Sharp mixing time bounds for sampling random surfaces

Pietro Caputo
Dipartimento di Matematica, Università Roma Tre
Rome, Italy
Email: caputo@mat.uniroma3.it

Fabio Martinelli
Dipartimento di Matematica, Università Roma Tre
Rome, Italy
Email: martin@mat.uniroma3.it

Fabio Lucio Toninelli
Laboratoire de Physique
CNRS and ENS Lyon
Lyon, France
Email: fabio-lucio.toninelli@ens-lyon.fr

***Abstract*—We analyze the mixing time of a natural local Markov Chain (Gibbs sampler) for two commonly studied models of random surfaces: (i) discrete monotone surfaces with "almost planar" boundary conditions and (ii) the one-dimensional discrete Solid-on-Solid (SOS) model. In both cases we prove the first almost optimal bounds. Our proof is inspired by the so-called "mean curvature" heuristic: on a large scale, the dynamics should approximate a deterministic motion in which each point of the surface moves according to a drift proportional to the local inverse mean curvature radius. Key technical ingredients are monotonicity, coupling and an argument due to D. Wilson [17] in the framework of lozenge tiling Markov Chains. The novelty of our approach with respect to previous results consists in proving that, with high probability, the dynamics is dominated by a deterministic evolution which follows the mean curvature prescription. Our method works equally well for both models despite the fact that their equilibrium maximal deviations from the average height profile occur on very different scales.**

***Keywords*-Monte Carlo Markov chains (MCMC), mixing time, lozenge tilings, monotone surfaces, mean curvature. spectral gap, Glauber dynamics.**

I. Introduction

This work is concerned with the analysis of local Markov chains for random planar combinatorial structures and other random surface models arising in statistical physics. To fix ideas, below we consider two models in detail: the lozenge tiling model and the Solid-on-Solid (SOS) model. We develop a unifying framework that allows us to obtain for the first time sharp mixing time bounds.

The first structures we consider are lozenge tilings of a finite region of the triangular lattice. Each lozenge is a unit rhombus with angles of $60°$ and $120°$, covering two adjacent triangles in the lattice, with three possible orientations. There is a well known correspondence between lozenge tilings and dimer coverings (perfect matchings) of the dual hexagonal lattice; see e.g. [10], [17]. Another equivalent description is in terms of plane partitions, i.e. arrays of integers with the property that all rows and columns are non-increasing. The latter is naturally interpreted as a monotone collection of stacks of unit cubes, that is a (discrete) *monotone surface* $\phi_x \in \mathbb{Z}$, $x \in \mathbb{Z}^2$, with $\phi_y \geq \phi_x$ if $x = (x^{(1)}, x^{(2)})$ and $y = (y^{(1)}, y^{(2)})$ satisfy $x^{(a)} \leq y^{(a)}$, $a = 1, 2$.

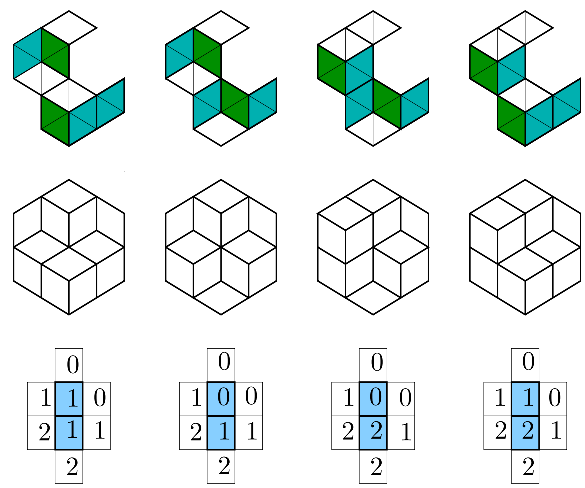

Figure 1. The top line displays the four possible tilings of a region V of the triangular lattice. The middle line shows the corresponding plane partitions. The bottom line represents (modulo a trivial rotation) the region $U \cup \partial U \subset \mathbb{Z}^2$ and the height of the monotone surface in each of the four cases. Note that here $U = \{(0,0),(0,1)\}$ consists of two points, and $\partial U = \{(-1,0),(-1,1),(0,2),(0,-1),(1,0),(1,1)\}$, with boundary condition $\phi_{(-1,0)} = 2$, $\phi_{(-1,1)} = 1$, $\phi_{(0,2)} = 0$, $\phi_{(0,-1)} = 2$, $\phi_{(1,0)} = 1$, $\phi_{(1,1)} = 0$.

Given a finite, tileable subset V of the triangular lattice, there exists a finite subset U of $\mathbb{Z}^2$ such that lozenge tilings of V are in one-to-one correspondence with monotone surfaces $\{\phi_x\}_{x \in U \cup \partial U}$, where ∂U is the external boundary of U in $\mathbb{Z}^2$. The height values $\{\phi_x\}_{x \in \partial U}$ are referred to as the *boundary condition*. The boundary condition is uniquely specified by the region V of the triangular lattice, i.e. the monotone surface representations of any two lozenge tilings of V will share the same boundary heights; see Figure 1.

The following algorithm is a standard way to sample uniformly among all monotone surfaces in U compatible with the boundary condition (or equivalently among all lozenge tilings of the region V). At each step, pick a random vertex $x \in U$ and update the value of ϕ_x to a new value $\phi'_x = \phi_x \pm 1$ according to a fair coin toss. If ϕ' (together with the boundary condition) is not a monotone surface on $U \cup \partial U$, then the update is rejected. For our purposes it is convenient to work with the continuous time version of this algorithm where each $x \in U$ is updated at the arrival times of an independent Poisson clock of parameter one. It is standard that the mixing time of the discrete time algorithm is of the order of the mixing time $\mathrm{T_{mix}}$ of the continuous time version multiplied by $|U|$, the size of the region; see e.g. [8].

It is generally expected that $\mathrm{T_{mix}} = O(L^2 \log L)$, where L is the diameter of U. Non-optimal polynomial bounds on $\mathrm{T_{mix}}$ were first obtained in the pioneering works [10], [17], [14] via coupling and comparison methods, by using an auxiliary non-local Markov chain. These methods are quite robust and can be applied to rather general boundary conditions, but the best they can provide is $\mathrm{T_{mix}} = \widetilde{O}(L^4)$[1]. In our work we restrict to the case of *approximately planar boundary conditions*. This corresponds to a choice of the boundary values such that the points $\{(x, \phi_x)\}_{x \in \partial U}$ approximately lie on a given plane of $\mathbb{R}^3$ (see Figure 2(b)). Such "planar" case is rather natural in the statistical physics context. It corresponds to the three-dimensional Ising model at zero temperature in the cylinder $U \times \mathbb{Z}$, with so-called Dobrushin boundary conditions, i.e. spins outside the cylinder are "+" above some plane and "−" below it.

The second structure that will be considered is the SOS model in a $L \times L$ box. A configuration is an assignment of an integer *height* $\phi_i \in \{0, 1, ..., L\}$ to each integer $0 \le i \le L+1$, with fixed boundary conditions $\phi_0 = \phi_{L+1} = 0$ (see figure 2(a)). The probability of a configuration is given by the *Gibbs distribution*

$$\pi(\phi) = \pi_L(\phi) := Z_L^{-1} \exp\left\{ -\sum_{i=1}^{L+1} |\phi_{i-1} - \phi_i| \right\}. \quad (1)$$

Here Z_L is a normalizing factor (the "partition function"). In statistical physics, π describes the distribution of the unique open contour in the two-dimensional Ising model in the box $\{1, \ldots, L\} \times \{1, \ldots, L\}$ with Dobrushin boundary conditions (spins outside the box are "−" along the bottom horizontal side and "+" elsewhere), in the limit where the couplings on vertical edges tend to infinity.

Configurations of the SOS model are sampled by means of a continuous time Markov chain defined as follows. Independently at each site i, the height ϕ_i is updated at the arrival times of a mean one Poisson clock and it is then replaced by $\phi_i \pm 1$ with probabilities $p_{i,\pm}$ given by the usual Gibbs sampler rule (see (4)). A long standing conjecture is that $\mathrm{T_{mix}} = O(L^2 \log L)$. Recently there has been important progress in this direction [11] with the bound $\mathrm{T_{mix}} = \widetilde{O}(L^{5/2})$.

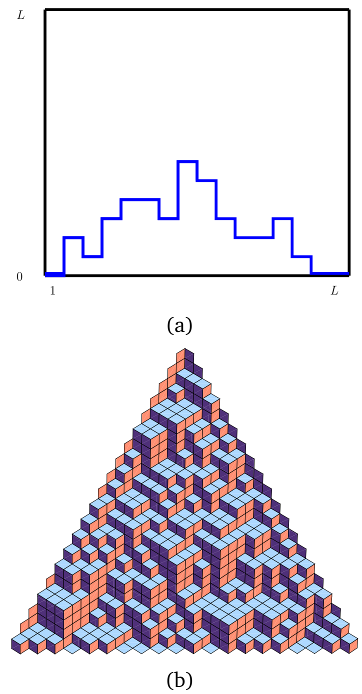

Figure 2. (a) A configuration of the SOS model in the $L \times L$ box. (b) An example of a monotone surface with planar boundary conditions on the plane $x + y + z = 0$. The picture is taken from [6]

Our main result on the mixing time for both structures is informally stated as follows.

Theorem. *(i) For the Solid-on-Solid model the mixing time satisfies* $\mathrm{T_{mix}} = \widetilde{O}(L^2)$. *(ii) For monotone surfaces with approximately planar boundary conditions the mixing time satisfies* $\mathrm{T_{mix}} = \widetilde{O}(L^2)$.

An important motivation for studying these problems is the fact in order to get optimal bounds one is forced to go beyond standard techniques such as coupling and comparison: new geometric insight on the typical equilibration process is required. Moreover, proving optimal mixing time bounds leads to a deeper understanding of the equilibration mechanism and, potentially, may bring important mathematical progress in certain problems of non-equilibrium statistical physics. Indeed, the conjectured behavior $\mathrm{T_{mix}} = O(L^2 \log L)$ is based on the idea that each point of the surface feels a drift proportional to the local inverse mean curvature radius. This *mean curvature* heuristic plays an important role in the analysis of the dynamics of interfaces in statistical physics models; see e.g. [16].

A high level description of the proof of the theorem, together with the connection with the above mentioned

[1]Throughout this paper the notation $\widetilde{O}(\cdot)$ hides factors of polylog(L).

heuristic, is given in Section III.

Our method can be potentially applied to a wide range of stochastic surface dynamics models where mean curvature motion is expected to occur macroscopically. An example that comes naturally to mind is the dynamics of domino tilings of the plane, for which at present only non-optimal polynomial upper bounds on $\mathrm{T}_{\mathrm{mix}}$ are available [10].

For the monotone surface model it is natural to ask what happens without the planarity assumption on the boundary conditions. In this case the typical equilibrium shape is not macroscopically flat (in contrast to the planar boundary case, cf. Proposition 1) and arctic circle type phenomena can occur [4]. While in principle one expects again that $\mathrm{T}_{\mathrm{mix}} = \widetilde{O}(L^2)$, what is missing here are precise, finite-L equilibrium estimates on height fluctuations and on the convergence of the equilibrium average height to its macroscopic limit.

We conclude this introduction with a short review of previous work and its relationship with the results presented here. The first polynomial bound on the mixing time for sampling monotone surfaces was established in [10] for a *non-local* dynamics in which large piles of unit cubes can appear/disappear in one move. A refined analysis of the non-local dynamics of [10] was later obtained in [17] who proved tight bounds for its mixing time. In [14] polynomial mixing time bounds for the *local* dynamics (i.e. the one we consider) were obtained by means of comparison with the non-local chain. Despite these remarkable progresses, sharp bounds on the mixing time of the local dynamics remained elusive, mostly due to the lack of precision of the comparison arguments. Concerning the SOS model, the best result so far is due to [11]. There, non-equilibrium geometric information was introduced for the first time to get non-trivial (but not tight) bounds. Similar ideas found later application to Glauber dynamics for the Ising model at low temperature [12], [9]. The present paper adds new important insight of the role of mean curvature, as opposed to [11] where less precise bounding dynamics were used.

II. Models and results

In this section we give precise definitions and formulate our main findings in detail.

A. Monotone surfaces with "planar" boundary conditions

A function $\phi : \mathbb{Z}^2 \mapsto (\mathbb{Z} \cup \{\pm\infty\})$ defines a (discrete) monotone surface if $\phi_x \geq \phi_y$ whenever $x^{(a)} \leq y^{(a)}, a = 1, 2$, where $x^{(a)}$ denotes the a^{th} component of x. The collection of all monotone surfaces is denoted by Ω. For $U \subset \mathbb{Z}^2$ we write ϕ_U for the restriction of ϕ to U. We consider the following local, continuous time Markov chain $\{\phi^\xi(t)\}_{t\geq 0}$ on monotone surfaces with initial condition ξ and fixed boundary conditions (b.c.) outside a finite region. Let U be a finite connected subset of $\mathbb{Z}^2$ (the finite region) and $\eta \in \Omega$ (the boundary condition). Without loss of generality we will always assume that U contains the origin. Given $\phi \in \Omega$ and $x \in \mathbb{Z}^2$, let $\phi^{(x,\pm)}$ be the monotone surfaces which coincide with ϕ outside x while $\phi^{(x,\pm)}_x = \phi_x \pm 1$ if compatible with the monotonicity constraints and $\phi^{(x,\pm)}_x = \phi_x$ if not. The state space of the Markov chain is the set

$$\Omega_{\eta,U} := \{\phi \in \Omega : \phi_x = \eta_x \text{ for } x \notin U\}.$$

The initial condition at time zero is some given $\xi \in \Omega_{\eta,U}$. To each $x \in U$ is assigned an independent exponential clock of rate 1. If the clock labeled x rings at time t, we replace $\phi(t)$ with $[\phi(t)]^{(x,+)}$ or $[\phi(t)]^{(x,-)}$ by tossing a fair coin. It is standard to check that such Markov chain is irreducible and reversible with respect to the uniform measure on $\Omega_{\eta,U}$, which we denote π^η_U or simply π. In the sequel we will write μ^ξ_t for the distribution of the chain started from ξ at time t. The mixing time is then defined as usual by

$$\mathrm{T}_{\mathrm{mix}} = \inf\big\{t > 0 : \sup_{\xi\in\Omega_{\eta,U}} \|\mu^\xi_t - \pi\| \leq (2e)^{-1}\big\}. \tag{2}$$

where $\|\cdot\|$ denotes the total variation distance; see e.g. [8].

As we mentioned in the introduction, it is conjectured that $\mathrm{T}_{\mathrm{mix}} = O(L^2 \log L)$, where L is the diameter of the region U. We prove this conjecture, up to poly-log corrections, under the assumption that the boundary conditions are "approximately planar". The planarity assumption on the boundary conditions is specified as follows:

Definition 1. *Given $\mathbf{n} \in \mathbb{R}^3$ with $\|\mathbf{n}\| = 1$, we write $\mathbf{n} > 0$ if $\mathbf{n}^{(i)} > 0, i = 1, 2, 3$. We let $\bar\phi^{\mathbf{n}} \in \Omega$ be the discrete monotone surface with slope $\mathbf{n}$:*

$$\mathbb{Z}^2 \ni x \mapsto \bar\phi^{\mathbf{n}}_x = \max\{z \in \mathbb{Z} : x^{(1)}\mathbf{n}^{(1)} + x^{(2)}\mathbf{n}^{(2)} + z\mathbf{n}^{(3)} \leq 0\}$$

We say that η is a good planar boundary condition *with slope $\mathbf{n} > 0$ if there exists $C > 0$ such that*

$$|\eta_x - \bar\phi^{\mathbf{n}}_x| \leq C \log(|x| + 1). \tag{3}$$

for every $x \in \mathbb{Z}^2$.

Note that our definition (3) allows a bit of flexibility with respect to the fully planar case $\eta_x \equiv \bar\phi^{\mathbf{n}}_x$. Roughly speaking, the reason why this cannot be relaxed beyond the logarithmic scale is because we only allow poly-logarithmic errors in the bounds of $\mathrm{T}_{\mathrm{mix}}$.

A key feature of the model is that, in the case of good planar boundary conditions, a typical surface at equilibrium is essentially flat. This fact, which we shall exploit throughout, is detailed in the following proposition.

Proposition 1. *Let η be a good planar boundary condition with slope $\mathbf{n} > 0$. For every $\epsilon > 0$ there exists $c > 0$ such that for every $a > 0$, for every connected $U \subset \mathbb{Z}^2$ containing the origin and with diameter L,*

$$\pi_U^\eta \left(\exists y \in U : |\phi_y - \bar\phi_y^{\mathbf{n}}| \geq a(\log L)^{1+\epsilon}\right) = O(e^{-c\,a(\log L)^{1+\epsilon}}).$$

The above estimate is based on the well-known representation of the height function in terms of the number of points of a determinantal point process whose kernel is explicitly known [6], [7]. A detailed proof requires a delicate use of monotonicity properties of the equilibrium measure, see [3, Theorem 1]. We can finally formulate our mixing time upper bound:

Theorem 1. *With the same assumptions as in Proposition 1 one has $\mathrm{T}_{\mathrm{mix}} = \widetilde{O}(L^2)$.*

We remark that Theorem 1 is essentially tight up to poly-log corrections. This can be seen by using the idea behind the mixing time lower bound of Theorem 3.1 in [2]. One could in fact show that $\mathrm{T}_{\mathrm{mix}} \geq L^2/(c\log L)$ for a suitable $c > 0$ for instance when $U = \mathcal{U}_L \cap \mathbb{Z}^2$, with $\mathcal{U}_L$ a smooth open connected set of $\mathbb{R}^2$ expanded by a factor L.

B. SOS model

The configuration space of the SOS model is $\Omega_L = \{0, 1, .., L\}^L$ and the equilibrium measure $\pi = \pi_L$ is given in (1). Standard random walk considerations show that a typical SOS interface is macroscopically flat. More precisely, one has (see [11, Appendix C]):

Lemma 1. *There exists some constant $c > 0$ such that, uniformly in $H \geq \sqrt{L}\log L$:*

$$\pi\big[\exists\, i = 1, \ldots, L : \ \phi_i > H\big] = O(\exp\left(-c\min\{H^2/L, H\}\right)).$$

The Markov chain to be considered is the standard continuous time Gibbs sampler. There are independent Poisson clocks with mean 1 at each site $i \in \{1, \ldots, L\}$. When the clock at site i rings, the height ϕ_i is updated to the new value $\max\{\phi_i - 1, 0\}$ or $\min\{\phi_i + 1, L\}$ with probabilities $p_{i,-}(\phi)$, $p_{i,+}(\phi)$ respectively, determined by:

$$p_{i,-}(\phi) = \frac{e^{-2}}{1+e^{-2}}\,\mathbb{1}_{\{\phi_i \leq a\}} + \frac{1}{2}\,\mathbb{1}_{\{b \geq \phi_i > a\}} + \frac{1}{1+e^{-2}}\,\mathbb{1}_{\{\phi_i > b\}}$$
$$p_{i,+}(\phi) = \frac{e^{-2}}{1+e^{-2}}\,\mathbb{1}_{\{\phi_i \geq b\}} + \frac{1}{2}\,\mathbb{1}_{\{b > \phi_i \geq a\}} + \frac{1}{1+e^{-2}}\,\mathbb{1}_{\{\phi_i < a\}} \tag{4}$$

where $a := \min\{\phi_{i-1}, \phi_{i+1}\}$ and $b := \max\{\phi_{i-1}, \phi_{i+1}\}$. With the remaining probability $1 - (p_{i,-}(\phi) + p_{i,+}(\phi))$, ϕ_i stays at its current value. It is not hard to check that the stationary reversible measure is π_L. The mixing time of this chain is defined as in (2).

Theorem 2. *The mixing time of the SOS model satisfies $\mathrm{T}_{\mathrm{mix}} = \widetilde{O}(L^2)$.*

A possible extension of this result is obtained by letting the constraints $0 \leq \phi_i \leq L$ be replaced by $0 \leq \phi_i \leq M$, where $M \in \mathbb{N}$ is an independent parameter. We refer to this as the $L \times M$ SOS model. One can extend the proof of Theorem 2 to get in this case the essentially sharp bound $\mathrm{T}_{\mathrm{mix}} = \widetilde{O}(L\max(L, M))$. On the other hand, one can obtain the upper bound $\widetilde{O}(L^2)$ on the *relaxation time*, i.e. the inverse of the spectral gap of of the $L \times M$ SOS model *uniformly* in M (see Theorem 4 in [3]). Establishing this behavior has been a long standing problem mainly because the equilibrium distribution is not strictly log-concave. Well known recursive approaches which fail here, would work if the function $|x|$ appearing in the definition of the SOS equilibrium distribution were replaced by one with strictly convex behavior at infinity.

III. High level description of the proof

As announced in the introduction, our bound $\mathrm{T}_{\mathrm{mix}} = \widetilde{O}(L^2)$ is proved by following a common strategy that we sketch here. First note that both models share the following monotonicity property. Define the partial order $\preceq$ on the state space by declaring $\phi \preceq \phi'$ iff at each vertex x one has $\phi_x \leq \phi'_x$. Clearly $\preceq$ admits unique maximal and minimal elements denoted respectively $\wedge$ and $\vee$. The key fact is that there exists a *grand coupling* for the dynamics which is monotone w.r.t. $\preceq$. Informally, this means that one can couple the dynamics corresponding to all possible initial conditions in such a way that if two initial conditions are ordered, then the same ordering is preserved at any later time (see e.g. [8, Chapter 5]). A key consequence is that it is enough to consider the evolution of the extremal configurations $\wedge, \vee$. The crucial step in the proof of the mixing time upper bounds is then the following (see Propositions 2 and 6 below for a precise formulation in the case of monotone surfaces and SOS model):

Step 1. *Starting from either $\wedge$ or $\vee$, after time $\widetilde{O}(L^2)$ the distance between the interface and the flat profile is not larger than its typical equilibrium value χ_L (thanks to Proposition 1 and Lemma 1, $\chi_L = \widetilde{O}(1)$ for monotone surfaces and $\chi_L = \widetilde{O}(L^{1/2})$ for the SOS model).*

Here, "flat profile" refers to the configuration $\bar\phi^{\mathbf{n}}$ (cf. Definition 1) in the monotone surface case and to the identically zero configuration $\phi_x = 0, i = 0, \ldots, L+1$ for SOS. Once Step 1 is accomplished, one concludes $\mathrm{T}_{\mathrm{mix}} = \widetilde{O}(L^2)$ provided that a result of the following type holds:

Step 2. *If the initial condition ξ is at distance at most χ_L from the flat profile, then $\|\mu_T^\xi - \pi\| \ll 1$ for some*

$T = \widetilde{O}(L^2)$.

The proof of step 2 is model-dependent: for the SOS model it was given in [11] and for monotone surfaces it follows from results in [2] (see Propositions 3 and 4 below) .

The key ingredient in establishing Step 1 is to show that, with high probability, the surface started from the maximal configuration $\wedge$ stays below a *deterministic interface evolution* which after time $\widetilde{O}(L^2)$ is at the correct distance $O(\chi_L)$ from the flat profile. It turns out that it is actually sufficient to specify the deterministic interface evolution at a sequence of deterministic times $t_n, 0 \le n \le M$. For clarity we describe the geometric construction only for the SOS model. At the end we will mention the modifications needed for the monotone surface case. At all times t_n, the deterministic interface is the boundary of $\mathcal{C}_{u_n}$ where, given $u > 0$, $\mathcal{C}_u$ is a circular segment of height u and base of linear size ρ_L roughly of order L, see Figure 3. The base of $\mathcal{C}_u$ lies on the line which contains the macroscopic flat profile. The evolution of $\mathcal{C}_u$, by a kind of "flattening process", in the time interval $(t_{n-1}, t_n]$ transforms $\mathcal{C}_{u_{n-1}}$ into $\mathcal{C}_{u_n}$. The sequence of increasing times $\{t_n\}_{0\le n\le M}$ and of decreasing heights $\{u_n\}_{0\le n\le M}$ will be introduced in a moment. The "domination statement" then is of the following type (see Propositions 5 and 8):

Claim 1. *For all $0 \le n \le M$, with high probability the following holds. For all times in $[t_n, L^3]$ the evolution started from the maximal configuration $\wedge$ stays below the boundary of $\mathcal{C}_{u_n}$.*

The initial height u_0 is taken to be proportional to L and one sets $t_0 = 0$; this guarantees that the statement of Claim 1 holds trivially for $n = 0$. In order to choose u_{n+1} given u_n, one uses the following procedure. Consider the circular segment $\mathcal{C}_{u_n}$ and choose a point on its curved boundary (e.g. the highest one). Move inward (i.e. inside $\mathcal{C}_{u_n}$) the tangent line at the chosen point by an amount Δ and call d_Δ the diameter of the intersection between the line with $\mathcal{C}_{u_n}$; see Figure 3. Then $u_n - u_{n+1}$ is chosen as the critical value Δ_n such that the equilibrium fluctuations on scale d_{Δ_n} are of order Δ_n (apart from logarithmic corrections), i.e. $\Delta_n = \widetilde{O}(\chi_{d_{\Delta_n}})$. Also, M is the smallest index such that $u_M \le \chi_L$, i.e. u_M is of the order of the equilibrium height fluctuations on scale L. As for the time sequence $\{t_n\}_{0\le n\le M}$, one sets $t_{n+1} - t_n$ to be of order $d^2_{\Delta_n}$ (again neglecting logarithmic corrections): that this is the correct choice is guaranteed by a careful use of Step 2, applied with $L = d_{\Delta_n}$. It is not difficult to see that $t_M = \widetilde{O}(L^2)$. Indeed $\chi_L \sim L^\gamma$ with $\gamma = 1/2$ (but the arguments below apply to *any* $0 \le \gamma < 1$, where $\gamma = 0$ means that $\chi_L \sim \text{polylog}(L)$). Then, simple geometric

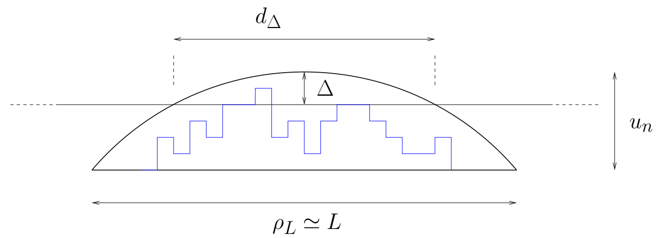

Figure 3. The circular segment $\mathcal{C}_{u_n}$. The deterministic evolution at time t_n coincides with the curved portion of the boundary of $\mathcal{C}_{u_n}$. At that time, the true stochastic SOS configuration (wiggled line) stays with high probability below it. Elementary geometry shows that $\Delta \simeq d^2_\Delta (u_n/\rho_L^2)$. The requirement $u_n - u_{n+1} = \Delta_n \simeq d^\gamma_{\Delta_n}$ then leads to $u_n - u_{n+1} \simeq (\rho_L^2/u_n)^{\gamma/(2-\gamma)}$.

considerations show that

$$u_n - u_{n+1} = \widetilde{O}\left(\left(\rho_L^2/u_n\right)^{\gamma/(2-\gamma)}\right),$$
$$t_{n+1} - t_n = \widetilde{O}\left(\left(\rho_L^2/u_n\right)^{2/(2-\gamma)}\right).$$

Approximating the recursion for u_n with a differential equation gives

$$u_n^{2/(2-\gamma)} \simeq u_0^{2/(2-\gamma)} - \rho_L^{2\gamma/(2-\gamma)} n$$
$$\simeq L^{2\gamma/(2-\gamma)}\left[L^{(2-2\gamma)/(2-\gamma)} - n\right]$$

since both u_0 and ρ_L are of order L. In particular, one has roughly $M = O(L^{(2-2\gamma)/(2-\gamma)})$. Then,

$$t_M = \sum_{n=0}^{M-1}(t_{n+1} - t_n) \simeq \rho_L^{4/(2-\gamma)} \sum_{n=0}^{M-1} \frac{1}{u_n^{2/(2-\gamma)}}$$
$$\simeq \frac{\rho_L^{4/(2-\gamma)}}{L^{2\gamma/(2-\gamma)}} \sum_{n=0}^{M-1} \frac{1}{L^{(2-2\gamma)/(2-\gamma)} - n} = \widetilde{O}(L^2)$$

since the last sum is of order $\log L$. Remarkably, the order of magnitude of t_M does not depend on the specific numerical value $1/2$ of the fluctuation exponent γ, while the sequence (t_n, u_n) and the value of M do.

The statement of Claim 1 for $n = M$ allows us to conclude Step 1. The evolution started from the maximal configuration, at time $t_M = \widetilde{O}(L^2)$, is below the deterministic evolution, which is within distance χ_L from the flat profile. A similar result holds starting from the minimal configuration $\vee$. In the monotone surface case, the above argument can be repeated without modifications provided that the circular segment $\mathcal{C}_u$ is replaced by a spherical cap of height u and base of size ρ_L on the plane containing the macroscopic flat profile and the exponent γ is set equal to zero.

Another way to understand the choice of the time scales t_n is the following. If one imagines that the boundary of $\mathcal{C}_u$ evolves by "mean curvature", i.e. feeling a inward drift proportional to the inverse of its instantaneous radius of curvature, then the time $t_{n+1} - t_n$ to transform $\mathcal{C}_{u_n}$ into $\mathcal{C}_{u_{n+1}}$ must be $O(R_n \times (u_n - u_{n+1}))$, where R_n is the radius of curvature of $\mathcal{C}_{u_n}$. One can easily check that, apart from logarithmic corrections, this coincides with the requirement $t_{n+1} - t_n \simeq d^2_{\Delta_n}$.

IV. Mixing time upper bound for monotone surfaces

A. Formalizing the strategy of Section III

Recall that $\phi^\xi(t)$ denotes the configuration at time t started from ξ, μ_t^ξ denotes its distribution and that $\wedge$ and $\vee$ denote the maximal and minimal configuration w.r.t. the partial order $\preceq$. The first key ingredient (corresponding to Step 1 in Section III) is a result saying that, after time of order $\widetilde{O}(L^2)$, the surface is at most at distance $(\log L)^{3/2}$ away from $\bar{\phi}^{\mathbf{n}}$ (cf. Definition 1). It is here that the new ideas based on mimicking the evolution by mean curvature play a crucial role.

Proposition 2. *There exists $c > 0, \beta > 0$ such that, for $T = L^2(\log L)^\beta$ one has*

$$\mathbb{P}\left(\max_{x\in U}(\phi_x^\wedge(T) - \bar{\phi}_x^{\mathbf{n}})) > (\log L)^{3/2}\right) = O\left(e^{-c(\log L)^{3/2}}\right)$$

The same estimate holds for the probability of the event $\{\min_{x\in U}(\phi_x^\vee(T) - \bar{\phi}_x^{\mathbf{n}}) < -(\log L)^{3/2}\}$.

The second result says that once the surface is within distance $(\log L)^{3/2}$ from $\bar{\phi}^{\mathbf{n}}$, it does not go much farther than that for a time much longer than T. This second step is more standard and its proof (see Appendix) combines monotonicity and reversibility together with the fluctuation bounds of Proposition 1. Define the set

$$\Omega' = \{\phi \in \Omega_{\eta,U} : \max_{x\in U} |\phi_x - \bar{\phi}_x^{\mathbf{n}}| \le 2(\log L)^{3/2}\}. \quad (5)$$

Proposition 3. *Let $\xi \in \Omega_{\eta,U}$ be such that*

$$\max_{x\in U} |\xi_x - \bar{\phi}_x^{\mathbf{n}}| \le (\log L)^{3/2}. \quad (6)$$

Then, there exists $c > 0$ such that

$$\mathbb{P}\left(\exists t < L^{10} : \phi^\xi(t) \notin \Omega'\right) = O\left(e^{-c(\log L)^{3/2}}\right).$$

The last step concerns the evolution constrained between two monotone configurations $\phi^- \preceq \phi^+$ (acting as "ceiling" and "floor" respectively). More precisely, the constrained dynamics is obtained by rejecting every move which would violate the condition $\phi^- \preceq \phi^\xi(t) \preceq \phi^+$.

Proposition 4. *[2, Theorem 4.3] Let $\phi^\pm \in \Omega_{\eta,U}$ with $\phi^- \preceq \phi^+$. For the dynamics constrained between ϕ^+ and ϕ^- one has*

$$\mathrm{T}_{\mathrm{mix}} = O(L^2(\log L)^2 H^2)$$

where $H = \max_{x\in U}(\phi_x^+ - \phi_x^-)$.

The combination of Proposition 3 and Proposition 4 correspond to Step 2 in Section III. We can now easily put together Propositions 2 to 4 to obtain the desired bound $\mathrm{T}_{\mathrm{mix}} = \widetilde{O}(L^2)$:

Proof of Theorem 1: It is known that

$$\max_{\xi\in\Omega_{\eta,U}} \|\mu_t^\xi - \pi\| \le 2L^3 \max(\|\mu_t^\wedge - \pi\|, \|\mu_t^\vee - \pi\|)$$

(see e.g. [2, Lemma 6.2] for a similar statement). Therefore, it is sufficient to prove that

$$\max(\|\mu_t^\wedge - \pi\|, \|\mu_t^\vee - \pi\|) \le 1/(4eL^3)$$

for $t = 2L^2(\log L)^\beta = 2T$. Let us consider e.g. the case of the maximal initial condition $\wedge$, the other case being analogous. Define $\Omega'' = \{\phi \in \Omega_{\eta,U} : \max_{x\in U} |\phi_x - \bar{\phi}_x^{\mathbf{n}}| \le (\log L)^{3/2}\}$ and, for $\xi \in \Omega''$,

$$\tau = \inf\{t > 0 : \max_{x\in U} |\phi_x^\xi(t) - \bar{\phi}_x^{\mathbf{n}}| \ge 2(\log L)^{3/2} - 1\}.$$

Let A be a subset of $\Omega_{\eta,U}$. Then, using Proposition 2,

$$\mu_{2T}^\wedge(A) = \mu_T^\wedge(\mu_T^\xi(A)|\xi \in \Omega'') + O\left(e^{-c(\log L)^{3/2}}\right).$$

Next, from Proposition 3, one has for every $\xi \in \Omega''$

$$\begin{aligned}\mu_T^\xi(A) &= \mathbb{P}(\phi^\xi(T) \in A; \tau > T) + O\left(e^{-c(\log L)^{3/2}}\right) \\ &= \mathbb{P}'(\phi^\xi(T) \in A; \tau > T) + O\left(e^{-c(\log L)^{3/2}}\right) \\ &= \mathbb{P}'(\phi^\xi(T) \in A) + O\left(e^{-c(\log L)^{3/2}}\right)\end{aligned}$$

where $\mathbb{P}'$ denotes the law of the dynamics restricted to the set Ω' of (5). Indeed, the two laws $\mathbb{P}$ and $\mathbb{P}'$ can be coupled so that the corresponding trajectories coincide up to the random time τ. In particular, τ has the same law under $\mathbb{P}$ and $\mathbb{P}'$. Finally, thanks to Proposition 4 and taking β sufficiently large, T is at least $(\log L)^2$ times the mixing time of the restricted dynamics (which is $\widetilde{O}(L^2)$). Therefore, from the standard sub-multiplicative property

$$\sup_{\xi\in\Omega} \|\mu_t^\xi - \pi\| \le e^{-\lfloor t/\mathrm{T}_{\mathrm{mix}}\rfloor},$$

and the fact that the invariant measure of the restricted dynamics is $\pi(\cdot|\Omega')$, one has

$$|\mathbb{P}^{\Phi^\pm}(\phi^\xi(T) \in A) - \pi(A|\Omega')| \le e^{-(\log L)^2}.$$

Thanks to Proposition 1 one has $\pi(\Omega') \ge 1 - O(\exp(-c(\log L)^{3/2}))$ and therefore $\pi(A|\Omega') = \pi(A) + O(\exp(-c(\log L)^{3/2}))$. Finally,

$$|\mu_{2T}^\wedge(A) - \pi(A)| = O\left(e^{-c(\log L)^{3/2}}\right)$$

for every event $A \in \Omega_{\eta,U}$, which implies $\|\mu_{2T}^\wedge - \pi\| \le 1/(4eL^3)$ for L large enough. ■

B. Proof of Proposition 2

Let $\bar\Pi$ be the plane of slope $\mathbf{n}$, obtained by translating upwards by a distance $C\log L$ the plane $\Pi^{\mathbf{n}} := \{z \in \mathbb{R}^3 : \mathbf{n}^{(1)}z^{(1)} + \mathbf{n}^{(2)}z^{(2)} + \mathbf{n}^{(3)}z^{(3)} = 0\}$, with C the same constant as in (3). Let W be a disk of radius

$$\rho_L = L \times \log L$$

on $\bar\Pi$ such that its projection V on the horizontal plane at height zero contains U and moreover the distance between ∂U and ∂V is at least $\rho_L/2$ (recall that U has diameter L).

Given $u > 0$, let $\mathcal{C}_u$ be the spherical cap whose base is the disk W and whose height is u. The radius of curvature R is related to u and ρ_L by

$$(2R - u)u = \rho_L^2 \tag{7}$$

and, since we will always work under the condition $u \ll \rho_L \ll R$, we have $R = \rho_L^2/(2u)(1 + o(1))$. For a point v on the curved portion of the boundary of $\mathcal{C}_u$, let $\mathbf{n}_v$ be the normal at v directed towards the exterior of $\mathcal{C}_u$. It is clear that, if $u \ll \rho_L$, one has $\mathbf{n}_v = \mathbf{n}+o(1)$; in particular, $\mathbf{n}_v > 0$ with the convention of Definition 1. Finally the height (w.r.t. the horizontal plane) of the spherical cap at horizontal coordinates $x \in V$ is denoted by

$$\psi_u(x) = \sup\{z \in \mathbb{R} : (x^{(1)}, x^{(2)}, z) \in \mathcal{C}_u\}.$$

We now define a sequence of spherical caps $\{\mathcal{C}_{u_n}\}_{n=0}^M$ with constant base W, decreasing height u_n and increasing radius of curvature R_n. More precisely, let $u_0 = 2L$ and $M := 2L - (\log L)^{5/4}$. Then we let $u_n = u_{n-1} - 1 = 2L - n$ and $(2R_n - u_n)u_n = \rho_L^2$. We also define

$$t_n = t_{n-1} + R_n(\log L)^\alpha, \quad t_0 = 0, \tag{8}$$

where $\alpha > 0$ is a constant to be determined later. It is worth noting that $R_0 \sim L(\log L)^2/4$, $R_M \sim L^2/(2(\log L)^{3/4})$ and $R_{n+1}/R_n = 1 + o(1)$ uniformly in the whole range $n = 0, \ldots, M$. Recalling that $R_n = \rho_L^2/(2u_n)(1 + o(1))$, where $o(1)$ is small uniformly in $1 \le n \le M$, it is immediate to deduce that

$$t_M = \rho_L^2(\log L)^\alpha \times O\Big(\sum_{n\le M} 1/u_n\Big) = \widetilde{O}(L^2). \tag{9}$$

With this notation the key step is represented by the next proposition, which corresponds to Claim 1 in Section III.

Proposition 5. *There exists a positive constant c' and a suitable choice of the parameter α in* (8) *such that the following holds for L large enough. For every $0 \le n \le M$ one has, with probability at least $1-n\exp(-c'(\log L)^{3/2})$, that, for every $x \in U$ and every $t \in [t_n, L^3]$,*

$$\phi_x^\wedge(t) \le \psi_{u_n}(x)$$

Since $u_M \ll (\log L)^{3/2}$, Proposition 5 together with (9) completes the proof. The proof of Proposition 5 is deferred to the appendix.

V. Mixing time upper bound for the SOS model

The proof of Theorem 2 follows the strategy described in Section III. Step 1 and Step 2 are now formalized as Proposition 6 and Proposition 7 respectively. Recall that $\phi^\wedge(t)$ denotes the evolution of the maximal initial configuration $\wedge$, given by $\wedge_i = L$ for every $i \in [1, L]$.

Proposition 6. *There exists $T_1 = \widetilde{O}(L^2)$ such that*

$$\mathbb{P}\big(\exists i \in \{1,\ldots,L\}\,,\ \phi_i^\wedge(T_1) > \sqrt{L}\,(\log L)^4\big) = O(L^{-3}).$$

Proposition 7. *There exist a constant $c > 0$ and $T_2 = \widetilde{O}(L^2)$ such that*

$$\sup_{\xi\in\Omega'_L} \|\mu_t^\xi - \pi\| \le e^{-c(\log L)^2}, \quad \forall t \ge T_2\,,$$

where $\Omega'_L = \{\xi \in \Omega_L :\ \xi_i \le \sqrt{L}\,(\log L)^4,\ i = 1,\ldots,L\}$.

As for the monotone surface dynamics, it is not difficult to conclude the proof of Theorem 2 once Propositions 6 and 7 are available. Concerning the proof of Propositions 7 we point out that it is essentially an application of the results in [11, Section 4]; see also [3].

We turn to a sketch of the proof of Proposition 6. Let $\rho_L = L\log L$ and, for $u > 0$, let $\mathcal{C}_u \subset \mathbb{R}^2$, denote the circular segment of height u and base a segment on the horizontal axis, with length $2\rho_L$ and centered at $L/2$. We shall use the notation $\phi \in \mathcal{C}_u$ for any configuration ϕ such that $0 \le \phi_i \le \psi_u(i), i = 1,\ldots,L$ where $\psi_u(i) = \sup\{y \in \mathbb{R} :\ (i, y) \in \mathcal{C}_u\}$. Note that the radius of curvature R of the circular segment $\mathcal{C}_u$ satisfies $(2R - u)u = \rho_L^2$. As in Section IV, we will always have $u \ll \rho_L \ll R$ so that $R = \rho_L^2/(2u)(1 + o(1))$. Define recursively $u_0 = 2L, t_0 = 0$ and

$$u_{n+1} = u_n - (\rho_L^2/u_n)^{1/3}(\log L)^2 \tag{10}$$

$$t_{n+1} = t_n + (\rho_L^2/u_n)^{4/3}(\log L)^{\alpha'}, \tag{11}$$

for some constant $\alpha' > 0$ to be fixed below, and let $M = \min\{n :\ u_n \le \sqrt{L}(\log L)^4\}$. Clearly, if R_n denotes the radius of curvature of the circular segment $\mathcal{C}_{u_n}$, then $R_n = \rho_L^2/(2u_n)(1+o(1))$. Note that the length and time scales (u_n, t_n) are quite different from their analogue in the monotone surface case. The main reason is the different order of magnitude of the maximal equilibrium fluctuations in the two models: $\log L$ versus $\sqrt{L}$. However their value is determined by the common recipe which was described in Section III. Crucially, as in the monotone surface case, $t_M = \tilde{O}(L^2)$. This can be proved along the lines of the argument given in Section III. The key step in the proof is analogous to Proposition 5.

Proposition 8. *For a suitable choice of the parameter α' in* (10)*, for every $0 \le n \le M$, with probability at least $1 - n\, e^{-(\log L)^{3/2}}$,*

$$\phi^\wedge(t) \in \mathcal{C}_{u_n} \text{ for every } t \in [t_n, L^3].$$

If we now apply Proposition 8 with $n = M$ and use the fact that $t_M = \widetilde{O}(L^2)$ and $u_M \le \sqrt{L}(\log L)^4$, we obtain the desired claim (**??**). The proof of Proposition 8 in turn is based on the same kind of arguments used for the proof of Proposition 5. Details can be found in the Appendix.

Acknowledgements

We would like to thank A. Sinclair for several advices on the presentation. This work has been carried out while FLT was visiting the Department of Mathematics of the University of Roma Tre under the ERC Advanced Research Grant "PTRELSS".

References

[1] N. Cancrini, F. Martinelli, C. Roberto, C. Toninelli, *Kinetically constrained models*, Probab. Theory Related Fields **140**, 459-504, 2008

[2] P. Caputo, F. Martinelli, F. Simenhaus, F.L. Toninelli, *"Zero" temperature stochastic 3D Ising model and dimer covering fluctuations: a first step towards interface mean curvature motion*, Comm. Pure Appl. Math. **64**, 778–831, 2011

[3] P. Caputo, F. Martinelli, F.L. Toninelli, *Mixing times of monotone surfaces and SOS interfaces: a mean curvature approach*, arXiv:1101.4190

[4] H. Cohn, M. Larsen, J. Propp, *The shape of a typical boxed plane partition*, New York J. Math. **4** (1998), 137–165.

[5] T. Funaki, *Stochastic interface models*, Lectures on Probability Theory and Statistics, 103–274, Lecture notes in Math. **1869**, Springer, Berlin, 2005.

[6] R. Kenyon , *Lectures on dimers*, Statistical mechanics, 191–230, IAS/Park City Math. Ser., **16**, Amer. Math. Soc., Providence, RI, 2009.

[7] R. Kenyon, A. Okounkov, S. Sheffield, *Dimers and amoebae*, Ann. Math. **163** (2006), 1019–1056.

[8] D. Levin, Y. Peres, E. Wilmer, *Markov Chains and Mixing Times*, American Mathematical Society, Providence, RI, (2009).

[9] E. Lubetzky, F. Martinelli, A. Sly, F. L. Toninelli, *Quasi-polynomial mixing of the 2D stochastic Ising model with "plus" boundary up to criticality*, arXiv:1012.1271

[10] M. Luby, D. Randall, A. Sinclair, *Markov chain algorithms for planar lattice structures*, SIAM Journal on Computing, **31** (2001), 167–192.

[11] F. Martinelli, A. Sinclair, *Mixing time for the Solid-On-Solid model*, Annals of Applied Probability (to appear), arXiv:1008.0125 (2010). A preliminary version appeared in Proc. 41st annual ACM symposium on Theory of computing, pp. 571–580 (2009).

[12] F. Martinelli, F. L. Toninelli, *On the mixing time of the 2D stochastic Ising model with "plus" boundary conditions at low temperature*, Comm. Math. Phys. **296** (2010), 175–213.

[13] Y. Peres, *Mixing for Markov Chains and Spin Systems*, available at www.stat.berkeley.edu/∼peres/ubc.pdf.

[14] D. Randall, P. Tetali, *Analyzing Glauber Dynamics by Comparison of Markov Chains*, J. Math. Phys. **41** (2000), no. 3, 1598–1615.

[15] S. Sheffield, *Random surfaces*, Astérisque No. 304 (2005).

[16] H. Spohn, *Interface motion in models with stochastic dynamics*, J. Statist. Phys. **71** (1993), 1081–1132.

[17] D. B. Wilson, *Mixing times of Lozenge tiling and card shuffling Markov chains*, Ann. Appl. Probab. **14** (2004), 274–325.

Appendix

A. Proof of Proposition 3

Let $\xi \in \Omega_{\eta,U}$ satisfy (6). By monotonicity, if $\eta' \in \Omega$ and $\xi' \in \Omega_{\eta',U}$ are such that $\xi \le \xi'$ and $\eta \le \eta'$, then

$$\mathbb{P}\left(\exists t < L^{10} : \max_{x \in U}\left(\phi^\xi_x(t) - \bar\phi^{\mathbf{n}}_x\right) \ge 2(\log L)^{3/2}\right) \le \mathbb{P}'\left(\exists t < L^{10} : \max_{x \in U}\left(\phi^{\xi'}_x(t) - \bar\phi^{\mathbf{n}}_x\right) \ge 2(\log L)^{3/2}\right) \tag{12}$$

where $\mathbb{P}'$ denotes the evolution with boundary condition η' (instead of η). In particular, this is the case if we set $\eta'_x = \eta_x + \lfloor (3/2)(\log L)^{3/2} \rfloor$ and ξ' is sampled from the measure $\pi^{\eta'}_U(\cdot|A)$, where A is the event $A = A_{\eta',\xi,U} = \{\phi \in \Omega_{\eta',U} : \phi_U \ge \xi_U\}$. From Proposition 1 (applied with $\epsilon = a = 1/2$) one sees that

$$\|\pi^{\eta'}_U(\cdot|A) - \pi^{\eta'}_U\| = O\left(e^{-c(\log L)^{3/2}}\right)$$

for some $c > 0$. This is because if $\Pi^{\mathbf{n}}$ denotes the plane $\{x \in \mathbb{R}^3 :\ x^{(1)}\mathbf{n}^{(1)} + x^{(2)}\mathbf{n}^{(2)} + x^{(3)}\mathbf{n}^{(3)} = 0\}$, and $\Pi^{\mathbf{n}}_h$ denotes the plane obtained shifting $\Pi^{\mathbf{n}}$ upwards by h, then η' is within distance $C \log L$ from $\Pi^{\mathbf{n}}_{(3/2)(\log L)^{3/2}}$, while ξ is within distance $(\log L)^{3/2}$ from $\Pi^{\mathbf{n}}$. Therefore, the probability in (12) is upper bounded by

$$\int \pi^{\eta'}_U(d\xi')\mathbb{P}'\Big(\exists t < L^{10} : \max_{x \in U}\big(\phi^{\xi'}_x(t) - \bar\phi^{\mathbf{n}}_x\big) \ge 2(\log L)^{3/2}\Big) + O\left(e^{-c(\log L)^{3/2}}\right). \tag{13}$$

The initial condition ξ' in (13) is sampled from $\pi_U^{\eta'}$, which is the invariant measure of the dynamics $\mathbb{P}'$, so that the distribution of $\phi^{\xi'}(t)$ coincides with $\pi_U^{\eta'}$ at all later times. Via a union bound over times and recalling the relation between η and η', the probability in (13) is upper bounded by

$$\pi_U^{\eta'}\left(\max_{x\in U}\left(\phi_x-\bar\phi_x^{\mathbf{n}}\right)\geq 2(\log L)^{3/2}\right)\times O(L^{10}\times L^2)$$
$$\leq \pi_U^{\eta}\left(\max_{x\in U}\left(\phi_x-\bar\phi_x^{\mathbf{n}}\right)\geq \frac{1}{2}(\log L)^{3/2}\right)\times O(L^{10}\times L^2)$$

which is of order $\exp(-c(\log L)^{3/2})$, see Proposition 1. The factor $O(L^{10}\times L^2)$ is just the average number of Markov chain moves within time L^{10}, since there are order of L^2 lattice points in U. Similarly one bounds the probability that $\min_{x\in U}(\phi_x^{\xi}(t)-\bar\phi_x^{\mathbf{n}})<-2(\log L)^{3/2}$ for some $t\leq L^{10}$ and the proposition follows. $\square$

B. *Proof of Proposition 5*

We choose the parameter α in (8) as $\alpha=17/2$. We prove the claim by induction on n. For $n=0$ this is trivial since we chose $u_0=2L$, which guarantees that the maximal configuration $\wedge\in\Omega_{\eta,U}$ is below the function $U\ni x\mapsto\psi_{u_0}(x)$.

Assume the claim for some n. For $x\in U$, define the event

$$A_x=\left\{\exists t\in[t_{n+1},L^3]:\phi_x^{\wedge}(t)>\psi_{u_{n+1}}(x)\right\},$$

so that we need to prove

$$\mathbb{P}(\cup_{x\in U}A_x)\leq(n+1)\exp(-c'(\log L)^{3/2}).$$

We have

$$\mathbb{P}(\cup_{x\in U}A_x)$$
$$\leq\sum_{x\in U}\mathbb{P}(A_x;\ \phi^{\wedge}(s)\leq\psi_{u_n}\text{ for every }s\in[t_n,L^3])$$
$$+n\,e^{-c'(\log L)^{3/2}}\tag{14}$$

where we write $\phi^{\wedge}(s)\leq\psi_{u_n}$ to mean that $\phi_y^{\wedge}(s)\leq\psi_{u_n}(y)$ for every $y\in U$.

Given $x\in U$, consider the plane $\widetilde\Pi$ tangent to $\mathcal{C}_{u_n}$ at the point $(x^{(1)},x^{(2)},\psi_{u_n}(x))$ and the plane $\widetilde\Pi'$ obtained by translating downwards $\widetilde\Pi$ by $(\log L)^{3/2}$. The intersection of $\widetilde\Pi'$ with $\mathcal{C}_{u_n}$ is a disk $\widetilde W$ of radius $O(\sqrt{R_n}(\log L)^{3/4})$, whose projection on the horizontal plane we call Z. Let $\widetilde U\subset\mathbb{Z}^2$ be such that $\widetilde U\subset Z$ and $\partial\widetilde U$ is at distance of order 1 from ∂Z, so that of course $diam(\tilde U)=O(\sqrt{R_n}(\log L)^{3/4})$.

Let $\wedge^{(n)}\in\Omega$ be the maximal monotone surface such that $\wedge_x^{(n)}\leq\psi_{u_n}(x)$ for every $x\in U$. Let $\tilde{\mathbb{P}}$ denote the law of the auxiliary monotone surface dynamics in $\widetilde U$, starting at time t_n from $\wedge^{(n)}$ and with boundary conditions $\wedge^{(n)}_{\partial\widetilde U}$. By monotonicity and the definition of $\wedge^{(n)}$ we have

$$\mathbb{P}\left(A_x;\ \phi^{\wedge}(s)\leq\psi_{u_n}\text{ for every }s\in[t_n,L^3]\right)$$
$$\leq\tilde{\mathbb{P}}\left(\exists t\in[t_{n+1},L^3]\text{ such that }\phi_x^{\wedge^{(n)}}(t)\geq\psi_{u_{n+1}}(x)\right).$$

As in the proof of Theorem 1, Propositions 3 and 4 show that after time

$$t_{n+1}-t_n=R_{n+1}(\log L)^{17/2}\geq const\times\text{diam}(\tilde U)^2(\log L)^7$$

the dynamics $\tilde{\mathbb{P}}$ has a variation distance of order $\exp(-c(\log L)^{3/2})$ for some $c>0$ from its equilibrium $\pi_{\tilde U}^{\wedge^{(n)}}$ (recall that $c_-\log L\leq\log R_n\leq c_+\log L$). Proposition 1 gives that

$$\pi_{\tilde U}^{\wedge^{(n)}}[\phi_x\leq\psi_{u_{n+1}}(x)]\geq 1-O\left(e^{-c(\log L)^{3/2}}\right).$$

Indeed, the point $(x,\psi_{u_{n+1}}(x))$ is at distance $(1+o(1))(\log L)^{3/2}$ from the plane $\widetilde\Pi'$ containing the "planar" boundary condition $\wedge^{(n)}_{\partial\tilde U}$.

Putting everything together (plus a union bound over times $t\in[t_{n+1},L^3]$) one gets

$$\mathbb{P}\big(A_x;\ \phi^{\wedge}(s)\leq\psi_{u_n}\text{ for every }s\in[t_n,L^3]\big)$$
$$=O\left(e^{-c(\log L)^{3/2}}\right)\times O(L^3\times L^2)=O\left(e^{-\frac{c}{2}(\log L)^{3/2}}\right).$$

Finally, provided that we choose $c'=c/2$, from (14) we get

$$\mathbb{P}(\cup_{x\in U}A_x)\leq(n+1)e^{-c'(\log L)^{3/2}}$$

(the union bound over $x\in U$ gives just an extra $O(L^2)$) which is the desired claim.

C. *Proof of Proposition 8*

As in the proof of Proposition 5 we proceed by induction in $n\leq M$. The initial step $n=0$ is obvious because the maximal configuration $\wedge$ is inside $\mathcal{C}_{u_0}$. Thus, let us assume the statement true for $n<M$ and let us prove it for $n+1$.

For $1\leq i\leq L$ define the event

$$A_i=\{\exists\,t\in[t_{n+1},L^3]:\phi_i^{\wedge}(t)>\psi_{u_{n+1}}(i)\}.$$

Using the inductive assumption we may write

$$\mathbb{P}\big(\cup_{i=1}^L A_i\big)$$
$$\leq\sum_{i=1}^{L}\mathbb{P}\big(A_i;\ \phi^{\wedge}(s)\in\mathcal{C}_{u_n}\forall s\in[t_n,L^3]\big)+n\,e^{-(\log L)^{3/2}}.$$

Fix $i=1,\dots,L$ and consider the line $\mathbb{L}$ tangent to $\mathcal{C}_{u_n}$ at the point $(i,\psi_{u_n}(i))$ and the line $\mathbb{L}'$ obtained by translating downwards (i.e. in the $-y$ direction) $\mathbb{L}$ by

$$2(u_n-u_{n+1})=2(\rho_L^2/u_n)^{1/3}(\log L)^2.$$

Let us denote by x_-,x_+ the horizontal coordinates of the leftmost and rightmost points of $\mathbb{L}'\cap\mathcal{C}_{u_n}$ and let I be the set of integers in $[x_-, x_+]$. Clearly $|I| = O((\rho_L^2/u_n)^{2/3} \log L)$.

Consider now the SOS dynamics in the interval I with boundary conditions equal to $\lfloor \psi_{u_n}(i_\pm) \rfloor$ at the left and right boundary $i_\pm$ of I and floor at zero height. This auxiliary evolution starts at time t_n from the maximal configuration $\wedge^{(n)}$ in the set of $\phi \in [0, L]^I$ such that $\phi_i \le \psi_{u_n}(i)$ for any $i \in I$. We denote by $\mathbb{P}'$ the law of this auxiliary chain. Observe that $\wedge^{(n)}$ is within distance $2(u_n - u_{n+1}) = O(\sqrt{|I|}(\log |I|)^{3/2})$ from the line $\mathbb{L}'$ so that Proposition 7 will be applicable with $\kappa = 3/2$. By monotonicity we have

$$\begin{aligned}&\mathbb{P}\left(A_i;\ \phi^\wedge(s) \in \mathcal{C}_{u_n} \forall s \in [t_n, L^3]\right)\\ \le\ &\mathbb{P}'\left(\exists t \in [t_{n+1}, L^3] \text{ such that } \phi_i^{\wedge^{(n)}}(t) \ge \psi_{u_{n+1}}(i)\right).\end{aligned}$$

Because of Proposition 7, for some $\alpha_1 > 0$, after time $|I|^2(\log |I|)^{\alpha_1} \le t_{n+1} - t_n$ (taking e.g. $\alpha' = 3 + \alpha_1$ in (10)) the dynamics $\mathbb{P}'$ has a variation distance of order $\exp(-c(\log L)^2)$ from its equilibrium which we denote by $\pi_I^{(n)}$. Since the distance between the point $(i, \psi_{u_{n+1}}(i))$ and the line $\mathbb{L}'$ is at least $c\sqrt{|I|}(\log |I|)^{3/2}$ for some $c > 0$, Lemma 1 gives that

$$\pi_I^{(n)}\left(\phi_i \le \psi_{u_{n+1}}(i)\right) \ge 1 - O\left(e^{-c(\log |I|)^3}\right) \ge 1 - e^{-c'(\log L)^3}.$$

As in the proof of Proposition 5, simple union bounds over $i \in [1, L]$ and $t \in [t_{n+1}, L^3]$ give

$$\begin{aligned}&\sum_{i=1}^{L} \mathbb{P}(A_i;\ \phi^\wedge(s) \in \mathcal{C}_{u_n} \forall s \in [t_n, L^3])\\ &\quad \le L^5 e^{-c''(\log L)^2} \le e^{-(\log L)^{3/2}}\end{aligned}$$

which finishes the proof of the inductive step. $\square$