n}\}$ —equal in distribution to κ_n i.i.d. copies of M_{L_n} under $\mu_{L_n}^\mp$ —is $O((\log n)^{-10})$, as required.

It remains to handle the case $R_n \neq 0$. Here, we will partition $\mathcal{L}_{0,n}$ into boxes

$$\mathcal{B}_1, \dots, \mathcal{B}_{\kappa_n}, \hat{\mathcal{B}}_{\kappa_n+1}, \dots, \hat{\mathcal{B}}_{\hat{\kappa}_n},$$

where the boxes $\mathcal{B}_1, \dots, \mathcal{B}_{\kappa_n}$ have side length $2L_n$, as before, and the remaining boxes $\hat{\mathcal{B}}_j$ ($j > \kappa_n$) have the shorter side length $2R_n$ (so that $\hat{\kappa}_n = (\lceil n/L_n \rceil)^2$). We use the same definition of \mathcal{B}_i^- also for $\hat{\mathcal{B}}_j^-$, that is,

$$\hat{\mathcal{B}}_j^- = \{x \in \hat{\mathcal{B}}_j : d(x, \partial \hat{\mathcal{B}}_j) \geq \log^2 n\} \quad (j = \kappa_n + 1, \dots, \hat{\kappa}_n),$$

noting that it may be the case that $\hat{\mathcal{B}}_j^- = \emptyset$ (whenever $R_n \leq \log^2 n$). However, we would still want to couple M_n to $\max_{i \leq \kappa_n} \bar{M}_{L_n}^{(i)}$, as we did in (7.1), ignoring the exceptional boxes $\hat{\mathcal{B}}_j$. To achieve this, apply Claim 7.2 with $S_0 = \mathcal{L}_{0,n} \setminus (\cup_{i=1}^{\kappa_n} \mathcal{B}_i^- \cup \cup_{j=\kappa_n+1}^{\hat{\kappa}_n} \hat{\mathcal{B}}_j^-)$, which, as before, has $|S_0| = O(\hat{\kappa}_n L_n \log^2 n) = O(n(\log n)^{2\gamma+2})$, and hence

$$(7.2) \quad \mu_n^\mp \left(M_n \neq \max_{x \in \mathcal{L}_{0,n} \setminus S_0} \text{ht}(\mathcal{P}_x) \right) < (\log n)^{-10}.$$

We may treat $S_1 = \cup_{j > \kappa_n} \hat{\mathcal{B}}_j^-$ as follows: recall from Corollary 5.2 that we have $\alpha_{m_n^* - \ell} \geq \alpha_{m_n^*} - (4\beta + e^{-4\beta})\ell$ and $\alpha_\ell \geq (4\beta - C)\ell$, and thus for $\ell > 0$ to be specified below,

$$\begin{aligned} \mu_n^\mp \left(M_n = \max_{x \in S_1} \text{ht}(\mathcal{P}_x) \right) &\leq |S_1| \max_{x \in S_1} \mu_n^\mp (\text{ht}(\mathcal{P}_x) \geq m_n^* - \ell) + \mu_n^\mp (M_n < m_n^* - \ell) \\ &\leq O(|S_1| e^{-\alpha_{m_n^* - \ell}} + e^{-\alpha_\ell}) \leq O\left(\frac{|S_1|}{n^2} e^{(4\beta + e^{-4\beta})\ell} + e^{-(4\beta - C)\ell}\right), \end{aligned}$$

where we used that $\mu_n^\mp(\text{ht}(\mathcal{P}_x) \geq h) = (1 - o(1))e^{-\alpha_n}$ for $x \in S_1$ and $h = m_n^* - \ell$ (as $h \ll \log^2 n \leq d(x, \partial\hat{\mathcal{B}}_j)$ for such x), the definition of $\alpha_{m_n^*}$, and Proposition 6.1. Choose

$$\ell = \ell_1 \wedge \ell_2 \quad \text{where } \ell_1 = \frac{3}{\beta} \log \log n \text{ and } \ell_2 = \frac{1}{8\beta} \log\left(\frac{n}{R_n}\right).$$

We see that $\ell = \ell_1$ implies $R_n/n \leq (\log n)^{-24}$, in which case, using $|S_1| = O(nR_n)$,

$$e^{-(4\beta-C)\ell} + \frac{|S_1|}{n^2} e^{(4\beta+e^{-4\beta})\ell} \leq (\log n)^{-12+\varepsilon\beta} + O(R_n/n)(\log n)^{12+\varepsilon\beta} = O((\log n)^{-12+\varepsilon\beta}).$$

On the other hand, when $\ell = \ell_2$ we have

$$e^{-(4\beta-C)\ell} + \frac{|S_1|}{n^2} e^{(4\beta+e^{-4\beta})\ell} \leq (R_n/n)^{\frac{1}{2}-\varepsilon\beta} + O((R_n/n)^{1-(\frac{1}{2}+\varepsilon\beta)}) = O((R_n/n)^{\frac{1}{2}-\varepsilon\beta}).$$

Altogether, for large enough β we find that

$$\mu_n^\mp\left(M_n = \max_{x \in S_1} \text{ht}(\mathcal{P}_x)\right) \leq O((R_n/n)^{1/3} + (\log n)^{-10}).$$

Combining this with (7.2), while noticing that $\mathcal{L}_{0,n} \setminus (S_0 \cup S_1)$ is nothing but $\bigcup_{i \leq \kappa_n} \mathcal{B}_i^-$, we obtain that

$$\mu_n^\mp\left(M_n \neq \max_{i \leq \kappa_n} \bar{M}_{L_n}^{(i)}\right) \leq O((R_n/n)^{1/3} + (\log n)^{-10}),$$

at which point the original analysis of the law of $\max_{i \leq \kappa_n} \bar{M}_{L_n}^{(i)}$, showing that it is coupled to the maximum of κ_n i.i.d. copies of M_{L_n} under $\mu_{L_n}^\mp$, completes the proof. \square

7.2. From multiscale coupling to Gumbel tails. We will first prove the sought bounds in the special case when the side length $2n$ is a power of 2. This will be extended to the general case at the end of Section 7.2.

7.2.1. Left tail. The following lemma establishes the doubly exponential left tail of the centered maximum.

LEMMA 7.3. *There exists $\beta_0 > 0$ such that, for every $\beta > \beta_0$, the following holds. For every fixed $\ell \geq 1$ and every large enough n that is a power of 2,*

$$\exp(-e^{(4\beta+e^{-4\beta})\ell+2}) \leq \mu_n^\mp(M_n \leq m_n^* - \ell) \leq \exp(-e^{\alpha_\ell-12\beta}).$$

PROOF. The proof of both inequalities will follow from coupling $\mu_n^\mp(M_n \in \cdot)$ to the maxima of smaller scales. We begin with the lower bound. Consider $n_1 = n2^{-j}$ and $n_2 = n2^{-(j+1)}$. Since $m_{n_i}^*$ is the minimal h such that α_h exceeds the threshold $2 \log(2n_i) - 2\beta$, the difference of these thresholds between $m_{n_1}^*$ and $m_{n_2}^*$ is precisely $2 \log 2$, whereas $\alpha_{h+1} - \alpha_h \geq \alpha_1 - \varepsilon\beta \geq 4\beta - C - \varepsilon\beta$ holds for every $h \geq 1$ by Corollary 5.2. Hence,

$$m_{n_1}^* - 1 \leq m_{n_2}^* \leq m_{n_1}^*,$$

and we may consider $L_{n,k} = n2^{-k}$ for the minimal $k \geq 0$ that would satisfy

$$m_{L_{n,k}}^* = m_n^* - \ell.$$

(The fact that $\ell \geq 1$ implies that $k > 0$.) We claim that this k satisfies

$$k \leq k_1 := \left\lceil \frac{(4\beta + e^{-4\beta})\ell}{2 \log 2} \right\rceil.$$

To see this, recall from Corollary 5.2 and the inequality below (5.4) that

$$2 \log(2n) - 2\beta < \alpha_{m_n^*} \leq \alpha_{m_n^* - \ell} + (4\beta + e^{-4\beta})\ell;$$

thus (using that $2 \log(2n) = 2 \log(2L_{n,k}) + k \log 4$ for every $k \geq 0$)

$$\alpha_{m_n^* - \ell} > 2 \log(2L_{n,k_1}) - 2\beta + k_1 \log 4 - (4\beta + e^{-4\beta})\ell \geq 2 \log(2L_{n,k_1}) - 2\beta,$$

so $m_{L_{n,k_1}}^* \leq m_n^* - \ell$ and $k \leq k_1$ by definition. Since $k_1 = O(1)$, we have $n/L_{n,k} = O(1)$, and Proposition 7.1 implies that

$$\begin{aligned} \mu_n^\mp(M_n \leq m_n^* - \ell) &= \mu_{L_{n,k}}^\mp(M_{L_{n,k}} \leq m_{L_{n,k}}^*)^{4^k} + o(1) \geq (1 - \varepsilon_\beta)^{4^{k_1}} + o(1) \\ &\geq \exp(-4^{k_1}) \geq \exp(-4e^{(4\beta + e^{-4\beta})\ell}), \end{aligned}$$

where we used that $1 - \varepsilon_\beta > e^{-1}$ for every large enough β in the transition between the lines, absorbing the $o(1)$ -term for n large enough in the process. This implies the desired lower bound.

For the upper bound, let $L_{n,k} = n2^{-k}$ for the minimal $k \geq 0$ that would satisfy

$$m_{L_{n,k}}^* = m_n^* - \ell + 2.$$

Further, assume for now that $\ell \geq 3$ (hence, $k > 0$); our resulting upper bound will hold trivially for $\ell = 1, 2$. We immediately note that $k \leq k_1$, since we saw above that $m_{L_{n,k_1}}^* \leq m_n^* - \ell$. We will need a lower bound on k to yield the required tail estimate. Once again, appealing to Corollary 5.2, we have

$$\alpha_{m_n^* - \ell + 2} \leq \alpha_{m_n^*} - \alpha_{\ell - 2} + \varepsilon_\beta \leq 2 \log(2n) + 2\beta - \alpha_{\ell - 2} + \varepsilon'_\beta$$

for some other sequence $\varepsilon'_\beta > 0$ vanishing as $\beta \rightarrow \infty$, where we used (6.1) and the relation between $\alpha_h, \tilde{\alpha}_h$ in (5.4). It now follows that

$$k \geq k_2 := \left\lfloor \frac{\alpha_{\ell - 2} - 4\beta - 1}{2 \log 2} \right\rfloor,$$

since, for β large enough so that $\varepsilon'_\beta < 1$, we have

$$\alpha_{m_n^* - \ell + 2} \leq 2 \log(2L_{n,k_2}) - 2\beta + (4\beta + \varepsilon'_\beta - \alpha_{\ell - 2} + k_2 \log 4) < 2 \log(2L_{n,k_2}) - 2\beta.$$

Applying Proposition 7.1 (recalling that $k \leq k_1$ and so $n/L_{n,k} = O(1)$ as before), we deduce that

$$\mu_n^\mp(M_n \leq m_n^* - \ell) = \mu_{L_{n,k}}^\mp(M_{L_{n,k}} \leq m_{L_{n,k}}^* - 2)^{4^k} + o(1) \leq ((1 + \varepsilon_\beta)e^{2\beta - \tilde{\alpha}_1})^{4^{k_2}},$$

using Proposition 6.1 for the last inequality. Using that $\tilde{\alpha}_1 \geq 4\beta - C$ and absorbing $\log(1 + \varepsilon_\beta)$ into this constant, we see that

$$\begin{aligned} (7.3) \quad \mu_n^\mp(M_n \leq m_n^* - \ell) &\leq \exp(-(2\beta - C)4^{k_2}) \\ &\leq \exp(-(2\beta - C)e^{\alpha_{\ell - 2} - 4\beta - 1}) \leq \exp(-e^{\alpha_{\ell - 2} - 12\beta}), \end{aligned}$$

where in the last inequality we used that $\alpha_{\ell - 2} \geq \alpha_\ell - 8\beta - \varepsilon_\beta$ (again by Corollary 5.2) and, thereafter, added the term $1 + \varepsilon_\beta$ to the exponent in exchange for the factor $2\beta - C$, which is valid for large β . (Note that, as promised above, the resulting bound holds also for $\ell = 1, 2$, becoming trivial since $\alpha_1 \leq \alpha_2 \leq 8\beta + \varepsilon_\beta$.) \square

7.2.2. *Right tail.* The exponential upper bound on the right tail of the centered maximum was established in Proposition 6.1, implying, via the relation (5.4) between $\alpha_n, \tilde{\alpha}_n$, that (with room to spare), for every $\ell \geq 1$,

$$\mu_n^\mp(M_n \leq m_n^* + \ell) \leq \exp(-\alpha_\ell + 3\beta).$$

It remains to provide a corresponding lower bound, as given by the following lemma.

LEMMA 7.4. *There exists $\beta_0 > 0$ such that, for every $\beta > \beta_0$, the following holds. For every fixed $\ell \geq 1$ and every large enough n that is a power of 2,*

$$\mu_n^\mp(M_n \leq m_n^* + \ell) \geq \exp(-(4\beta + e^{-4\beta})\ell - 4(\beta + 1)).$$

PROOF. The proof will follow from coupling i.i.d. copies of $\mu_n^\mp(M_n \in \cdot)$ to the maximum of a larger scale. As in the proof of Lemma 7.3—now viewing increasing rather than decreasing side lengths—we have that if $n_1 = n2^j$ and $n_2 = n2^{j+1}$, then $m_{n_1}^* \leq m_{n_2}^* \leq m_{n_1}^* + 1$, and, therefore, we may consider $L_{n,k} = n2^k$ for the minimal $k \geq 0$ (in fact, $k > 0$ necessarily) that satisfies

$$m_{L_{n,k}}^* = m_n^* + \ell + 1.$$

We claim that

$$k \leq k_1 := \left\lceil \frac{4\beta + 1 + (4\beta + e^{-4\beta})\ell}{2 \log 2} \right\rceil.$$

Too see this, recall from Corollary 5.2 and (6.1) (combined with the usual relation between $\tilde{\alpha}_h, \alpha_h$) that

$$\alpha_{m_n^* + \ell} \leq \alpha_{m_n^*} + (4\beta + e^{-4\beta})\ell \leq 2 \log(2n) + 2\beta + \varepsilon_\beta + (4\beta + e^{-4\beta})\ell.$$

Writing $2 \log(2n) = 2 \log(2L_{n,k}) - k \log 4$, we get that, for any $k \geq 0$,

$$\alpha_{m_n^* + \ell} \leq 2 \log(2L_{n,k}) - 2\beta + (4\beta + \varepsilon_\beta + (4\beta + e^{-4\beta})\ell - k \log 4),$$

and substituting $k = k_1$, as chosen above, now yields (for β large enough so that $\varepsilon_\beta < 1$)

$$\alpha_{m_n^* + \ell} < 2 \log(2L_{n,k_1}) - 2\beta,$$

and, therefore, $m_{L_{n,k_1}}^* > m_n^* + \ell$; that is, $m_{L_{n,k_1}}^* \geq m_n^* + \ell + 1$, implying that $k \leq k_1$ as claimed.

Since $k_1 = O(1)$, so $n/L_{n,k} = O(1)$, we may invoke Proposition 7.1 and find that

$$\mu_{L_{n,k}}^\mp(M_{L_{n,k}} < m_{L_{n,k}}^* - 1) = \mu_n^\mp(M_n < m_n^* + \ell)^{4^k},$$

and so

$$\mu_{L_{n,k}}^\mp(M_{L_{n,k}} \geq m_{L_{n,k}}^* - 1) \leq 4^k \mu_n^\mp(M_n \geq m_n^* + \ell).$$

Using Proposition 6.1 to bound the left-hand side from below by $1 - \varepsilon_\beta$, we obtain that

$$\begin{aligned} \mu_n^\mp(M_n \geq m_n^* + \ell) &\geq (1 - \varepsilon_\beta)4^{-k} \\ &\geq (1 - \varepsilon_\beta)4^{-k_1} \geq \left(\frac{1}{4} - \varepsilon'_\beta\right) \exp(-4\beta - 1 - (4\beta + e^{-4\beta})\ell) \\ &\geq \exp(-4\beta - 3 - (4\beta + e^{-4\beta})\ell) \end{aligned}$$

for large enough β , as required. \square

PROOF OF THEOREM 2. For n that is a power of 2, the bounds in Theorem 2 were all established: the lower bounds were obtained in Lemmas 7.3 and 7.4 for a choice of $\bar{\alpha} = 4\beta + e^{-4\beta}$; the upper bounds were obtained by Lemma 7.3 and by (7.3) (which followed from Proposition 6.1). It remains to extend the estimates in Lemmas 7.3 and 7.4 to general n , which will follow from the decorrelation inequalities of Section 5.

Let N be a power of 2 such that $N/4 < n < N/2$. By Corollary 5.6 we have that

$$\|\mu_n^\mp((\mathcal{P}_x)_{x \in \mathcal{L}_{0,n}^-} \in \cdot) - \mu_N^\mp((\mathcal{P}_x)_{x \in \mathcal{L}_{0,n}^-} \in \cdot)\|_{\text{TV}} = O(e^{-c \log^2 n}) = o(1).$$

Recall that $m_n^* \leq m_N^* \leq m_n^* + 1$ (since the scales changed by, at most, a factor of 2, whereas $\alpha_{h+1} - \alpha_h \geq 4\beta - C$, as explained in the proofs of Lemmas 7.3 and 7.4). Furthermore, if $M_n^- = \max\{\text{ht}(\mathcal{P}_x) : x \in \mathcal{L}_{0,n}^-\}$, then Claim 7.2 shows that $\mu_n^\mp(M_n \neq M_n^-) = o(1)$, and so

$$\mu_n^\mp(M_n \leq m_n^* - \ell) = \mu_n^\mp(M_n^- \leq m_n^* - \ell) + o(1) \geq \mu_N^\mp(M_N \leq m_N^* - \ell - 1) + o(1).$$

A lower bound on the probability in the right-hand side is given by Lemma 7.3, whereby (recalling $\bar{\alpha} = 4\beta + e^{-4\beta}$)

$$\mu_N^\mp(M_N \leq m_N^* - \ell - 1) \geq \exp(-e^{\bar{\alpha}(\ell+1)+3}) \geq \exp(-e^{\bar{\alpha}\ell+4(\beta+1)}),$$

a lower bound that extends to $\mu_n^\mp(M_n \leq m_n^* - \ell) + o(1)$ via the preceding inequality.

The remaining two inequalities (the upper bound on the right tail in (7.3) was already established for all n) will follow from a comparison of μ_n^\mp to μ_N^\mp where N is a power of 2 such that $2N < n < 4N$. The same coupling mentioned above shows that

$$\|\mu_N^\mp((\mathcal{P}_x)_{x \in \mathcal{L}_{0,N}^-} \in \cdot) - \mu_n^\mp((\mathcal{P}_x)_{x \in \mathcal{L}_{0,N}^-} \in \cdot)\|_{\text{TV}} = O(e^{-c \log^2 n}) = o(1).$$

As before, $\mu_N^\mp(M_N \neq M_N^-) = o(1)$, and now we have $m_N^* \leq m_n^* \leq m_N^* + 1$, whence

$$\mu_n^\mp(M_n \leq m_n^* - \ell) \leq \mu_N^\mp(M_N^- \leq m_N^* - \ell + 1) + o(1) = \mu_N^\mp(M_N \leq m_N^* - \ell + 1) + o(1).$$

By the upper bound in Lemma 7.3,

$$\mu_N^\mp(M_N \leq m_N^* - \ell + 1) \leq \exp(-e^{\alpha\ell-1-12\beta}) \leq \exp(-e^{\alpha\ell-(16\beta+1)}).$$

Similarly, we have

$$\mu_n^\mp(M_n \geq m_n^* + \ell) \geq \mu_N^\mp(M_N^- \geq m_N^* + \ell + 1) - o(1) = \mu_N^\mp(M_N \geq m_N^* + \ell + 1) - o(1),$$

whereas by Lemma 7.4,

$$\mu_N^\mp(M_N \geq m_N^* + \ell + 1) \geq \exp(-\bar{\alpha}(\ell + 1) - 4(\beta + 1)) \geq \exp(-\bar{\alpha}\ell - (8\beta + 5)),$$

thus concluding the proof. \square

7.3. *Nonconvergence of the centered maximum.* The following simple corollary of Proposition 7.1 will be used to derive Proposition 3.

COROLLARY 7.5. *Let M_n be the maximum height of \mathcal{I} under μ_n^\mp . Fix $\omega \in \mathbb{R}_+$. Then,*

$$\limsup_{n \rightarrow \infty} \sup_x |\mu_n^\mp(M_n \leq x) - \mu_{\lfloor n/\omega \rfloor}^\mp(M_{\lfloor n/\omega \rfloor} \leq x)|^{\omega^2} = 0.$$

PROOF. Fix any $\varepsilon > 0$, set $n_\omega = \lfloor n/\omega \rfloor$ for brevity and let

$$\Gamma = \Gamma_{\varepsilon,n} \cup \Gamma_{\varepsilon,n_\omega} \quad \text{where } \Gamma_{\varepsilon,n} = \{x \in \mathbb{R} : \mu_n^\mp(M_n \leq x) \geq \varepsilon\}.$$

By writing

$$\sup_{x_n} |\mu_n^\mp(M_n \leq x_n) - \mu_{n_\omega}^\mp(M_{n_\omega} \leq x_n)^{\omega^2}| \leq \varepsilon + \sup_{x_n \in \Gamma} |\mu_n^\mp(M_n \leq x_n) - \mu_{n_\omega}^\mp(M_{n_\omega} \leq x_n)^{\omega^2}|,$$

it will suffice to prove that the second term in the right-hand vanishes as $n \rightarrow \infty$ for any such choice of ε .

Applying Proposition 7.1 to $\mu_n^\mp(M_n \in \cdot)$ with $L_n = \lfloor n/\log n \rfloor$, whereby $\kappa_n = (1 + o(1)) \log^2 n$, we obtain

$$\sup_{x_n \in \Gamma} |\mu_n^\mp(M_n \leq x_n) - p_{n,x_n}^{\kappa_n}| = o(1) \quad \text{where } p_{n,x_n} = \mu_{L_n}^\mp(M_{L_n} \leq x_n).$$

Another application of Proposition 7.1, this time to $\mu_{n_\omega}^\mp(M_{n_\omega} \in \cdot)$ yet with the same $L_n = \lfloor n/\log n \rfloor$, gives

$$\sup_{x_n \in \Gamma} |\mu_{n_\omega}^\mp(M_{n_\omega} \leq x_n) - p_{n,x_n}^{\kappa'_n}| = o(1),$$

where p_{n,x_n} is as before, and $\kappa'_n = (\omega^{-2} + o(1))\kappa_n$ (recall that $\omega \in \mathbb{R}_+$ is fixed).

For $x_n \in \Gamma_{\varepsilon,n}$, we have that $p_{n,x_n}^{\kappa_n} \geq \varepsilon - o(1)$, whence $p_{n,x_n}^{(1+o(1))f_n} = p_{n,x_n}^{f_n} + o(1)$ if f_n has order $\log^2 n$, and the same conclusion holds if $x_n \in \Gamma_{\varepsilon,n_\omega}$ (as κ'_n is also of order $\log^2 n$). In particular,

$$\sup_{x_n \in \Gamma} |\mu_n^\mp(M_n \leq x_n) - p_{n,x_n}^{\log^2 n}| = o(1) \quad \text{and} \quad \sup_{x_n \in \Gamma} |\mu_{n_\omega}^\mp(M_{n_\omega} \leq x_n) - p_{n,x_n}^{\omega^{-2} \log^2 n}| = o(1),$$

and, combining these inequalities, yields that

$$\sup_{x_n \in \Gamma} |\mu_n^\mp(M_n \leq x_n) - \mu_{n_\omega}^\mp(M_{n_\omega} \leq x_n)^{\omega^2}| = o(1),$$

as required. \square

PROOF OF PROPOSITION 3. Suppose that m_n is a sequence such that $\{M_n - m_n\}$ weakly converges to a nondegenerate random variable M_∞ with a distribution function G . Since $m_n \bmod 1$ must converge (as M_n is integer valued), we may assume w.l.o.g. that $m_n \in \mathbb{Z}$ (whence G is also integer valued). Observe that the bounds in Theorem 2 (in fact, already those of Proposition 6.1) imply that we must have

$$\limsup_{n \rightarrow \infty} |m_n - m_n^*| < \infty.$$

Fix $k \geq 2$, and consider Corollary 7.5 with $\omega = \sqrt{k}$. By assumption, $\lim_{n \rightarrow \infty} \mu_n^\mp(M_n \leq m_n + x) = G(x)$ for all x , and the above corollary then implies that

$$\lim_{n \rightarrow \infty} \mu_{\lfloor n/\sqrt{k} \rfloor}^\mp(M_{\lfloor n/\sqrt{k} \rfloor} \leq m_n + x) = G^{1/k}(x).$$

However, $|m_n^* - m_{\lfloor n/\sqrt{k} \rfloor}^*| \leq C(\beta) \log k$ (see, e.g., the estimate (6.1)), and the above bound on $|m_n - m_n^*|$ thus implies that

$$a_{k,n} := m_n - m_{\lfloor n/\sqrt{k} \rfloor}$$

satisfies $\limsup_{n \rightarrow \infty} |a_{k,n}| < \infty$. Let $(a_{k,n_j})_{j \geq 1}$ be a converging subsequence of $(a_{k,n})_{n \geq 1}$, and denote its limit by a_k , whereby, recalling that $m_n \in \mathbb{Z}$, we must have $a_{k,n_j} = a_k$ for all sufficiently large j . Thus, for every large enough j ,

$$\mu_{\lfloor \frac{n_j}{\sqrt{k}} \rfloor}^\mp(M_{\lfloor \frac{n_j}{\sqrt{k}} \rfloor} \leq m_{n_j} + x) = \mu_{\lfloor \frac{n_j}{\sqrt{k}} \rfloor}^\mp(M_{\lfloor \frac{n_j}{\sqrt{k}} \rfloor} \leq m_{\lfloor \frac{n_j}{\sqrt{k}} \rfloor} + x + a_k) \xrightarrow{j \rightarrow \infty} G(x + a_k),$$

where the last equality is again by our weak convergence assumption. Together, this implies that, for every $x \in \mathbb{R}$, we have $G^k(x + a_k) = G(x)$; having established this for every $k \geq 2$, we find that G is max-stable, yet G is discrete, contradicting the fact that the only (nondegenerate) max-stable distributions are continuous ones, belonging to one of the three classes of extreme value distributions (see, e.g., [21], Theorem 1.3.1). \square

Acknowledgments. We are grateful to an anonymous referee for valuable comments. E.L. was supported in part by NSF Grant DMS-1812095.

REFERENCES

- [1] AIZENMAN, M. (1980). Translation invariance and instability of phase coexistence in the two-dimensional Ising system. *Comm. Math. Phys.* **73** 83–94. MR0573615
- [2] BOLTHAUSEN, E., DEUSCHEL, J.-D. and GIACOMIN, G. (2001). Entropic repulsion and the maximum of the two-dimensional harmonic crystal. *Ann. Probab.* **29** 1670–1692. MR1880237 <https://doi.org/10.1214/aop/1015345767>
- [3] BRAMSON, M., DING, J. and ZEITOUNI, O. (2016). Convergence in law of the maximum of the two-dimensional discrete Gaussian free field. *Comm. Pure Appl. Math.* **69** 62–123. MR3433630 <https://doi.org/10.1002/cpa.21621>
- [4] BRAMSON, M. and ZEITOUNI, O. (2012). Tightness of the recentered maximum of the two-dimensional discrete Gaussian free field. *Comm. Pure Appl. Math.* **65** 1–20. MR2846636 <https://doi.org/10.1002/cpa.20390>
- [5] BRICMONT, J., EL MELLOUKI, A. and FRÖHLICH, J. (1986). Random surfaces in statistical mechanics: Roughening, rounding, wetting, *J. Stat. Phys.* **42** 743–798. MR0833220 <https://doi.org/10.1007/BF01010444>
- [6] BRICMONT, J., LEBOWITZ, J. L. and PFISTER, C. E. (1979). Nontranslation-invariant Gibbs states with coexisting phases. II. Cluster properties and surface tension. *Comm. Math. Phys.* **66** 21–36. MR0530914
- [7] CAPUTO, P., LUBETZKY, E., MARTINELLI, F., SLY, A. and TONINELLI, F. L. (2014). Dynamics of (2+1)-dimensional SOS surfaces above a wall: Slow mixing induced by entropic repulsion. *Ann. Probab.* **42** 1516–1589. MR3262485 <https://doi.org/10.1214/13-AOP836>
- [8] CAPUTO, P., LUBETZKY, E., MARTINELLI, F., SLY, A. and TONINELLI, F. L. (2016). Scaling limit and cube-root fluctuations in SOS surfaces above a wall. *J. Eur. Math. Soc. (JEMS)* **18** 931–995. MR3500829 <https://doi.org/10.4171/JEMS/606>
- [9] CERF, R. and KENYON, R. (2001). The low-temperature expansion of the Wulff crystal in the 3D Ising model. *Comm. Math. Phys.* **222** 147–179. MR1853867 <https://doi.org/10.1007/s002200100505>
- [10] DING, J. (2013). Exponential and double exponential tails for maximum of two-dimensional discrete Gaussian free field. *Probab. Theory Related Fields* **157** 285–299. MR3101848 <https://doi.org/10.1007/s00440-012-0457-9>
- [11] DOBRUSHIN, R. and HRYNIV, O. (1997). Fluctuations of the phase boundary in the 2D Ising ferromagnet. *Comm. Math. Phys.* **189** 395–445. MR1480026 <https://doi.org/10.1007/s002200050209>
- [12] DOBRUSHIN, R., KOTECKÝ, R. and SHLOSMA, S. (1992). *Wulff Construction: A Global Shape from Local Interaction. Translations of Mathematical Monographs* **104**. Amer. Math. Soc., Providence, RI. MR1181197 <https://doi.org/10.1090/mmono/104>
- [13] DOBRUSHIN, L. R. (1973). An investigation of Gibbs states for three-dimensional lattice systems. *Teor. Veroyatn. Primen.* **18** 261–279. MR0325076
- [14] DOBRUSHIN, R. L. (1972). The Gibbs state that describes the coexistence of phases for a three-dimensional Ising model. *Teor. Veroyatn. Primen.* **17** 619–639. MR0421546
- [15] DOBRUSHIN, R. L. (1972). Asymptotical behavior of Gibbsian distributions for lattice systems and their dependence on the form of the container. *Teoret. Mat. Fiz.* **12** 115–134. MR0479233
- [16] GHEISSARI, R. and LUBETZKY, E. Maximum and shape of interfaces in 3D Ising crystals. *Comm. Pure Appl. Math.* To appear.
- [17] HIGUCHI, Y. (1981). On the absence of non-translation invariant Gibbs states for the two-dimensional Ising model. In *Random Fields, Vol. I, II (Esztergom, 1979). Colloquia Mathematica Societatis János Bolyai* **27** 517–534. North-Holland, Amsterdam. MR0712693
- [18] HRYNIV, O. (1998). On local behaviour of the phase separation line in the 2D Ising model. *Probab. Theory Related Fields* **110** 91–107. MR1602044 <https://doi.org/10.1007/s004400050146>
- [19] IOFFE, D. (1994). Large deviations for the 2D Ising model: A lower bound without cluster expansions. *J. Stat. Phys.* **74** 411–432. MR1257822 <https://doi.org/10.1007/BF02186818>

- [20] IOFFE, D. (1995). Exact large deviation bounds up to T_c for the Ising model in two dimensions. *Probab. Theory Related Fields* **102** 313–330. MR1339736 <https://doi.org/10.1007/BF01192464>
- [21] LEADBETTER, M. R., LINDGREN, G. and ROOTZÉN, H. (1983). *Extremes and Related Properties of Random Sequences and Processes*. Springer Series in Statistics. Springer, New York–Berlin. MR0691492
- [22] LUBETZKY, E., MARTINELLI, F. and SLY, A. (2016). Harmonic pinnacles in the discrete Gaussian model. *Comm. Math. Phys.* **344** 673–717. MR3508158 <https://doi.org/10.1007/s00220-016-2628-5>
- [23] PFISTER, C.-E. and VELENIK, Y. (1997). Large deviations and continuum limit in the 2D Ising model. *Probab. Theory Related Fields* **109** 435–506. MR1483597 <https://doi.org/10.1007/s004400050139>
- [24] PFISTER, C.-E. and VELENIK, Y. (1999). Interface, surface tension and reentrant pinning transition in the 2D Ising model. *Comm. Math. Phys.* **204** 269–312. MR1704276 <https://doi.org/10.1007/s002200050646>
- [25] ZEITOUNI, O. (2016). Branching random walks and Gaussian fields. In *Probability and Statistical Physics in St. Petersburg. Proc. Sympos. Pure Math.* **91** 437–471. Amer. Math. Soc., Providence, RI. MR3526836